

Data centre waste heat:  
applications, societies, metrics



**Petter Terenius**

**This dissertation is submitted for the degree of**

**Doctor of Philosophy**

**June 2023**

**School of Computing and Communications**

*I dedicate this thesis to my mother, an agricultural scientist who helped food security endeavours in the Global South through her travels, lab establishments and dissemination of the (now) low-tech lab practices she set up many years prior.*

*As she explains: what good does a fancy piece of lab equipment do, if there is no availability of spare parts, let alone reliable power?*

*Her work is a reminder that successful implementation of technology must be situated in a geographical and societal context, that is, where the material meets the social.*

*And more than that, she's my loving mum and a close friend.*

In principle, most of us would accept that to see something, say x, as a system means to see x as a whole (or a system-whole) – not as a mere aggregate of parts; that is, a systems approach is about seeing things globally. Hence, a systems approach to world problems is a way of studying and treating world problems globally. But how else could world problems be tackled? The notions of a “systems approach” and that of “world problems” seem to belong to each other.

Ramsés Fuenmayor,

“The inseparability of systems thinking and world issues in the modern epoch”

Here again we are reminded that in nature nothing exists alone.

Rachel Carson,

*Silent Spring*

If it’s interesting, it is worth writing about.

Damian Borowiec

## **Declaration**

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words, including footnotes, but excluding the bibliography, appendices and the front matter. A rough estimate of the word count is 69,000.

Petter Terenius

June 2023

## **Abstract**

In the near future, a few percent of world electricity may be needed to power data centres around the world. This energy ultimately becomes waste heat, and the thesis investigates ways to use it. But for selected uses in cold regions, previous research has not addressed this issue. Neither have the growing data needs of low-income countries been discussed much from an environmental perspective.

The thesis argues there exists a bond between technology, societal progress and environmental sustainability, and that this bond can be used to solve the energy problems of the rapidly growing data centre industry. In fact, a society in need of data exchange and a planet unable to cope with unsustainable energy use turn out to be good bed-fellows, as an evidently holistic problem calls for an equally holistic, systems science-based solution.

Through three cases studies (Malaysia, Costa Rica, Sweden), research is carried out relating to dehydration of commodities such as coffee beans, wooden pellets and seaweed, as well as to energy storage solutions. The concepts are then evaluated using a developed analytical framework and novel data centre energy efficiency metrics. The work is underpinned by a literature review, interviews and ethnographic studies. Crucial to the evaluation has been the possibility to compare the three contrasting cases, where the Arctic meets the tropics and where city meets countryside.

The results show that a systems science-based view and a high-level metric open up new possibilities for data centre waste heat use worldwide.

## **Thank you,**

Richard, for your supervision and for fascinating discussions, but even more for believing in my ability to sculpt and execute this multi-faceted thesis;

Peter, for your supervision and for sharing your technical expertise, but even more for the way in which you led the Experimental Distributed Systems Laboratory through my time in Lancaster and through the covid closure;

my fellow members of the B38 squad – Gingfung, Dominic, Damian, James, Daria, Will, Matt, Stefan, Yao, Marta, Lewis, Weijia, John – for several years of discussions and for a simply enjoyable time;

my fellow members of the MSF – especially Colin, Pip, Maria, Agnieszka, Tommy, Beth and Jade – for your insights and for your friendship;

the Leverhulme Foundation and the managing team of the Material Social Futures program – Rob, Richard, Louise and others – for making this opportunity come true and for your sincere interest in a sustainable world;

Chengcheng, for our discussions, but so much more for being my housemate and my supportive friend;

my other friends at the S-CC: Ludwig, Gwen, Christina, Kelly, Miriam, James, Kathy, Lucy, Adrian, Darren, Mark, Claire-Anne, Nige and others, for being supportive and kind;

Dr Sathiabama T. Thirugnana, director of UTM-OTEC, for your help, for your time, and for your ever-so-kind spirit;

Hariharan Jai Karthikeyan (Hari), my fellow researcher at UTM-OTEC, for your help, smart thinking, smile and kindness;

Dr Lars G. Golmen from the Norwegian Institute for Water Research for OTEC discussions;

all other people that are somehow connected to this thesis: Shukur, Prof. Dato Bakar Jaafar, Assoc. Prof. Syuhaida Ismail and others at Universiti Teknologi Malaysia for expertise and for help arranging meetings in peninsular Malaysia, Prof. John Summers and Assoc. Prof. Mattias Vesterlund at Research Institutes of Sweden for research insights, Dr Andy Lawrence and Max Smolaks at the Uptime Institute for perspectives on metrics, Dr Sönnich Dahl Sönnichsen at Copenhagen Business School for our discussion on circular economy, the International Coffee Organization for sharing working data, entrepreneur Caroline Prah for discussions on fodder dehydration and other matters, Dr Drew Lilley at UC Berkeley for pointing to salt hydrates, staff at EcoDataCenter for showing the premises, CEO Ulf Lindvall at Lindvalls Kaffe for sharing experiences from the international coffee industry, Dr Tangku Mohd Azzman Shariffadeen for sharing the story of Malaysia's transformation, Operations Manager Daniel Widman at Falu Energi & Vatten AB for presenting pellet dehydration work, Maja Kehic and Christophe Lillo of DeepSquare and Samuel Perren of Industrielle Werke Basel for presenting the liquid immersion cooled data centre in Basel, Dr Ly Lindman from the Swedish University of Agricultural Sciences for introducing me to ecology methods and hosting our forestry fieldwork; Steve Naumann, Senior Advisor at the Office of Government-wide Policy, U.S. General Services Administration, for discussions on data centre waste heat use implementation in the USA; Prof. Phang and Dr Yeong at Universiti Malaya for discussions on seaweed farming, Claude Egli from

optimiQ for networking assistance in Switzerland, WWF Malaysia and the Boston Consulting Group for agreeing to our meeting regarding report findings, Jon Laban, Jaime Comella Gómez-Aller, Cosimo Pecchioli and others at OCP Foundation for introducing me to open standards, the OCP data centre energy efficiency metrics working group for past and future discussions;

my family – Malin, Lars, Olle – for your support as well as for discussions on methodology and sustainability over the years;

my extended family and friends in Malexander – particularly Mats Åkerblom, Björn Åkerblom, Karin Åkerblom and Eva Kaijser – for much of the same thing;

Henrik, for inspiring discussions on, well... everything relating to scientific endeavours;

Ira, for taking me to the cradle of Western thought;

Sofia, Ross, Kristina, Calle, Camilla, Matouf, Saxi and the NHS, for keeping me going;

Astrid, for brightening up my day, and for being a hundred percent you.



# Preface

Much has been written about data centre energy efficiency from a hardware perspective. This thesis addresses the same topic, but at a conceptually higher level, viewing a data centre as a part of the societal web. The reader should therefore be prepared to meet quite diverse topics in this text – ranging from energy metrics to the tsunami risk of the Indonesian coastline, from the writings of Dilthey to nitrogen-filled data centres on the ocean floor, from governmental environmental legislation to the drying speed of *Elodea canadensis*.

Though data centre energetics may seem straightforward, this thesis will treat such matters as a riddle. Unpacking, examining and ultimately solving this riddle necessitates new perspectives and broader strokes than those previously painted.

I wish the reader an interesting journey.

Lancaster – Uppsala – Malexander June 2023

Petter Terenius

# Contents

|  |             |
|--|-------------|
| <b>Preface .....</b>   | <b>viii</b> |
| <b>Contents .....</b>  | <b>ix</b>   |
| <b>List of figures.....</b>  | <b>xvi</b>  |
| <b>List of tables .....</b>  | <b>xix</b>  |
| <b>List of equations.....</b>  | <b>xxix</b> |
| <b>Refereed publications .....</b>                                       | <b>xxix</b> |
| <b>Glossary .....</b>  | <b>xxix</b> |
| <b>1 Introduction .....</b>  | <b>1</b>    |
| 1.1 Chapter outlines and contributions.....                              | 6           |
| 1.2 Assumptions .....  | 9           |
| 1.2.1 The warm airflow from data centre exhausts can be repurposed ..... | 9           |

|   |           |
|---|-----------|
| 1.2.2 The citizens of the local community have access to smartphones and/or computers (desktops, laptops, tablets)..... | 9         |
| 1.2.3 The society has sufficient access to electricity to run the data centre.....                                      | 10        |
| 1.2.4 The society has a need for dehydration, energy storage or energy production as envisaged in this thesis.....      | 10        |
| 1.3 Purpose.....  | 10        |
| 1.4 Research questions.....   | 11        |
| 1.5 Limitations.....  | 12        |
| <b>2 Related work.....</b>  | <b>13</b> |
| 2.1 Data centre research – an overview.....   | 13        |
| 2.2 The anthropology view.....  | 16        |
| 2.3 The duality of the data centre cosmos.....  | 22        |
| 2.4 Chapter conclusion.....   | 25        |
| <b>3 Methodological perspective.....</b>  | <b>27</b> |
| 3.1 Methodological approaches of this thesis.....   | 28        |
| 3.2 The material social perspective.....  | 29        |
| 3.3 Systems science and systems engineering.....  | 34        |
| 3.4 Mixed methods.....  | 37        |
| 3.4.1 Mixed methods in cross-disciplinary sustainability projects.....  | 39        |
| 3.4.2 Apparent weaknesses of mixed methods.....   | 41        |
| 3.4.3 Team collaboration versus the lone researcher.....  | 41        |
| 3.4.4 Lack of unified language.....   | 41        |
| 3.4.5 Mixed methods in the thesis.....  | 44        |

|   |           |
|---|-----------|
| 3.4.6 All connected .....                                       | 48        |
| 3.5 Contributions and associated methods .....                  | 49        |
| 3.5.1 Literature review .....                                   | 50        |
| 3.5.2 Investigating uses for data centre waste heat.....        | 54        |
| 3.5.3 Cultural factors .....                                    | 57        |
| 3.5.4 Metrics .....   | 58        |
| 3.6 Chapter conclusion.....                                     | 59        |
| <b>4 Data centre energetics .....</b>                           | <b>63</b> |
| 4.1 Defining clean energy .....                                 | 64        |
| 4.2 Use and reuse within the circular economy .....             | 70        |
| 4.3 Data centre energy use and waste heat.....                  | 75        |
| 4.4 A dive into liquid immersion cooling .....                  | 77        |
| 4.5 The data centre as an energy grid regulating force .....    | 81        |
| 4.6 Metrics, used and abused .....                              | 84        |
| 4.7 Chapter conclusion.....                                     | 86        |
| <b>5 A framework for data centre placement .....</b>            | <b>89</b> |
| 5.1 Energy quality and use .....                                | 95        |
| 5.1.1 Local definitions of clean energy.....                    | 96        |
| 5.1.2 Alternatives for energy use.....                          | 97        |
| 5.2 Preferable data centre configuration .....                  | 98        |
| 5.2.1 Traditional, visionary and specialised data centres ..... | 99        |
| 5.2.2 Size preference.....                                      | 102       |
| 5.3 Geographical prerequisites.....                             | 105       |
| 5.3.1 Physical.....   | 110       |

|   |            |
|---|------------|
| 5.3.2 Infrastructure access .....                           | 113        |
| 5.4 Deployment and operation .....                          | 116        |
| 5.4.1 Deployment.....                                       | 116        |
| 5.4.2 Operation .....                                       | 118        |
| 5.5 Waste heat use capabilities.....                        | 121        |
| 5.5.1 Reuse feasibility at specific site.....               | 122        |
| 5.5.2 Metrics for waste heat use .....                      | 123        |
| 5.6 Society profile .....                                   | 124        |
| 5.6.1 Society structure .....                               | 124        |
| 5.6.2 Population .....                                      | 125        |
| 5.6.3 Culture .....   | 131        |
| 5.7 Chapter conclusion.....                                 | 135        |
| <b>6 Applications for data centre waste heat .....</b>      | <b>142</b> |
| 6.1 Heating buildings, greenhouses and swimming pools ..... | 143        |
| 6.2 Commodity dehydration.....                              | 144        |
| 6.2.1 Large-scale commodity dehydration .....               | 147        |
| 6.2.2 Small-scale commodity dehydration .....               | 149        |
| 6.3 Energy storage solutions .....                          | 150        |
| 6.3.1 Seaweed for biofuel production.....                   | 151        |
| 6.3.2 Salt hydrate charging .....                           | 154        |
| 6.4 Industrial symbiosis with OTEC plants .....             | 155        |
| 6.5 Challenges and chapter conclusion .....                 | 158        |
| <b>7 Data centres and society .....</b>                     | <b>161</b> |
| 7.1 The three actors of a society.....                      | 161        |

|   |            |
|---|------------|
| 7.2 Site selection .....  | 164        |
| 7.2.1 Rationale for the selection of countries .....                | 165        |
| 7.2.2 Introducing ICRI.....   | 167        |
| 7.3 Building a sustainable future for Malaysia .....                | 169        |
| 7.4 Drying coffee beans in Costa Rica.....                          | 174        |
| 7.5 Sweden and the new Klondike.....                                | 179        |
| 7.6 Three research approaches to society.....                       | 185        |
| 7.7 Chapter conclusion.....   | 186        |
| <br>  |            |
| <b>8 Data centre energy efficiency metrics .....</b>                | <b>190</b> |
| 8.1 Metrics, used or proposed .....                                 | 191        |
| 8.1.1 PUE.....  | 191        |
| 8.1.2 The quest for new data centre energy efficiency metrics ..... | 194        |
| 8.1.3 Some proposed metrics for data centre work.....               | 196        |
| 8.1.4 CPE, Compute Power Efficiency.....                            | 197        |
| 8.1.5 DPPE, Datacenter Performance Per Energy .....                 | 197        |
| 8.1.6 DCeP, Data Center Energy Productivity .....                   | 198        |
| 8.1.7 FLOPS and other server energy measurement metrics.....        | 198        |
| 8.1.8 Other IT equipment energy metrics.....                        | 199        |
| 8.1.9 Contextualising metrics .....                                 | 200        |
| 8.2 Introducing the Datacenter Energy Sustainability Score .....    | 202        |
| 8.3 PUE, ERE, ERF and DESS – a comparison .....                     | 206        |
| 8.4 The EU Energy Efficiency Directive .....                        | 209        |
| 8.5 Chapter conclusion.....   | 212        |

|  |            |
|--|------------|
| <b>9 Analysis</b> .....  | <b>215</b> |
| 9.1 Testing the analytical framework as a checklist for identifying potential data centre sites .....  | 217        |
| 9.1.1 Energy quality and use.....  | 217        |
| 9.1.2 Preferable data centre configuration .....   | 218        |
| 9.1.3 Geographies .....  | 219        |
| 9.1.4 Deployment, operation.....   | 221        |
| 9.1.5 Heat use capabilities .....  | 222        |
| 9.1.6 Society profile.....   | 224        |
| 9.1.7 Commentary.....  | 224        |
| 9.2 PUE and DESS.....  | 225        |
| 9.3 Reflection on the two analysis methods.....  | 230        |
| 9.3.1 The analytical framework .....   | 230        |
| 9.3.2 PUE and DESS .....   | 231        |
| 9.4 Chapter conclusion.....  | 231        |
| <b>10 Discussion</b> .....   | <b>233</b> |
| 10.1 This thesis in retrospect.....  | 236        |
| 10.2 Reflections on the research questions and purpose.....  | 239        |
| 10.2.1 RQ 1: How can data centre waste heat be used in different geographical and societal contexts?.....  | 240        |
| 10.2.2 RQ 2: What are the societal implications and environmental benefits of implementing the ideas proposed? .....                               | 240        |
| 10.2.3 RQ 3: Which is the optimal design for a sustainable data centre?.....   | 240        |
| 10.2.4 RQ 4: Can there actually be such a thing as a one key metric for data centre energy efficiency, and if so, how should it be designed? ..... | 241        |

|  |            |
|--|------------|
| 10.2.5 RQ 5: To what extent is material social studies suitable for<br>investigating cross-disciplinary problems?..... | 243        |
| 10.3 Further research.....   | 244        |
| 10.4 Final remarks.....  | 247        |
| <b>Post scriptum: a long and winding road?.....</b>  | <b>251</b> |
| <b>References.....</b>   | <b>253</b> |
| <b>Appendices .....</b>  | <b>272</b> |
| <b>Appendix A: Engagement .....</b>  | <b>273</b> |
| <b>Appendix B: Research visits and fieldwork .....</b>   | <b>275</b> |
| <b>Appendix C: Interview questionnaire for Dr Azzman.....</b>  | <b>276</b> |



# List of figures

*All screenshots, diagrams, graphs and photos by the author, unless otherwise specified.*

|   |    |
|---|----|
| Figure 1. Google search on “datacenter”.....  | 20 |
| Figure 2. The “cloud”.....  | 25 |
| Figure 3. A methodological view of the thesis.....  | 29 |
| Figure 4. The doughnut model and the Rings of Sustainability model. ....                            | 33 |
| Figure 5. The methodological foundation of the thesis. ....   | 35 |
| Figure 6. The thesis’s relation to engineering, sustainability science and the social sciences..... | 45 |
| Figure 7. The two tracks of the thesis.....   | 46 |
| Figure 8. The analytical framework of the thesis. ....  | 51 |
| Figure 9. A systematic view of the circular economy. ....   | 72 |
| Figure 10. Energy flows in Europe on 27 October 2022.....   | 88 |

|   |     |
|---|-----|
| Figure 11. The analytical framework of the thesis (reintroduced). .....             | 94  |
| Figure 12. Crypto mining. ....  | 104 |
| Figure 13. The thousand largest seismic events (earthquakes) since 1970. ....       | 112 |
| Figure 14. Terrestrial and subsea Internet cables, and groups of data centres. .... | 114 |
| Figure 15. Population growth in sub-Saharan Africa. ....                            | 127 |
| Figure 16. Map of Africa with subsea Internet cables and groups of data centres.... | 130 |
| Figure 17. A few facets of culture in Malaysia. ....                                | 132 |
| Figure 18. A Google Trends search for “datacenter” as a topic. ....                 | 135 |
| Figure 19. Means and goals to enable sustainable societies.....                     | 147 |
| Figure 20. Dried fish sold in a supermarket in Kuala Lumpur.....                    | 148 |
| Figure 21. Large-scale commodity dehydration.....                                   | 149 |
| Figure 22. Small-scale coffee beans dehydration.....                                | 150 |
| Figure 23. The dissipation of moisture in three seaweed batches. ....               | 153 |
| Figure 24. The working principle of CC-OTEC.....                                    | 156 |
| Figure 25. A tropical data centre, OTEC and dehydration plants. ....                | 157 |
| Figure 26. Income for the fifteen largest coffee exporters. ....                    | 167 |
| Figure 27. The Kuala Lumpur skyline.....  | 170 |
| Figure 28. Cyberjaya’s centre.....  | 172 |

|   |     |
|---|-----|
| Figure 29. Earthquakes zoomed in on Malaysia.....   | 174 |
| Figure 30. Coffee farms, data centres, dehydration facilities, traders/retailers and ICT actors (reintroduced)..... | 177 |
| Figure 31. Maps showing dimensions of concern for data centre placement.....  | 179 |
| Figure 32. Erecting Facebook’s fourth data centre building in Luleå. ....   | 182 |
| Figure 33. A promotional catalogue for data centres in Luleå and surroundings.....                                  | 183 |
| Figure 34. Human Development Index. ....  | 188 |
| Figure 35. PUE issues.....  | 194 |
| Figure 36. The three main concerns of energy use relating to a data centre. ....                                    | 203 |
| Figure 37. Focus areas of some of the metrics pertaining to data centre energy.....                                 | 207 |
| Figure 38. More than 800 million people live under the poverty line.....  | 217 |
| Figure 39. Thesis writing and road and railroad infrastructure in Kuala Lumpur.....                                 | 222 |
| Figure 40. Comparison cases for PUE and DESS. ....  | 229 |
| Figure 41. Components of a data centre. ....  | 234 |
| Figure 42. A “home for the cloud”. ....   | 239 |
| Figure 43. Presenting DESS at the OCP European Summit, Prague 2023. ....  | 242 |

# List of tables

|   |     |
|---|-----|
| Table 1. Literature theme distribution among chapters. ....                             | 53  |
| Table 2. Some of the themes and cases. ....   | 137 |
| Table 3. Duration, time of day, temperature, humidity and batch weight. ....            | 152 |
| Table 4. Actors with concerns relating to a local data centre placed in a society. .... | 164 |
| Table 5. Comparison of some key metrics for the three countries. ....                   | 165 |
| Table 6. Country infrastructure and consumer readiness. ....                            | 169 |
| Table 7. PUE example. ....  | 192 |
| Table 8. Constituents of DESS. ....   | 206 |
| Table 9. The annual PUE and DESS cases mentioned in the thesis. ....                    | 225 |

# List of equations

|  |     |
|--|-----|
| Equation 1: ICRI, Infrastructure/Consumer Readiness Index .....          | 168 |
| Equation 2: Estimated installed data centre capacity in Costa Rica ..... | 178 |
| Equation 3: Estimated addition from container-sized data centres .....   | 178 |
| Equation 4: PUE, Power Usage Effectiveness .....                         | 191 |
| Equation 5: CPE, Compute Power Efficiency .....                          | 197 |
| Equation 6: DPPE, Datacenter Performance Per Energy .....                | 197 |
| Equation 7: DCeP, Data Center Energy Productivity .....                  | 198 |
| Equation 8: Energy Reuse Effectiveness .....                             | 200 |
| Equation 9: Compute Carbon Efficiency .....                              | 201 |
| Equation 10: Datacenter Energy Sustainability Score .....                | 204 |

# Refereed publications

P. Terenius, P. Garraghan, and R. Harper, “Using data centre waste heat to dry coffee whilst supplying small-scale farmers with ICT,” in International Conference for Sustainable Development, New York, The Earth Institute, 2020: Columbia University. [Revised and published in the *Journal of Strategic Innovation and Sustainability*; see below.]

P. Terenius, L. G. Golmen, P. Garraghan, and R. Harper, “Heat energy from datacenters: an opportunity for marine energy,” International Conference on Ocean Energy 2021, Washington, DC, 2021.

P. Terenius, P. Garraghan, and R. Harper, “Using data centre waste heat to dry coffee whilst supplying small-scale farmers with ICT: A novel idea and a case study based on a systems approach,” *Journal of Strategic Innovation and Sustainability*, vol. 16, no. 2, 2021.

P. Terenius, P. Garraghan, and R. Harper, “Novel Strategies for Data Centre Waste Heat Use,” in Towards a Cleaner Earth: 18<sup>th</sup> International Conference on Clean Energy, Kuching, 2022. (Accepted/In press) In: AIP Conference Proceedings. 6 p.

P. Terenius, P. Garraghan, and R. Harper, “A material social view on data center waste heat: Novel uses and metrics,” *Frontiers in Sustainability*, vol. 3, 2023, doi: <https://doi.org/10.3389/frsus.2022.1008583>.

# Glossary

**actor** – in systems science and systems engineering terminology, an “actor” is a representation for a group of users, stakeholders or similar. In this thesis, actors discussed are mainly presumed to be industrial actors (say, members of Big Tech). However, the term is also used when referring to three categories involved with, and affected by, data centre placement: Industry, Consumer and the Authorities.

**Big Tech** – here viewed as Google (the parent company is named Alphabet, which also includes YouTube and Android), Meta (parent company for Facebook, WhatsApp, Messenger and more brands), Microsoft, Amazon (including AWS) and Apple. For convenience and readability, the thesis uses “Google” and “Meta”. Meta’s offerings are relatively diverse, so “Facebook” would be a wrongful and misleading metonym. With YouTube and Android as exceptions, Google, on the other hand, uses its brand name more consistently (such as “Google Search” and “Google Docs”).

**CapEx and OpEx** – Capital expenditure (CapEx) is money paid to obtain a product of some sort. Operating expenditure (OpEx) is the money paid to keep it running. For a data centre, CapEx is the cost of the building material, the erection of the building,



cabling infrastructure, cooling infrastructure etc, and optionally the servers within it. OpEx is the cost of running the data centre, that is, salaries, electricity bills and so on.

**circular economy** – suggested by Boulding as the optimum model for a sustainable society, the circular economy focuses extended use of materials, mainly through the three Rs: Reduce, Reuse, Recycle. A true circular economy, where materials are recycled ad infinitum, is unattainable. This is because materials degrade over time; there are only so many times a paper cup can be recycled since its structure needs a certain amount of virgin paper. That being said, the circular economy is an ideal many strive for, to get away from consumerism’s *linear* economy, the antithesis of the circular.

**co-location data centre** – a co-location data centre rents out space to clients’ servers. Since co-location data centres are not tied to a specific customer, they compete over prospective clients’ businesses. Typical competitive factors are uptime, price and stability, but today, sustainability is becoming a factor as well, especially for clients bound by CSR policies.

**CSR (corporate social responsibility)** – the ambition for an organisation to do more than what is required by law. Many, if not most, international enterprises today promote their CSR endeavours, and to live up to them, they may be using a sustainability framework such as the UN SDGs. Environmental, Social and Governance (ESG) concerns, pointing back to the environmental, social and financial sustainability goals promoted by the United Nations in 1987 (see chapter 1), are linked to CSR.

**data centre** – a large collection of servers that store and compute the data we use every day. “Hyper-scale” data centres have servers in the thousands, and are typically owned and run by companies such as Google, Amazon, Meta or Alibaba.

**edge data centre** – a small data centre used to compute and store locally used data. The benefit of edge data centres is that data do not require transmission to faraway data centres. This minimises latency and puts less strain on transmission networks.

**embodied [energy, GHG emissions or similar]** – energy or other matter of concern used or accumulated during production, distribution and installation for the product at hand. One example is the embodied energy of a server. It can be compared to the capital expenditure (CapEx) of the server.

**energetics** – (the study of) concerns in the academic discourse associated with energy.

**entropy** – in physics, entropy, a term connected to the second law of thermodynamics, is seen as the degree of disorder. Throughout the lifetime of the universe, the disorder will increase for spontaneous processes; stars and planets are but temporary agglomerates of matter which will eventually disintegrate into ultimate equilibrium. This scenario implies that ordering of matter requires added energy. In the context of this thesis, the term is used also in a derived sense: as the degree of energy previously available for use, but now unavailable to perform work.

**fossil gas** versus **natural gas** – “natural gas” is strictly a correct term since methane and similar gases appear in nature. On the other hand, adding “natural” conceals what this mixture of gaseous hydrocarbons (as opposed to, for example, the gaseous mixture we call “air”) really is: fossil fuel. Hence, the term fossil gas is used in this thesis.

**GDP and GDP (PPP)** – Gross Domestic Product (GDP) is the monetary value of all products and services produced within a country over a year. Power purchase parity (PPP) is used to reflect the differences in monetary value between markets: a euro in one country may buy less than it would have done in another. To obtain an

understanding of the comparable wealth of citizens of certain countries, it often makes sense to study the GDP (PPP) *per capita*.

**geographies** – this term is used for concerns associated with (physical) geography, for example, how people’s everyday lives are affected by topography or distance to sea.

**GHG (greenhouse gases)** – the greenhouse gases discussed most today are mainly carbon-based. Carbon dioxide (CO<sub>2</sub>) and methane (CH<sub>4</sub>) – which is the main constituent of fossil (or “natural”) gas and a much more potent GHG than carbon dioxide – are the most problematic of these. Other GHGs include nitrous oxide (N<sub>2</sub>O) and ozone (O<sub>3</sub>). It should be said that also water vapour is a greenhouse gas, needed to heat Earth’s climate to the correct temperature levels.

**GHG emissions** – greenhouse gas emissions except for water vapour. Here, this term has been chosen over “carbon emissions”, since not all GHG emissions include carbon. The term “GHG footprint” is used rather than “carbon footprint”, for the same reason.

**hyper-scale and large-scale data centres** – these are warehouses with big collections of servers. There are no physically bound definitions of the sizes of these buildings. Consequently, the scales may well differ with time, especially as data centres in general tend to be built bigger with time. In this thesis, following current trends, a hyper-scale data centre uses 20 MW or more when run at full power.

**ICT (Information and Communications Technology)** – ICT involves both data and (mobile) telephony, that is, end-users’ computers and phones, data transmission networks, data centres and similar technologies and infrastructure.

**Internet backbone** – the main data routes, or highways, of the Internet, and its supportive infrastructure. For long-distance data transmission, it is important to reach

the Internet backbone quickly. For example, data from and to an end-user in Poland connecting to a website hosted in the USA will rapidly be transmitted through the Internet backbone.

**Internet exchange point (IXP)** – a “station” at which local ISPs (Internet Service Providers) connect to the wider Internet. To maximise speed and stability, the distance from the ISP to the IXP should ideally be short.

**Internet Service Provider (ISP)** – an organisation that (above all) provides access to the Internet. An ISP thereby serves as a gateway to the wider Internet. ISPs are connected to data centres in different ways depending on location, ownership and similar variables.

**IT (Information Technology)** – in the vernacular, “IT” has typically excluded telephony (the “C” in ICT), but with today’s smartphones, those lines have become blurred. As a consequence, the thesis makes little distinction between the two terms.

**latency** – here viewed as the time it takes for the roundtrip between a client and a host, or put differently, between the end-user’s request and the data centre’s response. In the following, it is assumed that acceptable latency should be lower than a second – preferably much, much lower, say, a few milliseconds for a traffic system.

**material social studies** – a cross-disciplinary academic field proposed in the thesis, where material and social/societal enquiries are given equal weight, and together answer complex questions with higher reliability than a uni-disciplinary academic field can do in isolation.

**net zero** – indicates that an entity (often a business segment, a specific business or a branch) has no net GHG emissions. Just about all human activity comes with GHG

emissions, so to achieve net zero, one must both limit greenhouse gas emissions and somehow lower the GHGs in the atmosphere (say, through reforestation).

**on-premise** – hosting a small data centre within an organisation’s walls. This is often unfavourable today, due to added costs and maintenance risks (see section 5.2.2).

**rack density** – how much power (in kW) a server rack can use. In other words, the rack density partially sets the maximum limit for server performance (of course, how efficient the server does computation is another factor).

**rebound effects** – a phenomenon where improvements on one end increases demand, thus obliterating the initial benefits. For example, improvements in car batteries does not mean electric vehicles are much lighter or less expensive than those five years ago; instead, their driving range expands.

**SDG** – see the United Nations’ Sustainable Development Goals (SDGs).

**server** – a server is a computer used to provide a client computer with data and computation. A data centre is filled with servers. A server can be configured for specific tasks, but apart from performance capabilities, there is in principle little that differs a server from a regular PC.

**service level agreement (SLA)** – an agreement of minimum delivered service. For example, a data centre can have an SLA with one of its clients assuring 99.99% access to services.

**strong sustainability** – the idea that humanity is a subset to nature. This idea implies that whatever activities humanity undertakes, nature’s limits cannot be crossed. In

contrast, “weak sustainability” promotes more of a human-centric attitude to the world and its resources.

**system** – a set of related parts and the relations themselves. Connected to a system are external entities. In systems science, a system has at least one input, internal process(es) and at least one output. A system can typically be divided into sub-systems. For example, a city can be divided into many households, themselves forming systems.

the **United Nations’ Sustainable Development Goals (SDGs)** – sixteen goals involving societal and environmental challenges (poverty, gender equality, infrastructure, climate action and so forth) and one overarching goal to oversee and facilitate the implementation of the others.

**uptime** – a measurement of stability of a data centre, typically measured over the course of a year. A data centre is of little use unless it can be reached by its clients. As a consequence, uptime is a key metric, and often described in the “number of nines”, say, 99.999% uptime. “Availability” is a synonym to uptime, though the term “uptime” signals that the conversation involves data centres specifically. As the name suggests, the Uptime Institute is an authority on data such concerns.

**waste heat use** versus **waste heat reuse** – the industry term is heat reuse. In the thesis, “waste heat use” and “energy reuse” are preferred. The reason why is that waste heat is a delivery of a data centre; it has not been used in its current form before. Hence, as a by-product it cannot be “reused” the first time around, only “used”.

# 1 Introduction

Until the middle of the nineteenth century, the world was imagined more or less as a cornucopia to be used to humanity's desire. Sugar, silk and other commodities could be shipped all over the world, and allowed for rapid economic growth. Innovations shaped the latest centuries, and with steam engines and the harnessing of the somewhat magical electrical energy, a vision emerged of a society vastly different from what had previously been thinkable. Mary Shelley's *Frankenstein* and Jules Verne's body of works illustrate how these scientific ideas overflowed and found their ways into the humanities.

With the twentieth century came the realisation that treating the world as an ever-filled larder came at a cost: the London observatory moved out from the city in 1948 as the smog made observations difficult (Macdonald 2020). On a larger scale, Rachel Carson's *Silent Spring* (1962) raised the general public's awareness about the ecological consequences the world was about to face, and what, in fact, had already started to happen. *The Limits to Growth* (1972) acknowledged Carson's findings and efficiently illustrated the fact that the world's natural resources are very much finite.

1987 saw the advent of the Brundtland Commission's *Our Common Future*, whereupon not only academics and governments but also corporations took initiatives for a cleaner and more sustainable future. A climate debate followed, where financial or political interests often weighed heavier than environmental concerns – George W Bush's refusal to ratify the Kyoto Protocol 2001, with the unconvincing motivation that China and India were exempted from some of its requirements (Bush 2001), comes to mind. More recent examples are the EUs questionable decision in June 2022 to include fossil (natural) gas among sustainable energy resources (Hancock 2022), and Germany's decision to – despite massive protests – start a new coal mine in January 2023, “*bulldozing*” a town in the process (Mansoor 2023).

Thus, we have a history of warnings about climate change, as well as a history of ignoring those warnings – in fact, the greenhouse effect was proven more than 120 years ago (Arrhenius 1896). It is therefore not so strange that we continue to consume on a scale much larger than what the environment can possibly handle. Rare earth metals, much needed in the *ICT* (Information and Communications Technology) sector are approaching depletion (Rockström, Klum, and Mania 2016), and each day, massive quantities of coal, oil and fossil gas are burned.

The thesis is written with this paradoxical situation in mind: we, as a species, continue to strive for everlasting growth, even though the environmental footprint of everyday life is now threatening not only growth but humanity itself. Many in the scientific community have recognised this dilemma, also after *The Limits to Growth*. The works of Johan Rockström (2009) and Kate Raworth (2012) and their associated “donut model” are based on the same philosophy and conclusions: we need to stay within the planetary boundaries (and as Mike Berners-Lee shows (2021), moving to



another planet is simply not possible, just looking from an energy perspective). Since it is now proven that everlasting material growth threatens our way of life, the Keynesian idea of spiralling consumption as a driver for future societal and individual wellbeing must be retired. We can only hope that it is not too late to undo the damage humans have already done to the planet.

Relating to computer science, a significant contribution to climate change is the power draw of computers. Norbert Wiener foreshadowed this problem in his foundational work *Cybernetics* (1949). In this book, he essentially outlined the principles for building a human being, not entirely unlike Dr Frankenstein a century earlier, but making an effort to do this with modern science. With regards to the brain, he wrote:

*“... let me point out that a large computing machine, whether in the form of mechanical or electric apparatus or in the form of the brain itself, uses up a considerable amount of power, all of which is wasted and dissipated in heat. ... unless adequate ventilating and cooling apparatus is provided, the system will suffer from what is the mechanical equivalent of pyrexia, until the constants of the machine are radically changed by the heat, and its performance breaks down.”* (Wiener 1948)

Of course, the machines Wiener was referring to – ENIAC and EDVAC – had very different energy use than today’s computers. That being the case, in his design of the human brain, he was correct to pinpoint energy use as a coming problem (and possibly he was the first to do so), as well as cooling and ventilation as a necessary measure to prevent the *“mechanical equivalent of pyrexia”*. Modern computers are magnitudes less power-hungry. For example, the laptop this thesis is written on uses between 15 and 45 watts, less than a traditional tungsten light bulb. Still, with the

millions (or billions, depending on one's definition of the term) of computers in use today, the accumulated use of electricity is high enough to be a part of the climate problematics.

This PhD work circles sustainability and *data centres*, large collections of computers (called “servers”, as opposed to “client” computers used by people) that store and compute the data we use every day, be it Google Maps, family photos or the classified files of the FBI. Computation of these data comes at an energy cost. Data centres' high energy needs mainly originate from processor-intensive computations on the servers. Today's data centres use approximately 1% of world electricity (Masanet et al. 2020), and it is likely this use will increase further (Freitag et al. 2021). For EU, the energy use of data centres was approximately 80 TWh in 2018, and the EU expects the figure to approach 100 TWh in 2030 (European Parliament 2023a).

Over the last fifteen years, research for energy-saving measures in data centres has intensified, a notable change from the earlier primary focus on computational performance (e.g. (Kooimey et al. 2008)). There are two reasons for this change of focus:

1. Moore's law, which states that the number of transistors on a microchip doubles about every two years and has proven correct for decades, is finally reaching an end. One reason for this deceleration is that with wiring as thin as a few nanometres, we are approaching physical limits for accurately pushing electrons around. (The undesirable effect, “quantum tunnelling”, means that a particle moves through an energy barrier. In microelectronics, it refers to the risk of electrons not reaching its destined path (gate) in the transistor. Such a fault would cause the electronic device to act erroneously.) Another reason Moore's law is difficult to maintain is that

with higher transistor densities, processor units (CPUs and GPUs) become hotter. Without sufficient cooling, they would soon malfunction.

2. In the mid '00s, data centre energy use was primarily regarded a (huge) financial problem (e.g. (Mills et al. 2008)), but shortly thereafter, environmental concerns grew bigger (see chapter 9). With financial and environmental concerns coalescing, energy savings are now high on the agenda for any data centre owner, including industry giants such as Google, Meta (the parent company of Facebook, Instagram, WhatsApp, Messenger and more companies), Apple, Amazon and Microsoft, collectively referred to as *Big Tech*. The energy use of data centres has also received gradually higher interest by legislative bodies, such as the EU. It is telling that the union's energy directive is currently in the process of being rewritten to include data centres (European Parliament 2023b). In contrast, an anthology on renewable energy law in the EU from 2014 does not mention data centres, IT or ICT (Peeters and Schomerus 2014).

To conclude: server energy use is a problem today, for financial, environmental and also technical reasons. Accordingly, much has been written about energy savings *within* the data centre (see chapter 8). However, there exists a strong case for moving beyond it: due to *rebound effects* – a phenomenon where improvements on one end increases demand, thus obliterating the initial benefits – energy savings inside the building enable more data centres.

In the IT and entertainment industries, rebound effects are commonplace. An ordinary computer game for the PC may take up *a hundred thousand* times more storage than the Commodore C-64 games of the 1980s. Still, many users play “retro-style”

games today, and while today's PC games are better in every aspect than the games of the 1980s, they might not be a hundred thousand times more fun to play. Likewise, as computing, storage, and delivery technology have evolved in the 2010s, YouTube movies have moved from 720p image quality to HD, 4k and even 8k. As anthropologist Steven Gonzalez Monserrate puts it,

*“virtual reality constructs powered by machine learning with ultrahigh bandwidth (8k definition) have accelerated anthropogenic climate change. Demands for cloud services have outpaced sustainable growth, leading to cascading systemic failures, data rationing, and more.”* (Gonzalez Monserrate 2022)

The rebound effects that Gonzalez Monserrate identifies are part of data centre *energetics* – the concerns associated with energy. However, as will be apparent in this thesis, data centre energetics need investigation from many perspectives. Therefore, instead of investigating energy savings on the processor, server, or rack level, or purely from a societal perspective, this PhD work involves energy use of data centres *holistically*, put within an infrastructural, societal and environmental context. But perhaps surprisingly, the principal aim of the thesis is not to limit the energy use of data centres, but to question the very idea of energy consumption, and think of the used energy as a commodity with potential value.

Indeed: *what if we do not view data centre waste heat as consumed energy but as a resource?* The thesis investigates that question.

## **1.1 Chapter outlines and contributions**

The thesis spans several disciplines, and due to its complexity, it has its roots in systems science, the most cross-disciplinary academic field in existence (Shim and Bellomy

2018). Moreover, the work is grounded in exploratory research. The aim is to a) investigate the very rationale for and feasibility of data centre waste heat use, b) provide a framework for potential implementation of sustainable data centres and – through the framework and new metrics – c) establish a baseline for further research.

After the introduction, the literature overview and the methodology perspective chapters, the thesis is divided in three parts as follows: *Part 1: 4. Data centre energetics, 5. A framework for data centre placement; Part 2: 6. Applications for data centre waste heat, 7. Data centres and society, 8. Data centre energy efficiency metrics; Part 3: 9. Analysis and 10. Discussion.* Conceptually, the main chapters are grouped into theory, case investigations and analysis. This section briefly presents each group and its individual chapters.

### ***Part 1: Framing the problem and identifying dimensions of concern for data centre placement***

Chapter 4. *Data centre energetics* introduces the problematics of energy, clean energy, the circular economy, energy use in data centres and how data centre waste heat fits into a societal context. The chapter lays the theoretical foundation for the rest of the thesis, and situates the research in a larger discussion.

Chapter 5. *A framework for data centre placement* builds on the previous chapter and investigates more specific issues pertaining to the data centre context. Among the topics dealt with are energy efficiency measurements, the future of server cooling and the societal value of data centres. At the same time, this long chapter identifies and examines the particulars of data centres in need of investigation for the research to be carried out and properly evaluated. The chapter results in the thesis's analytical framework, based on these identified particulars.

## ***Part 2: Applications, societies and metrics***

Chapter 6. *Applications for data centre waste heat* discusses use cases for data centre waste heat under different geographical conditions. The chapter is based on findings presented in papers and at conferences, and it discusses novel uses for this waste heat.

Chapter 7. *Data centres and society* illustrates how some of the proposed ideas from the previous chapter can be implemented to benefit a society. It does so through illuminating the feasibility of the applications for three carefully chosen societal contexts: large-scale commodity drying in Malaysia, coffee bean drying in Costa Rica and district heating and wooden pellet dehydration in the Nordics.

Chapter 8. *Data centre energy efficiency metrics* discusses the problem space of current energy efficiency metrics, based on the old management saying that “you can’t manage what you can’t measure”; for the ideas discussed in chapters 8 and 9 to become reality, management and the authorities must be able to recognise and evaluate the potential gain. The chapter then explores diverging ways to define replacement metrics, presents a formula for a new possible key metric, and discusses the two options in light of existing metrics. Last, the importance of data centre energy metrics is shown through a discussion of the update of the EUs Energy Efficiency Directive, which is currently being formed.

## ***Part 3: Analysis and discussion***

Chapter 9. *Analysis* brings together findings from the three investigative and experimental chapters with the context framing of chapters 4-5. The analytical framework is used to discuss the merits of the cases of chapters 7 and 8, after which the metrics of chapter 9 are used for more quantitative evaluation of the cases.

Chapter 10. *Discussion*, finally, draws on the conclusions from the analysis to present an expanded view of the contributions of data and heat energy to a society. Thereafter, the findings are tied to the research questions introduced below. The chapter closes with suggestions for further research and further action by industry.

## **1.2 Assumptions**

This thesis is built around a set of general assumptions. Some physical assumptions are obvious, such as that the world faces an energy shortage problem as well as increasingly severe climate change effects, effects that, in turn, are energy-induced and a result of too *much* available energy in the biosphere. A similar assumption is that communities in low-, mid, and high-income countries alike should preferably be developed with respect to social and environmental sustainability.

In addition to those general assumptions, there are four specific assumptions being made in the thesis, as described in the following.

### **1.2.1 The warm airflow from data centre exhausts can be repurposed**

The first assumption has been proven for large-scale data centres in several cold regions, such as the Nordics (see section 7.5). Since reclaiming data centre waste heat in warmer climates is a novel idea, it has yet to be proven that repurposing data centre waste heat in those contexts is indeed possible and financially sound. This is why the thesis investigates site-specific uses and establishes a framework for such reclamation.

### **1.2.2 The citizens of the local community have access to smartphones and/or computers (desktops, laptops, tablets)**

The second assumption is met by just about every society in the world (GSM Association 2020a). In fact, even in 2015, *more people in low-income Sub-Saharan countries had access to mobile telephony than to electricity* (The World Bank 2016),

and typically, though network coverage may be unreliable and data transmission costly (Houngbonon, Le Quentrec, and Rubrichi 2021), the phones themselves come with Internet access capabilities (GSM Association 2020a).

### **1.2.3 The society has sufficient access to electricity to run the data centre**

The third assumption is required mostly for sustainability reasons. Whilst it is possible to run a data centre on diesel, adding to fossil-fuel dependency is not in alignment with the philosophy, values and goals underpinning this thesis. On that note: also where the power grid is based on fossil fuel – be it a diesel generator for a remote community in Papua New Guinea (Atteridge and Savvidou 2019) or Canada (Pinto and Gates 2022), or a coal-fired power plant serving millions in China or Poland – the argument for waste heat reclamation persists.

### **1.2.4 The society has a need for dehydration, energy storage or energy production as envisaged in this thesis**

That the society has a use for the waste heat is not obvious, as will be seen in chapters 4 and 5. Hopefully though, a society not currently reusing waste heat may gain from doing so, which is why a holistic perspective on societal value and monetary and environmental cost is needed. Hence, as argued throughout the thesis, with good planning and an understanding of site-specific barriers and opportunities, a data centre can be an integral part of a sustainable society.

## **1.3 Purpose**

Based on assumptions relating to data need, waste heat use options in a few selected societies and advances in data centre energy efficiency metrics, the purpose of this thesis is to investigate potential uses for data centre waste heat under different conditions. More specifically, the purpose is to explore potential data centre waste heat



uses as well as implications for the local society, the environment, data centre configuration and data centre metrics through a holistic approach to data centre energetics.

## 1.4 Research questions

To fulfil the purpose of this thesis, the following research questions (RQ) are investigated:

- RQ 1: How can data centre waste heat be used in different geographical and societal contexts?
- RQ 2: What are the societal implications and environmental benefits of implementing the ideas proposed?
- RQ 3: Which is the optimal design for a sustainable data centre?
- RQ 4: Can there actually be such a thing as a one key metric for data centre energy efficiency, and if so, how should it be designed?

These research questions form the foundation of chapters 2-9 and their themes percolate throughout the argumentation. The questions are reflected on in the concluding chapter.

The research questions above need multi-disciplinary investigation, but none of the existing academic cross-disciplinary fields has been found suitable to address the questions fully. Hence, the thesis is grounded in *material social studies*, a field proposed in the thesis, where material and social/societal enquiries are given equal weight. Hence, an ancillary research question, pertaining to methodology, is:

- RQ 5: To what extent is material social studies suitable for investigating cross-disciplinary problems?

## 1.5 Limitations

The scope of this study is limited in terms of access to data centre energy data. This is a general problem for data centre industry research, as such information is not readily published on regular basis by datacenter providers. However, the many types of data centre configurations and server (computer) models used, and the great variability in computation workloads run on these servers, give very different results. Thus, because of different server and data centre configurations, such data would have limited value for this particular thesis. In addition, though much research has been carried out on energy savings within the data centre (see section 8.1), this is – as shown in chapter 9 – not where the lion’s share of future energy savings can be made.

Another limitation is that the thesis only relates to *energy use*. Thus, it is not concerned with all aspects of a data centre, such as building accreditation certificates. Neither is it concerned with greenhouse gas (GHG) emissions directly (only indirectly through energy use), water use for cooling systems, land use of the facilities, or, for that matter, with the lifecycle of servers – itself an important topic. Hopefully, the thesis can inspire to – and serve as a template for – such future work.

Without a doubt, more investigated cases would provide a more detailed view of the problems addressed, and hopefully indicate patterns for future strategies with increased granularity. To compensate for this issue, the three societal cases have been chosen to reflect the diversity of data centre and society energy concerns.

## 2 Related work

This chapter highlights data centre research, and especially works on data centres seen in a societal context.

### **2.1 Data centre research – an overview**

The data centre industry is a big and highly exposed business sector that changes rapidly, and academic research relating to data centres' place in society is met by an abundance of news coverage and marketing material. The social sciences have largely been absent from data centre research (Hu 2015), meaning engineering enquiries have dominated the topic. This lack of interest (which is explored by Ipsen (2018)) is notable, as inside an often unspectacular and understated physical appearance, a large-scale data centre is an entity with many ties to society. To be fair, the social sciences have not had a long time to study this sector, which, as a consequence, is "*among the least studied areas of digital culture*" (Hu 2015). Moreover, whereas engineering is by nature forward-looking, much of the social sciences investigates the present, or the past. This is because – as Hegel put it – the owl of Minerva spreads its wings only with the falling of dusk: we can only know in hindsight. Indeed, given their importance to today's

societies, the history of Big Tech and large-scale data centres is astonishingly short: Google was founded in 1998, Facebook in 2004. The iPhone, dependent on its communication with data centres, was released in 2007. The web itself – basically the enabler of the data centre industry – became widespread first during the very final lustrum of the previous millennium.

But before discussing social and societal aspects of data centres, it is reasonable to investigate engineering concerns of data centre matters, since data centre energy matters are ultimately bounded within a physical context. The various constituents of the data centre – servers and their performance, server racks, cooling systems, and so forth – have been studied at length individually. However, their integration with other parts of a data centre’s infrastructure has been less explored. Some researchers, such as Cheng et al. (2021), Ebrahimi et al. (2014), Levy (2019), Li et al. (2018), Guitart (2017) and Manganelli et al. (2021) are among those who have, in fact, examined the relationships between data centres’ internal components, but also they stop at the data centre walls: societal integration has not been part of engineering enquires.

Fan et al. (2016) have tried to lower energy use through optimising computation workloads between data centres globally, depending on latency, network energy use and data centre energy efficiency. Today, researchers explore the use of AI as a strategy to lower energy use (e.g. (Mahbod et al. 2022)).

The literature also includes research relating to energy savings for the data centre as a building. Even at an early stage, Jamalzadeh and Behravan sought to establish a holistic framework of data centre-related metrics (Jamalzadeh and Behravan 2012). Some followed, such as Zakarya (2018). Tozer et al. (2018) proposed a metric based primarily on the life cycle analysis (LCA) of data centre components, energy use and

water use, but do not recognize waste heat use. Lei and Masanet (2020) explored energy use for free data centre cooling scenarios and a few have turned their focus to waste heat integration (e.g. Wan et al. (2021) and Wahlroos et al. (Wahlroos et al. 2017; Wahlroos et al. 2018)), but again, from an engineering viewpoint.

In addition to the academic literature, industry recommendations exist, from Google (Barroso, Hölzle, and Ranganathan 2019) and from accreditation bodies such as the Uptime Institute. However, more research of data centre infrastructure is needed, not least with regards to “future” technologies now being deployed, such as liquid immersion cooling (see section 4.4).

Written before the advent of social media, the smartphone and other major software and hardware components of the e-business ecosystem, Star’s “The Ethnography of Infrastructure” (Star 1999) is an important source for the comprehension of infrastructure and its role in society. Not only does she investigate infrastructure through an ethnographer’s lens, but she also proposes a methodology for ethnographers to undertake investigations on infrastructure.

As it happens, some anthropologists have investigated the *societal impact* of data centres. The anonymity of data centres has been studied by Vonderau (2021) and Hu (2015), and how they underpin AI activities, with a slew of ethical problems, was recently dealt with by Crawford (2021). A similar approach was taken by Lucivero (2019). Big Tech’s market dominance of IT, through e-commerce (e.g. Amazon) and closed ecosystems (e.g. Apple) is a theme that has been explored many times, most recently by Lehdonvirta (2022). Broader society-related connotations have been investigated by Sovacool et al. (2022), Velkova (2016) and, again, Vonderau (2019).

Since social and societal views on data centres are not so obviously tied to a thesis on waste heat, the views are discussed in the following.

## **2.2 The anthropology view**

Through novel software and hardware technologies as well as a long trend of moving data and computation online, the data centre industry quickly changes our personal lives and work-lives. Its expansion is determined through negotiation over energy, land use and water use with local and national authorities. Some of the concerns of the servers that data centres host – for example, regarding AI computation, social media, privacy, crypto-mining, copyright issues and antitrust laws – further limit their growth through restrained access to the more politically delicate services hosted within them. The data centres are therefore an integral part of society, shaping it, and simultaneously, being shaped by it.

Whilst there exists an abundance of literature relating to data centres, as well as to society, there are not many works discussing the relationship between data centres and society per se. The few academic works published in these quite overlapping fields are written by anthropologists, ethnographers and geographers. In the following, authors writing in this style are referred to as the “anthropologists”.

Kate Crawford’s acerbic *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* projects many facets of data centres and their output in terms of societal gain... or rather, societal risk. For example, Crawford discusses how San Francisco once changed due to the exploitation of its surroundings, and how it now does so again:

*“As San Francisco drew enormous wealth from the mines, it was easy for its populace to forget where it all came from. The mines were located far from the city they enriched, and this remoteness allowed city dwellers to remain ignorant of what was happening to the mountains, rivers, and laborers that fed their fortunes. ... The pulley systems that carried miners down into the mine shafts were adapted and turned upside down to transport people in elevators to the top of the city’s high-rises. ... The ores extracted from holes in the ground were sold to create the stories in the air; the deeper the extractions went, the higher the great towers of office work stretched into the sky.*

*San Francisco is enriched once more. Once it was gold ore that underwrote fortunes; now it is the extraction of substances like white lithium crystal. The technology industry has become a new supreme interest, and the five biggest companies in the world by market capitalization have offices in this city: Apple, Microsoft, Amazon, Facebook, and Google. Walking past the start-up warehouses in the SoMa district where miners in tents once lived, you can see luxury cars, venture capital-backed coffee chains, and sumptuous buses with tinted windows running along private routes ... But only a short walk away is Division Street ... where rows of tents have returned to shelter people who have nowhere to go. In the wake of the tech boom, San Francisco now has one of the highest rates of street homelessness in the United States. The United Nations special rapporteur on adequate housing called it an ‘unacceptable’ human rights violation, due to the thousands of homeless residents denied basic necessities of water, sanitation, and health services in contrast to the record number of billionaires who live nearby. The greatest benefits of extraction have been captured by the few.” (Crawford 2021)*

The rhetoric in the above passage is not unusual for anthropology: giant corporations or corrupt country leaders exploit humans as well as natural resources (as exemplified by the rhetoric of Velkova and Crawford in this chapter). This view may be challenged, but it is not without merit, in that anthropologists and others in social sciences have established themselves as the academia watchdogs. And in a way, the

rhetoric makes sense: after all – why write about well-functioning societal constructs?

And to solve issues, is not identification of the issues required?

Crawford's text points to a number of problems that can be solved. Several of them, such as mining activities, do not relate to the problem space associated with the data centre industry (energy use, water use etc, as will be discussed throughout this thesis) directly. On the other hand: though rare-earth materials needed for AI computation is not part of the data centre structure, a data centre cannot sustain without the servers, and so, mineral extraction problems are closely linked to the sector.

Velkova's *Data that warms: Waste heat, infrastructural convergence and the computation traffic commodity* (2016) investigates the transformation of data centres from information carriers only to now also become energy providers. "*What are the processes by which the data processing industry becomes an energy supplier, and with what societal consequences?*" she asks (Velkova 2016). Doing so, Velkova expands on Star's research and its focus on anonymous infrastructure, claiming that

*"by paying close attention to the ways in which urban infrastructures are rewired to transport waste heat, we could understand better the processes of ongoing renegotiations of the meaning, materiality and the value of digital data. Not least, we can see how the data industry contributes to extend [sic] the sphere of capitalist production and the digital economy by redefining waste into a desirable commodity."*  
(Velkova 2016)

Interestingly, Velkova here integrates Star's ideas on infrastructure generally with the more particular, such as the books by Crawford and Hu. In her article, she also investigates the strong capitalistic tendencies of the data centre industry: consumers are



made into “*symbolic raw materials*”, “*fed to the algorithms*” and so, “*audiences for advertisers to mine*” (Velkova 2016).

Velkova further notes how waste heat opportunities say nothing about the source of the energy used. She finds that

*“valorising waste heat arguably integrates the data centre industry with the energy sector. Yet, this integration does not happen by generating truly green energy. Even if the data industry claims to be an active agent against global warming that reduces its GHG footprint by creating infrastructural loops of renewable heat, none of the approaches discussed in this article provide an actual alternative to polluting energy sources that power data centres with electricity, such as coal. Rather, the data centre industry relies on the existing sources of power available in each specific location, and these can differ substantially.”* (Velkova 2016)

Her finding has linkages to a key point in this thesis, in the development of new data centre energy sustainability metrics. In fact, one of the metrics proposed in chapter 8 incorporates both the source of energy and the waste heat.

From a very young age, people are bombarded with symbolic images. The same way many children worldwide know what a (culture-specific) dragon or a unicorn look like, a data centre has its own symbolic language in everyday culture. As Taylor writes, the image “*of the server-cabineted corridor ... typically bathed in blue neon, has become the canonical icon of the data centre*” (Taylor 2017). Matrix and other sci-fi is not far away (see Figure 1). This symbol says little about the data centre as a building. For one thing, one gets the idea that all data centres are extremely large, which is not the case at all. Second, the often restrained façade of the building is less often shown.

Third, the infrastructure tying the data centre to the outside world through supplies of electricity and water and through cables for data transmission, is missing.

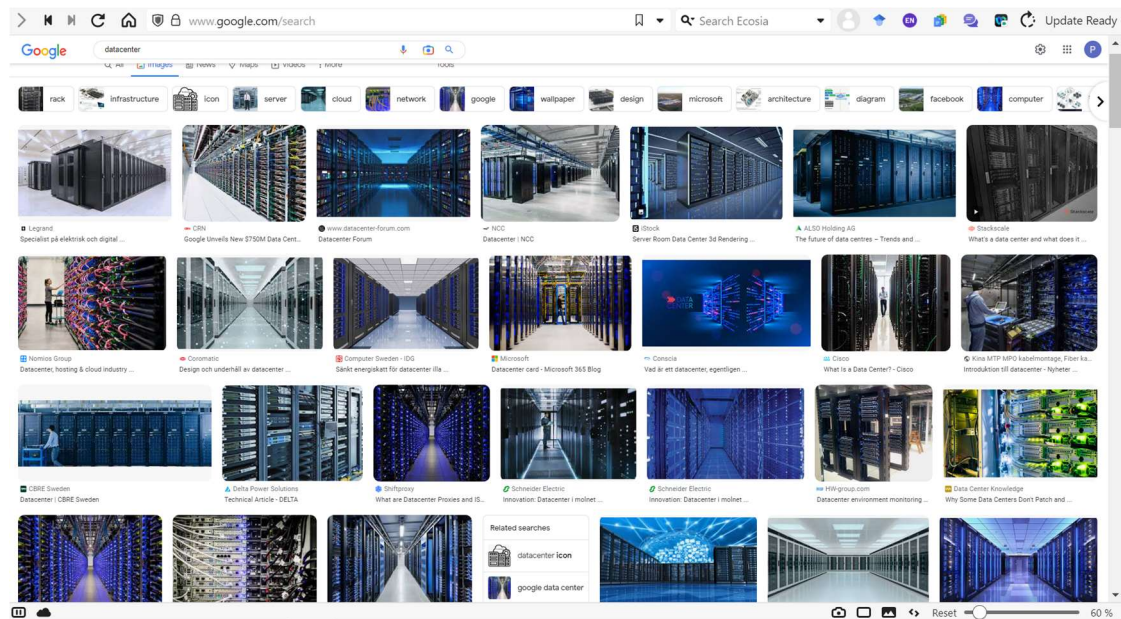


Figure 1. What is a data centre? This Google search on “datacenter” shows that whatever it is, its colour is neon blue. (Screenshot in February 2023.)

For researchers in media studies and architecture, the data centre provides many opportunities for comparisons to other objects: nineteenth century factory layouts, depictions of spaceships in sci-fi movies and games and much more. These “*virtual*”, “*future-proof*”, “*sterile*”, “*transparent*”, “*revealed*”, and, ultimately, “*clouded*” spaces (Taylor 2017) shown in Figure 1 are discussed by Taylor (2017); how the data centre architecture can better interplay with its surroundings by Pfeiffer (2020). Jennifer Holt and Patrick Vonderau reveal a deliberately constructed image: the colourful and clean environment of Google (holding on to millions if not billions of family photos) starkly contrasts to a Stockholm-based nuclear-bomb-proof data centre (that, aptly, used to host WikiLeaks (Baltzer 2010)):

*“how can we overlook the polity-building processes implied in Google’s infrastructure design – its lively colored pipes, well-organized lines of glowing server racks in shades of blue and green, and brightly illuminated architectural spaces – as compared to*

*Bahnhof's underground cold war bunker setting and historical engine for backup power?"* (Holt and Vonderau 2015)

Yet again crossing the border between the material and the social worldviews, Velkova adds that “data centre operators do not offset the environmental problems that the industry generates, but rather reshape the discourse around it” (Velkova 2016). And reshaping the discourse is not that difficult: if the waste heat is used, the origin of the electricity may matter less to media.

From the works mentioned above, it may seem as if the thesis proposes a view of the data centre industry founded in the social sciences. Such a presupposition would only be half-true: there are limits to what the social sciences can do from a methodology viewpoint. Star writes that it is easy

*“to stay within the traditional purview of field studies: talk, community, identity, and group processes, as now mediated by information technology. There have been several good studies of multiuser dungeons (MUDs), or virtual role-playing spaces, distance-mediated identity, cyberspace communities, and status hierarchies. There are fewer on the effect of standardization or formal classification on group formation, [on] the design of networks and their import for various communities, or on the fierce policy debates about domain names, exchange protocols, or languages.*

*Perhaps this is not surprising. The latter topics tend to be squirreled away in semi-private settings or buried in inaccessible electronic code. Theirs is not the usual sort of anthropological strangeness. Rather, it is an embedded strangeness, a second-order one, that of the forgotten, the background, the frozen in place.”* (Star 1999)

Provided access is gained (which, by the way, is unlikely), an ethnographer can certainly survey these strange, forgotten places. However, as Star states, the researcher’s

comprehension would be limited to topics such as identity. Likewise, though data centres are inevitably material constructs – products of steel and concrete with extremely tangible concerns such as physical security, service level agreements (SLAs) and easily measurable performance data – they are rooted in and dependent on a societal context. Therefore, researchers would be helped by a broader understanding of the Internet, computers, engineering and management. Hence, crossing the borders between the material and the social disciplines can become advantageous to future research – and how to cross the border is discussed in the next chapter.

### **2.3 The duality of the data centre cosmos**

A key point in this work is the *duality* of data centres: they are both material constructs and carriers of information, thus “*both cultural and technological*” (Gonzalez Monserrate 2022). How they are viewed depends mainly on one’s relation to them. For people in Luleå in northern Sweden, where Meta’s (“Facebook” at the time) first European data centre was built, a data centre is quite physical. For many others, it is just an intangible object where data is stored: “*catching the cloud and pinning it down*”, as Ipsen set out to do (Ipsen 2018), is easier said than done.

Pointing to the “*paradoxes of infrastructure as both transparent and opaque*” (Star 1999), Star writes that infrastructure

*“is both relational and ecological – it means different things to different groups and it is part of the balance of action, tools, and the built environment, inseparable from them.”* (Star 1999)

In other words, infrastructure in general has many facets. In turn, these facets can be divided into immaterial (action) and material (tools and the built environment).

However, there is a big difference between a data centre and, say, a sewage system: The inputs and outputs of the sewage system are material, but *information* – that is, structured and meaningful data – a principal part of the data centre’s input and output, is not.

The duality of data centres and their problem space has not gone unrecognised by the anthropologists. Digital media scholar Tung-Hui Hu notes that “*data centers exist at the border between the dematerialized space of data and the resolutely physical buildings they occupy*” (Hu 2015). Crawford addresses the duality in more detail:

*“The mining that makes AI is both literal and metaphorical. The new extractivism of data mining also encompasses and propels the old extractivism of traditional mining. The stack required to power artificial intelligence systems goes well beyond the multilayered technical stack of data modeling, hardware, servers, and networks. The full-stack supply chain of AI reaches into capital, labor, and Earth’s resources – and from each, it demands an enormous amount. The cloud is the backbone of the artificial intelligence industry, and it’s made of rocks and lithium brine and crude oil.”* (Crawford 2021)

In a similar fashion, Oxford scholar Federica Lucivero mentions that

*“it is well acknowledged that ICT in general, and data centres in particular, have an ambiguous relationship with the vision of sustainability. Despite the promises of delivering a more sustainable world, ICT[s] also compromise this very vision through their high energy consumption and carbon footprint.”* (Lucivero 2019)

Asta Vonderau, now an anthropology professor in Germany, has studied social, societal and environmental concerns around the establishment of Big Tech’s establishment in Sweden. Like Lucivero, Vonderau notes that the duality is partly

orchestrated by Big Tech. Comparing to Big Oil, she discusses *entanglement*, arguing that

*“IT companies such as Facebook work towards a disentanglement – or at least controlled entanglement – of the cloud from its infrastructural localities, scaling the cloud as a global and virtual technological zone.”* (Vonderau 2019)

And the “*controlled entanglement*” makes sense from a business perspective: as long as one views the entity hosting all our data as a cloud – *such a fluffy, harmless object!* – delicate questions regarding sustainability and customer privacy may be overlooked by the end-user. Under scrutiny, end-users could “*disturb free data and profit flows*” (Vonderau 2019), and sometimes they actually do (e.g. (Moss 2022)). In addition, from an interaction design perspective, the cloud metaphor is somewhat of a strain, since the Internet consists of connected nodes and not of a singular entity. Using clouds as a symbol for distributed computation and storage may therefore seem like a deliberate decision to avoid too much reflection on the material and social costs of data.

As it happens, the “business perspective” also stretches to the authorities. Luleå, the town hosting the first European Facebook data centre (see section 7.5), defined itself as offering a “City as a Service”, with the aim of transforming the old steel shipping town into an international IT nexus. According to Vonderau, there existed a “*collective aspiration for effectivity and regional relocation [that] directs public attention away from problematic aspects of the new industry, such as its enormous energy and water needs, potential long-time environmental consequences, or the limited possibilities of regulation and control*” (Vonderau 2019).

Today, harmless or not, the term lives its own life (see Figure 2), puffing around on our computers, smartphones, tablets... and in the cloud.

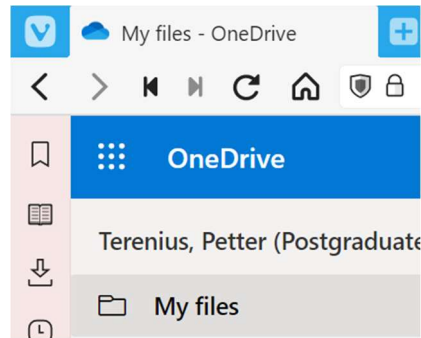


Figure 2. Naturally, this thesis is saved to the “cloud”, a term reflected in the OneDrive icon (to the left of “My files”). The fact that the cloud, in this case, happens to be a building located on the Lancaster University campus, is another story.

## 2.4 Chapter conclusion

This chapter has presented some works of relevance to the thesis; many more are integrated in the text, not least statistics from organisations such as the United Nations (UN) and the International Energy Agency (IEA). Of the academic works, many belong to other fields than computer science, such as systems science, management science, marine energy, food science and sustainability science. What is clear is that the literature of direct relevance to the core idea – to investigate data centre energetics from both material and societal aspects – is scarce.

In the literature concerning the environmental sustainability of data centres, energy use has been the primary topic of study. It is understandable, as energy use is problematic, but also as there are ways to make such research easily defined, clear-cut and often measurable. Data centre social (or societal) issues, on the other hand, are anything but.

In the overview of works written by the “anthropologists” above, examples have been provided of the given rhetoric of some anthropology research directions: a trope-filled language that stresses the dilemma of technological “progress” (though at least Lucivero acknowledges that the *“Big Data initiatives can be a resource for modern and*

*sustainable economies*” (Lucivero 2019)). The engineering discipline has a rhetoric too: here, one does the opposite – innovate as fast as possible, often leaving out a discussion of the possible consequences for people and for the environment. It seems these are opposing forces, where one side innovates and the other scrutinises this innovation. To remedy the situation, the next chapter introduces a new constructive and responsible paradigm, where the two existing views are complemented by *more reflective material views* and *more constructive social views*.



# 3 Methodological perspective

The universe expands at great and accelerating speed, and so does human knowledge. During the high renaissance, Leonardo da Vinci – the epitome of the polymath – excelled equally in painting and engineering, and moved freely between a range of scientific and artistic disciplines. In the 1700s, Linnaeus was simultaneously a botanist, a medical doctor, a zoologist and an explorer. Today, the body of knowledge is so large that it is usually not possible for one researcher to span more than his or her own scientific discipline.

Against this background, sections 3.1-3.4 of this chapter present a methodological toolbox meant to structure the problems of data centre sustainability and light the way to possible solutions where unidisciplinary approaches are not enough. The toolbox – containing *the material social perspective*, *systems science*, *systems engineering and mixed methods* – also makes up the foundation for this thesis. How the methods of this toolbox have been applied to the main constituents of the work is discussed in section 3.5.

### **3.1 Methodological approaches of this thesis**

This thesis aims to tackle the challenges described in chapter 1 partly through viewing several fields of research collectively as a system. The goal is to combine subsets of these fields of research, which is logical given that the problems discussed in the thesis are holistic in nature. Thus, a cross-disciplinary, holistic and systematic approach would seem reasonable for this discussion. The case studies supply the research with a higher degree of reliability (ability to reproduce the results across time and among researchers). Case study selection and site selection have been performed through consulting the literature and sources such as maps, meteorological charts, demographic data and ICT access data. How this is done is shown in chapter 7. The case studies are then collectively run against an analytical framework (see section 3.5, chapter 5 and section 9.1). Thus, the framework can be used to check the reliability of the methods used. In addition, the thesis includes a new metric for data centre sustainability evaluation. The three cases are partly used to validate the metric, meaning the qualitative and the quantitative meet here (see section 9.2).

Figure 3 shows the thesis structure from a methodological viewpoint. Here, after the literature review defines theory and gaps within theory, systems approaches are suggested as a tool to attack the complexity of the initial problem. To make use of systems science and systems engineering principles, material social studies is proposed. With the methodological stage set, a framework based on the literature review is constructed, and after that, the exploration begins. The mainly quantitative work on data centre energy metrics proceeds in parallel with the search for applications. Based on the applications found, vastly contrasting societal cases are studied. The analysis chapter runs proposed solutions for each case study against the existing and proposed data

centre sustainability evaluation metrics. A discussion brings the thesis to an answer to the initial set of questions, and the work to a close.

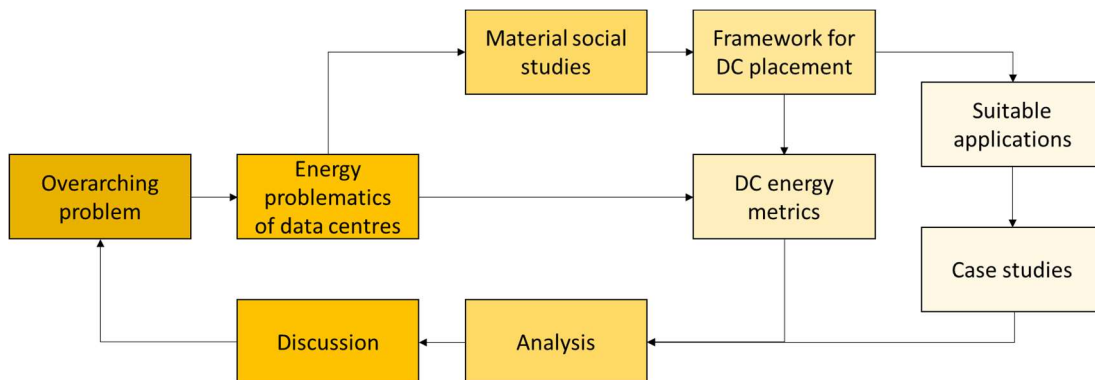


Figure 3. A methodological view of the thesis.

The thesis circles environment, society and technology. These entities can all be viewed as *systems*, which is why systems science and its derivate, systems engineering, are discussed below. First, a new paradigm for how to address and ultimately solve multi-disciplinary sustainability-related questions is presented in the following.

### 3.2 The material social perspective

Presently, humanity faces enormous problems concerning climate change and ecology as a whole. We must also deal with challenges regarding increased populations and sometimes accelerating social problems in cities and rural areas. In response, the United Nations' *Sustainable Development Goals* (SDGs) try to address such problems via a hierarchical view of them. This view is somewhat intertwined. For example, “SDG 5 Gender equality”, a goal in its own right, also exists as a subset of “SDG 10 Reduced inequalities” (more specifically Target 10.2).

In an increasingly connected world, many of today’s issues are multi-faceted and intricate, or “*wicked*” (Churchman 1967). To give but one example: energy use in a municipality is not only about limiting energy in households and industries, but about

energy forms, properly built infrastructure for district heating or other solutions, legislation regarding insulation, investigating of renewable energy sources, how to teach and persuade the population not to waste energy, the advocating for nationwide sustainable energy policy, and many other tasks. With local and global environmental implications (pollution and climate change being two examples) and societal problems on top of these, energy use in a municipality may seem a straight-forward problem but is in fact incredibly complex, comprising many interrelated problems, many technical solutions and – to further complicate things – many actors, in turn with their many wishes, desires and open and hidden agendas.

Undeniably, some problems of major concern to the world today need expanded views. Solving the problems of energy use within the municipality in the example above would require research in both engineering and social sciences. The *material social* perspective, stemming from (and taught at) Lancaster University, recognises this need:

*“In short, the material and the social go hand in hand in complicated and dynamic ways.*

*Until recently, however, research tools and techniques have tended to separate these concerns, with material scientists working without reference to social matters, and social and humanities researchers knowing little about materials science. As a result, researchers from any discipline have not been able to make the future in ways that bring the benefits of materials science alongside better understandings of social and human values.”<sup>1</sup>*

---

<sup>1</sup> <https://www.lancaster.ac.uk/material-social-futures/>, accessed 10 November 2022.

There is a constant interplay between the material and the social – and also the environmental. As noted by Huber (2015), environmental and economic historians now recognise that in addition to technical systems, fossil fuels have reshaped our political, cultural, and environmental relations for the last two centuries. Evans et al. see similar patterns for plastics (Evans et al. 2020), and Alan Blackwell, a Cambridge professor of interdisciplinary design, for societal perspectives on ICT (Blackwell 2021). Through applying the material social perspective proposed above, data centre energy is here subject to related concerns. Similar to Rachel Carson’s system-based view of nature (see the quote that opens this thesis), Crawford notes how Big Tech deliberately pushes environmental concerns and the problem of conflict minerals to third-party contractors:

*“Just like the mines that served San Francisco in the nineteenth century, extraction for the technology sector is done by keeping the real costs out of sight. Ignorance of the supply chain is baked into capitalism, from the way businesses protect themselves through third-party contractors and suppliers to the way goods are marketed and advertised to consumers. More than plausible deniability, it has become a well-practiced form of bad faith: the left hand cannot know what the right hand is doing ...*

*While mining to finance war is one of the most extreme cases of harmful extraction, most minerals are not sourced from direct war zones. This doesn’t mean, however, that they are free from human suffering and environmental destruction. The focus on conflict minerals, though important, has also been used to avert focus from the harms of mining writ large.*

*If we visit the primary sites of mineral extraction for computational systems, we find the repressed stories of acid-bleached rivers and deracinated landscapes and the extinction of plant and animal species that were once vital to the local ecology.”*  
(Crawford 2021)

One argument that follows from Crawford's argumentation is that by claiming "conflict-free" mining of cobalt, gold, diamond or lithium, the commodity is presumed sustainably mined. Other aspects of concern, such as hazardous waste from the mining process or the move of the local population to start a mine, may therefore be invisibilised. Moreover, as Crawford shows here, technology, society and nature are connected. This is also the case for *strong sustainability*, where society and its doings can only exist within the limitations given by nature (Figure 4). Worryingly, we are transgressing many of these limitations today (Rockström et al. 2009; Rockström, Klum, and Mania 2016).

In academia, the humanities can be regarded the opposite of the core sciences (Foley 2018). The social sciences are somewhere in-between, or rather: much depending on the mix of quantitative and qualitative methodologies, the research tends to be closer to the humanities or to the core sciences. Foley writes that

*"the overall propensities of the sciences and the humanities are poles apart. Assimilation and simplification on the one hand; diversification and complication on the other. Each has its place. There is potentially enormous power in simplicity and potentially great richness in complexity."* (Foley 2018)

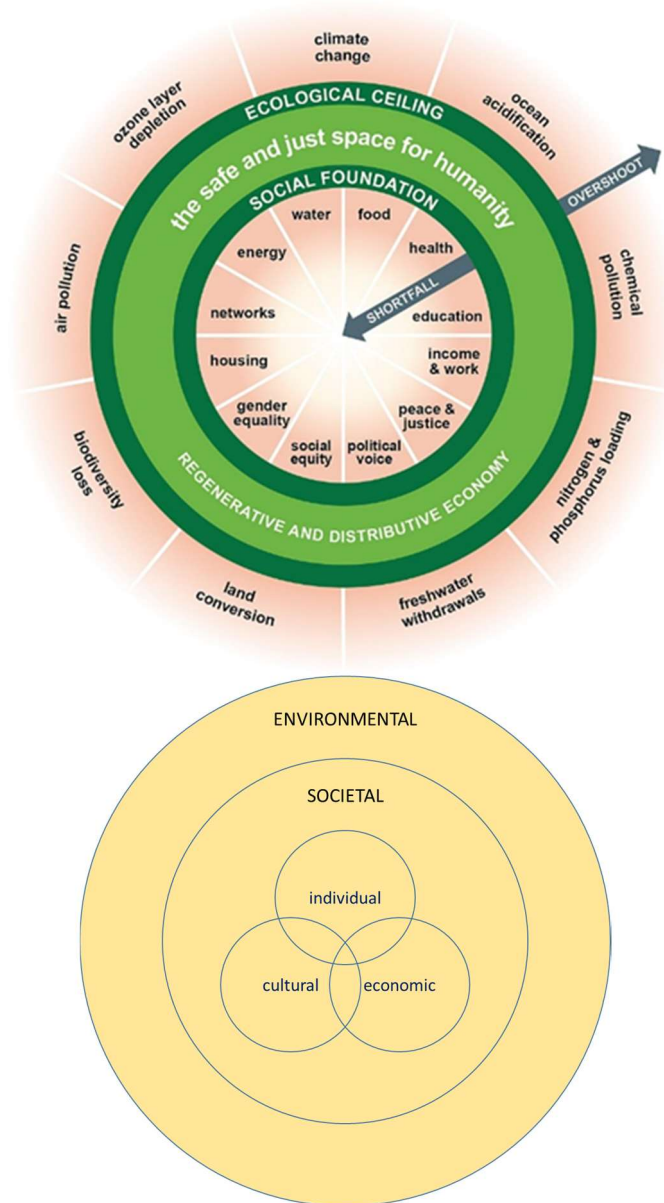


Figure 4. The doughnut model (left), by Kate Raworth (2012) after Johan Rockström and Will Steffen’s definition of the “nine planetary boundaries” (Rockström et al. 2009), relating to land use, water use, food and many other human interests or activities. According to Rockström and Steffen, we need to stay inside these to not threaten the environment’s ability to self-regulate. Hence, this is one example of strong sustainability. Less sophisticated, the Rings of Sustainability model (right) developed for this thesis also displays a worldview of strong sustainability. In addition, the model recognises that a society’s degree of sustainability depends on its infrastructure and of people and their norms – or put in systems terminology: Societal sustainability is a superset of Individual, Cultural and Economic sustainability. (Cultural, Economic, Environmental and Societal sustainability form the “four pillars” of sustainability, advocated for by UNESCO and others (Sabatini 2019).)

From Foley's words, it seems plausible that research projects that take advantage of both the logical precision and generalisability of the core sciences, and the real-world lens of the social sciences and the humanities, may utilise some of this power.

The following sections describe ways to join the strengths of these two paradigms.

### **3.3 Systems science and systems engineering**

*"Finding fresh methods for understanding the deep material and human roots of AI systems is vital at this moment in history, when the impacts of anthropogenic climate change are already well under way. But that's easier said than done."* (Crawford 2021)

Crawford's words above indicate that solving data centre environmental concerns – energy included – is not trivial. She calls for *"fresh [research] methods"*, assumedly since the methods we have today are insufficient. In other words, today's perspective is too detail-oriented to solve the problems. Crawford further notes that we need to *"be connecting across multiple systems to understand how they work in relation to each other"*. That is precisely the approach of this thesis.

The main purpose of the PhD work is to combine various actors, themes and real-world problems to create sustainable and beneficial real-world solutions. The study of relationships such as these is at the heart of *systems science* and its derivative *systems engineering* (Figure 5). Indeed, systems science – the study of entities and the relationships between them – can provide a valuable tool when defining how the constituents of the data centre energetics bond interplay.

For decades, systems engineering has served as an enabler for the identification, study, design and undertaking of processes in business, engineering, sociology,



bioinformatics, ecology and sustainability science. The “waterfall model”, “iterative development”, “object-oriented analysis and design”, “systems thinking”, “holistic approaches” and “hierarchies” are all terms either established in or related to systems engineering.

In turn, systems science is the science behind – and *the art of* – such models, and it gives birth to such mindsets. To the systems scientist, a model is much more than a set of connected entities; it is a carefully thought out process, expressed as easily as possible, with the clear intent to serve disciplines such as the ones described in the above passage: in Watson’s words, “*systems engineering [and so, systems science] exists to develop a solution to meet a need*” (Watson 2019). Thereby, a model has the power to enable an understanding of complex phenomena, make for better decision-making, improve consultant-customer communication and successfully execute large-scale projects.

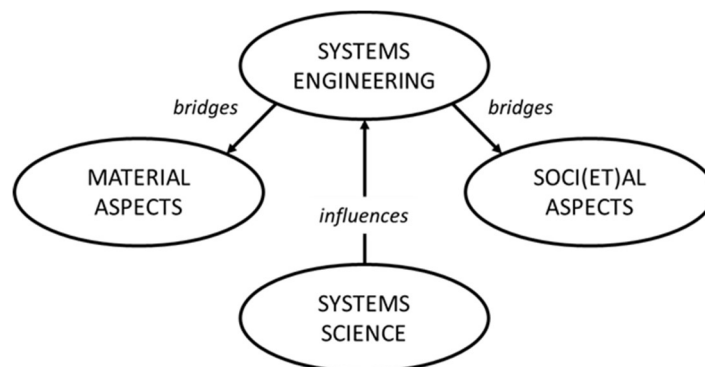


Figure 5. The methodological foundation of the thesis.

However, models do not have limitless power. For one thing, many of them show only entities and their dependencies, not time or triggers. From a more philosophical viewpoint, models can also be challenged as they focus on objects or *things*. With another attitude to systems, Star and Ruthleder hold that “*infrastructure is*

*fundamentally and always a relation, never a thing*” (Star and Ruhleder 1994). Their idea is to de-emphasise objects and focus on the connections between them. The authors exemplify with the way that railroads, timetables and management systems are related, writing that a focus on objects “*inverts traditional historical explanations*” (Star and Ruhleder 1994).

Whilst the idea has its merits, there is no reason why one or the other of these two component types should be emphasised: objects are pointless without relations, and relations need objects to exist (a railroad with no stations has no value). Complementing the strategy of Star et al. with more traditional strategies would mean to study the same problem from two perspectives, “object-first” and “relation-first”. Sometimes, the two approaches can cross-fertilise: this thesis is written from an object-first viewpoint (data centre problematics and components), but it is rooted in a relation-first attitude (societal needs and opportunities).

A related problem (and virtue) of systems science models is their immediate explainability, and so, their ability to persuade. For example, the diagrams and charts provided by systems engineers often work through inheritance or dependencies. In reality, many relations are multidimensional. As a consequence, one of its greatest powers – the ability to on just one sheet of paper explain a complex environment – is perhaps also its Achilles heel. As a researcher, one must therefore remember that what a model depicts is not the truth but a simplification. (An exception is a strict implementation of iterative development in software projects, where the complete system is moulded through a series of iterations. Going from a helicopter view of the system in the first iteration to more and more detailed specifications and finally code or engineered structures, there is no clear distinction between the system as an idea and

the developed system. As Larman and Basili note: in the latter, ultimately, “*the model becomes the system*” (Larman and Basili 2003.)

As one of the most challenging problems humanity has ever faced, tackling “*climate change is a complex systems engineering problem*” (Lu and Xu 2017). However, the systems engineer needs field expertise from every discipline: not least do we need natural scientists to understand nature, and economists, sociologists and psychologists to understand societies and individuals. This is especially true as – in the words of Vickers (1983) – a system made by humans is not fully technical but also a result of societal norms. Thus, insights in both technology and society are needed to understand systems that have positive or negative effects on our climate. The next section will show a way to blend such insights.

### **3.4 Mixed methods**

There is an evident dichotomy at play concerning natural and social sciences. This dichotomy was investigated in the 1800s by researchers such as Max Weber (1864-1920) and Wilhelm Dilthey (1833-1911), speaking of “*explorative*” (natural) science, where one searches for absolute truths, and “*descriptive*” (social) sciences, in which truths are based on one’s experiences (Tapper 1925). Thus, the latter truths may be absolute, but they are not universal. As Tapper observed in the early 1900s,

*“Dilthey is very explicit in the distinction which he makes between begreifen [comprehend] and verstehen [understand]. It is the ideal of the natural sciences to comprehend, but the ideal of the mental sciences is to understand.”* (Tapper 1925)

Nine decades later, systems scientists Lin, Duan, Zhao & Xu (2012) saw “*a barrier between the natural scientific world and the physical world with people that is*

*studied in social science*” and argue one reason for this dichotomy is the lack of a “*unified methodology for scholars to investigate natural and social systems jointly*”.

Ipsen, too, seems to recognise this gap, and calls for ways to overcome it:

*“The very nature of data centers as both real and virtual spaces requires that research crosses disciplinary boundaries. Their impacts are neither just physical or social, nor spatial or temporal, but a mix of everything at once. Much of the existing literature is technical, where scholars note the intensive energy use of data centers, but fail to engage with the associated social or political dimensions this infrastructure can bring about ...”* (Ipsen 2018)

Therefore, to follow Ipsen’s argument, investigating the problematics of the data centre industry requires a cross-disciplinary approach. As mentioned in the introduction to this chapter, undertaking cross-disciplinary studies is challenging. However, it is also rewarding, in that one is exposed to many sorts of methodologies, methods, writing styles and – not least – mindsets. Fully embracing cross-disciplinarity can mean working according to the *mixed methods* paradigm, useful when one needs a “*large toolkit of methods and designs to address complex, interdisciplinary research problems*” (Creswell and Garrett 2008).

Mixed methods, also referred to as “*multi-method*” or “*multi-strategy*” (Bryman 2006), is an umbrella term, indicating that a research project uses both qualitative and quantitative approaches. It can be argued that to some extent, most projects do. For example, a project on quantum physics could well include some sort of unstructured interview with a fellow researcher or a visit to another lab. A project on depression amongst millennials, in turn, would likely not only depend on deep interviews but also

use statistics and questionnaires with quantifiable responses to confirm one's claims and set them in a wider context.

What differs a “true” mixed methods project from others is the degree of *embracement* of both quantitative and qualitative methods: they must both be used in a formal manner and with adequate ambition. In other words, an informal lunch with a colleague would not make a quantitative project a mixed methods project. Moreover, the project must accept and recognise the inclusion and value of other methodologies. In an overview of earlier research, Bryman (2007) notes that many research projects were in fact using both quantitative and qualitative research methods, but that only one of these paradigms was used for the analysis stage, or at least, that one of them was given much greater attention than the other. He also notes that even in cases where both quantitative and qualitative methods were applied, findings from these projects were often presented in parallel. In other words, the research projects failed to integrate quantitative and qualitative findings.

Based on the assumptions of section 1.2, it is clear this is a multi-faceted thesis, involving topics from many fields of research. It makes sense that the topics should be dealt with in accordance with the traditions of each field; after all, there should be a reason why the fields use the research methods they, indeed, use.

### **3.4.1 Mixed methods in cross-disciplinary sustainability projects**

In the computing, sustainability and social sciences, researchers may use mixed method approaches in different ways and to varying degrees. The more qualitative aspects might involve people, whereas quantitative approaches can be useful for information-gathering and experimentation regarding the study topic (be it people, flora and fauna or municipality infrastructure).

Many problems addressed in mixed methods studies are what Churchman (1967), a pioneer in systems science, calls “*wicked*”, meaning they are difficult to solve, deeply entangled. Wicked problems usually refer to social problems; Skaburskis uses racial problems in American cities in the 1960s as an example (Skaburskis 2008). In more recent times, these often fall under the political ecology term, where social and ecological problems are either on the same side (and money on the other) or opposing each other (as in lithium mining, which would enable more electric cars but may come with local social issues (Bonds and Downey 2012)). Though political ecology tends to juxtapose two opposing forces, it is likely to assume problems such as these need to be viewed from many perspectives, in turn calling for several research methods. In concordance with this assumption, and as alluded to above, mixed methods research is found in several recent papers dealing with complex societal issues (e.g. Sioen, Terada et al. (2016) and Seele and Lock (2017)).

In a recent ethnographic study relying on mixed methods (Dencer-Brown et al. 2021), the data collection contained three stages. First, a literature review was carried out for general information about mangrove ecology and management. In parallel, the researchers investigated mangrove plantation managers’ attitudes and values regarding their work. Next, semi-structured interviews were held with a variety of actors. In a third stage, the researchers undertook fieldwork in mangrove plantations. Throughout the research process, there was a constant interchange between the various actions taken. The same strategy is used in this PhD work (see section 3.4.5).

Mixed methods can be used within disciplines, but also to bridge two or several disciplines. The main point here, is that both qualitative and quantitative methods are used to investigate a specific problem (Johnson, Onwuegbuzie, and Turner 2007). The

mangrove example above exemplifies how mixed methods can help to merge research from sustainability science (itself a multi-disciplinary field) and the social sciences.

### **3.4.2 Apparent weaknesses of mixed methods**

Every research project should choose its research methods carefully, based on the research question. Mixed methods, with its smorgasbord of available research methods, offers a plenitude of solutions to methodological struggles. However, it has its own set of weaknesses, flip sides of its versatility and encompassment. Since they may impede on a research project, two major weaknesses of mixed methods are presented and discussed in the following.

### **3.4.3 Team collaboration versus the lone researcher**

Creswell and Garrett bring up several problems associated with the use of mixed methods, such as the difficulty of working together in a cross-disciplinary team (Creswell and Garrett 2008). Since this is a “*lone researcher*” thesis (Creswell and Garrett 2008), that particular problem is not of immediate concern, but working cross-disciplinarily still poses problems: a project that spans both the quantitative and the qualitative is difficult to find a home for in the literature and among scholars. As a consequence, it is also research that may be difficult for a journal to review, to present at a conference and to attract funding for.

### **3.4.4 Lack of unified language**

Another drawback is the lack of a unified language. In Creswell and Garrett’s words, the mixed methods approach is “*bilingual*”, “*neither quantitative nor qualitative*” (Creswell and Garrett 2008). Similarly, the home of the discourse associated with mixed methods is not obvious. Whereas there have been some efforts to form a methodologically stringent research practice based on mixed methods, the fact remains

that it is a meta-paradigm with a great amount of elasticity. Thus, not surprisingly, the greatest strength of mixed methods as a methodological paradigm is its greatest flaw: where anything goes, stringency suffers.

In his enlightening essay, Johnson (2017) aims to solve this dilemma by settling on some ideals that “*should cut across all people*” – and by “people”, he means researchers. One such element is ethics, and, as is evident from the following, he is not only referring to research ethics.

Elevating mixed methods to a meta-research-paradigm which he refers to as “*dialectical pluralism*”, Johnson proposes a number of principles required to establish the needed “*intellectual process*” of this new paradigm, including these:

- *“Pay careful and deep attention to multiple sides of issues.*
- *Use respectful dialog with multiple perspectives and mental models.*
- *Rely on multiple methods.*
- *Attempt to produce negotiated, thoughtful positional adjustment and [a] win-win solution.*  
...
- *Strive for social betterment and social justice.*  
...
- *Use syncretism (attempting to reconcile or produce a new union of different and opposing principles and practices).*
- *Use synechism (antidualism) which is the philosophical doctrine rejecting dualisms and stressing continua.*  
...
- *Strive for dialectical integration and intellectual growth (thesis, antithesis, synthesis, ... continually).*



- *Connect theory and practice; produce practical theory.*
- *Interrelate or integrate into workable wholes ...” (Johnson 2017)*

In these bullet points, several patterns can be discerned. First, there is a humility and recognition of “the other”, that is, the fellow researcher or the fellow researcher’s method preference. Second, this humility stretches to the societies/groups that are under study and to the explicit desire to better the world. Third, there is an emphasis of interweaving findings from the qualitative and the quantitative, and the use of a multitude of methods is stressed. Fourth, as in systems engineering, an unspoken desire of dialectal pluralism is to build a whole, or at the very least to establish parts of a system in which the whole is greater than the sum of its parts. Fifth, the concept seems designed for solving today’s “big problems”, say, climate change. As Johnson somewhat amusingly puts it: one of the principles is to connect theory and practice into “*practical theory*”.

If there is anything to object to in Johnson’s set of principles, it may be the “dialog” of the quantitative and the qualitative. A view based on the opposing forces of two paradigms could come across as stagnant and anything but the synthesis he wishes for. However, the “dialog” (or, to move beyond Plato, dialectism), that Johnson envisions would be non-binary. Johnson refers to this dialog as “*multidimensional*”.

Some other researchers choose to recognise the plurality of methods – regardless of what “camp” they belong to. In fact, this recognition goes further than to Johnson’s lexicon: for better or worse, Thorén and Persson (2013) prefer to speak of “*methodological pluralism*” rather than use the term “mixed methods”. The future will have to decide, but hopefully we will be able to move away from the qualitative and

quantitative, as cementing their dichotomy may prove counterproductive within this promising methodological paradigm.

Leaving the above discussion aside, also Dencer-Brown et al. (2021), whose mangrove research is cited above, stress the equal worth of the quantitative and the qualitative. Johnson, in turn, notes that *“dialectal pluralism fits perfectly with equal-status mixed research”*. Analogous to Johnson, Dencer-Brown et al. apparently believe in addressing a research enquiry from different angles to form a truth more complete than what individual studies can: *“Using an equal-priority mixed methods approach opened the researchers up to new ways of knowing and observing the ecosystems and communities involved as a complex whole”* (Dencer-Brown et al. 2021).

### **3.4.5 Mixed methods in the thesis**

Grounded in computer science, the writings involve technology, environment and people. More specifically, the work has ties to engineering, sustainability science and the social sciences, as shown in Figure 6. Positioning the key terms of the thesis on a triangular landscape gives an indication that although to smaller or larger degrees, all publications (including this thesis) belong to these three fields. Some insights in all three fields are essential for this work to be undertaken, and the methodology must reflect this mix.

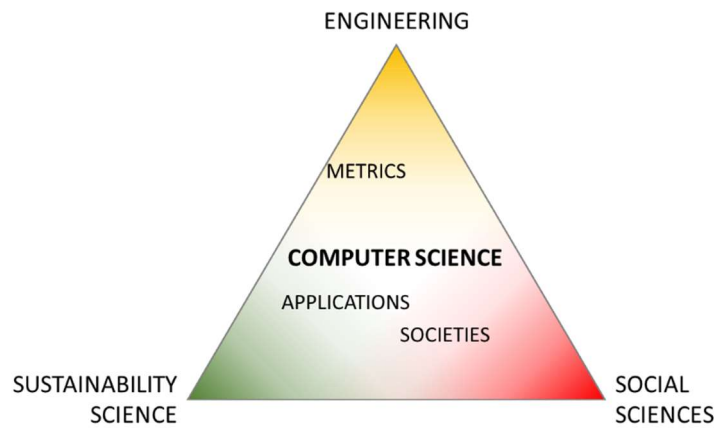


Figure 6. The thesis, rooted in computer science but related to engineering, sustainability science and the social sciences.

As the thesis involves factors spanning a wide range of disciplines (see above), integrating quantitative and qualitative findings is key to answering the research questions (found in section 1.4). The thesis follows two tracks, both of which integrate quantitative methods such as computation or statistics and qualitative methods such as interviews and observation. For each track, the activities point toward the next (and, to some extent, to activities belonging to the other track). Finally, the two tracks come together, and general conclusions are drawn based on this synthesis. Figure 7 shows this progression, visualised as the two parallel tracks.

The top track in Figure 7 looks at new data centre energy efficiency metrics, required to correctly reflect improvements done in data centre energy efficiency. Combining quantitative and qualitative methods, but emphasising the quantitative, these studies belong mostly to the natural science side of computer science. The lower track, in turn, investigates opportunities and hindrances for implementation of sustainable data centres in different countries and cultural settings. Also here, research methods from both sides are essential.

A paper published for this thesis (Terenius, Garraghan, and Harper 2023) incorporates findings from both tracks. In so doing, it marries the dominantly quantitative and the dominantly qualitative, and makes the thesis as a whole a true mixed methods-endeavour.

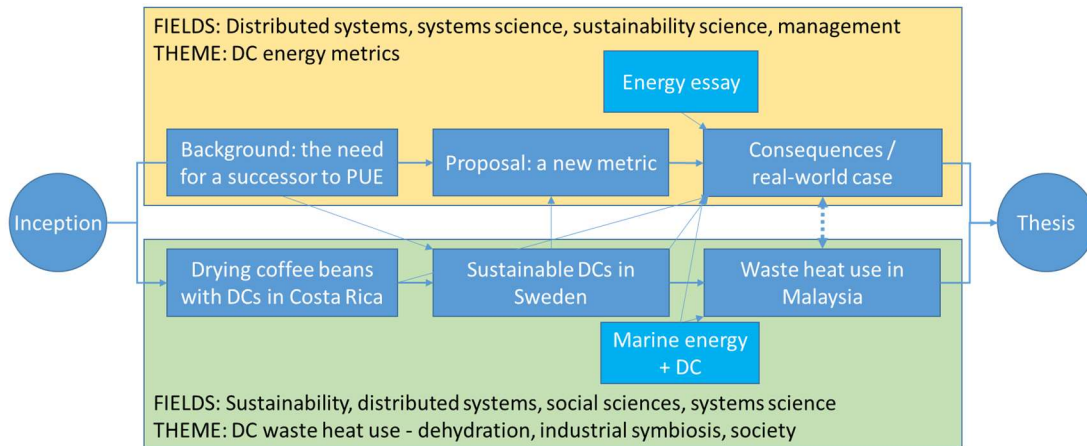


Figure 7. The two tracks of the thesis. The top track investigates data centre energy efficiency metrics and the lower track sustainability-related matters for data centre implementation in different regions. (“DC” = data centre.)

As observed in section 1.1, the research journey has been explorative in nature, guided by questions such as “What if...?” and “Why can’t...?”. Nonetheless, the choice to use an exploratory research has not meant research endeavours have been undertaken haphazardly or that findings and interests have changed much over time. Rather, the themes, enquiries and assumptions have remained fairly consistent throughout the scope of the undertaking, following the figure above.

This work also touches upon geography, anthropology and ethnography studies. As shown by Randall, Harper and Rouncefield (2007), through observation of people and societies in a certain context, such studies finds “*the glue and the glitches*” not found when only investigating very specific issues. Here, a cornerstone is visual observation. Generally, visual observation can provide the researcher with a “*sense of place*” (Wit 2003). A sense of place is, in turn, “*central to place identity and therefore*

*an important consideration for scholars who hope to grapple effectively with regional character*” (Wit 2003). For the argumentation of the thesis, visual observation has been important to understand many things that office work alone cannot easily convey:

- How hot does tropical heat really feel: what does it do to the body, and how does the heat affect the local population in general? And what does it do to infrastructure?
- Beyond tourist photographs, what does a city of four million people in a mid-income country look like? What implications does its size have for infrastructure?
- How easy is it to implement ideas in different societal contexts?
- How big is a 20MW data centre? What does it *sound* like? And *smell* like? Actual footprint aside, is it advisable to incorporate such a building (or set of buildings) in a residential area?

Such questions require visual observation, which is why this research method has been an integral part of the PhD work. Moreover, going places also reveals many things that would be very difficult to find from one’s desktop. For example, during this work, an entire outdoor Bitcoin-mining data centre was found in an isolated, snowy, region.

Visual observation has served the PhD work especially well, as it complemented its engineering-focused explorations (cf. section 5.6.3).

### 3.4.6 All connected

Computer science, at its core, is linked to the natural sciences, not least math (for programming concepts) and engineering disciplines (for the mechanics of hardware and for the understanding of processor-intensive programming and its consequences). The social sciences, on the other side, start and end with people, and uses concepts such as culture, emotion and society.

Seen as a both physical and social construct, a data centre involves both engineering related to servers and their performance – such as fluid mechanics and algorithmics – and people – hardware specialists, security guards, managers, investors, municipality officials and others. Indeed, if, as Zola claimed, Les Halles should be viewed as “*le ventre de Paris*” – the stomach of Paris – the connected data centres of today are the brain of our world, using processing and storage to serve its global customers, and terrestrial and subsea cables, satellites and local wired and wireless solutions to reach and connect them. However, whilst Zola wrote about a city, the data centre network spans the entire Earth. This fact brings even more complexity to the thesis, the concept of locality. To understand the peculiarities concerning a) the physical data centre structures, b) the people involved and not least c) the questions of locality pertaining to commodity production, dehydration needs and potential for increased ICT infrastructure, the work undertaken would arguably be a good match for mixed methods research.

From the argumentation in the above sections, it should be clear that in large and cross-disciplinary research projects, a systems science perspective *makes sense*. These preceding two words have a double meaning: a) It is usually wise to view a problem

with a wider angle, so it is sensible doing so, and b) a bigger picture provides additional insights, thus gathers additional information.

As a discipline, systems science, with branches such as cybernetics and information theory, grew in popularity after World War II (Lin et al. 2012). However, a strong focus on specialised knowledge in academia and a vast increase of required expert knowledge to be successful in natural science and engineering later led to a decline for systems science as an own discipline (Lin et al. 2012; Klir 1991; Bailey 1994). At the same time, as Fuenmayor (2001) notes, systems science has shown good potential for addressing sustainability problems, due to its holistic nature. Climate change and associated problems are the most difficult and severe challenges humanity faces today, and the problems are likely to expand in severity. As these are wicked and holistic problems, it follows that understanding systems may be helpful in solving them, and that by incorporating elements from the social sciences as well as from systems science and engineering, material social studies can provide a way to solve the riddle of data centre energetics.

### **3.5 Contributions and associated methods**

This section accounts for both the inner and the outer journeys of the PhD work. For each activity, the section briefly discusses why and how it was undertaken, with the hope of situating this multi-faceted research within a reasonably logical and methodical sequence of activities.

To show the significance of the activities carried out, the analytical framework (see Figure 11) is referenced throughout this chapter. The formats [AFx], [AFxx] and [AFxxx] are used. For example, “[AF32]” points to infrastructure access, since, in the figure, “Geographical prerequisites” is topic 3, and “infrastructure access” subtopic 2.

In addition, the exact same structure and numbering are used for section headings in chapter 5 and in the analysis of the three cases (see section 9.1). Therefore, the reader can find the reference from this chapter in the analytical framework, read more on the subject in chapter 5, and view the discussion in the analysis of the three cases in section 9.1.

### **3.5.1 Literature review**

The work commenced in October 2019 with an initial literature review, using Barroso et al.'s *The Datacenter as a Computer* (Barroso, Hölzle, and Ranganathan 2019) as a point of departure. Having started with insights in data centre problematics – not least in relation to enquiries of energy [AF11, AF2] – the literature review grew throughout the PhD work to include more books, papers and grey literature such as industry reports. The topics of these works related to data centres as facilities [AF2, AF42], to energy [AF1], and to data centres seen within societal contexts [AF61]. In December 2019 and January 2020, the investigation expanded to the very concept of energy [AF1], a topic central to this thesis. Later, other themes followed, so that the literature review came to relate to data centres, energy, metrics and metrology, commodity production, drivers for commercialisation, societal progression, environmental sustainability and academic fields. In addition, newsletters from industry associations, media and data centre infrastructure specialists have been monitored. This reading deepened the understanding of the latest trends within the computer science, engineering, societal, environmental and managerial dimensions of the work, and provided insights in today's and tomorrow's possibilities and barriers for successful data centre energy improvements in specific societal and environmental contexts.






Figure 8. The analytical framework of the thesis.

The literature review has been conducted using Lancaster University's OneSearch article search feature. This is a meta-search engine, exploring all of the databases that the university subscribes to. Additional searches have been carried out using Google Scholar, to locate also sources not indexed at Lancaster University. In addition, data have been gathered through subscriptions to news channels and research institutes (e.g., Reuters Sustainable Switch, Uptime Institute, and DCD Sustainability), and then following leads from those media. About a thousand of the found sources have been regarded for inclusion in the thesis. After the initial data collection, the sources deemed of high relevance (ca 530) have been entered in the reference system used for the thesis, EndNote. Of these sources, a little more than 300 have been used in the actual thesis.

Table 1 shows the distribution of the sources added to EndNote. These have been divided into five coarse categories. For each category, the number of sources is accounted for, as well as the category of each source: Academic work, defined as journal articles, books, book sections, PhD theses, conference papers and conference proceedings, are labelled "A". Reports (typically from bodies such as the UN or the IEA, or from trade organisations and similar industry associations such as the International Coffee Organization) are labelled "R". Other material, such as news items, governmental documents, patents and datasets, are labelled "O".

The exact distribution figures should be read with some caution: since the PhD work is multi-disciplinary, many of the sources found relate to several of the categories used for classification. For example, a conference paper called "Multiple use of deep-sea water: Industrial cooling, ocean thermal energy conversion and blue economy

development in La Réunion island” refers both to the *Energy* and to the *Policy, infrastructure and development* categories, but is only classified in the latter category.

The table also indicates the chapter(s) in which each category has mainly been of use. Finally, it provides examples of keywords used for the searches. A matter which has made many searches more time-consuming is the spelling variations of the term “data centre”. Thus, in exact searches, three terms have been used: “data centre[s]”, “data center[s]” and “datacenter[s]”. For searches in Swedish, two terms had to be employed: “datacenter” and “serverhall” (including “datacentret”, “datacentren”, “serverhallen”, “serverhallar” and “serverhallarna” for exact searches).

Table 1. Literature theme distribution among chapters. The list also shows examples of keywords used for the searches within the context of the corresponding chapter.

| Theme                                  | Number of sources        | Primarily used in | Examples of keywords used  |
|--|--------------------------|-------------------|--|
| Commodities, Dehydration               | 50<br>A:30; R:13; O:7    | chapter 6         | coffee bean drying, commodity dehydration                                  |
| Data centres, Metrics                  | 198<br>A:102; R:23; O:73 | chapters 5 and 8  | data centre waste heat, sustainable data centres, PUE, energy efficiency   |
| Energy                                 | 87<br>A:48; R:13; O:26   | chapter 4 and 5   | renewable energy, energy mix   |
| Policy, infrastructure and development | 135<br>A:72; R:22; O:41  | chapters 2 and 7  | Africa, Sweden, Malaysia, Costa Rica, climate change, industrial symbiosis |
| Systems science, Methodology           | 60<br>A:60; R:0; O:0     | chapter 3         | wicked problems, mixed methods, systems science, systems engineering       |

The sources’ distribution is an indicator of the status of the topics under investigation:

- Comparatively many sources relating to *Data centres*, *Metrics* are from other fields than academic work. This is because the industry expands and changes so quickly that academia has difficulty keeping up, especially concerning data centres and societal integration.
- *Policy, infrastructure and development* features a large number of reports. These concern the competitiveness of nations, ICT coverage, energy use of different regions, and so forth. Its Other category is also substantial, as it includes many news items relating to Malaysia, Sweden and Costa Rica.
- *Systems science, Methodology* features only academic work. This seems logical, as the development of academic disciplines is mostly a matter of concern to academia itself.

### **3.5.2 Investigating uses for data centre waste heat**

As will be discussed in chapter 4, when processors turn electrical power to heat, the same amount of energy is still present after use, though in a different and typically less desirable state. Thus, uses for this lukewarm heat [AF51] began, in early 2020.

In cold regions, where cooling is free and waste heat can be reused to heat buildings, the environmental footprint of data centres does not have to be high (see section 7.5). To view data centres in cold climates, visits were made to a few RISE research facilities Luleå in northern Sweden, 25-28 October 2020, to EcoDataCenter in Falun in mid-Sweden, 23 August 2021, and to a residential building featuring a small-scale data centre in Basel, Switzerland on 3 March 2023. In Luleå, a meeting was held with researchers from RISE and from Luleå Technical University, and a subsequent tour to a datacenter in the nearby town Boden was carried out. At EcoDataCenter, a

meeting was held with a few members of the marketing and operating teams, and a walkthrough of the facility followed. In Basel, the chief technology officer of a company called DeepSquare and a representative for the energy company Industrielle Werke showed the setup.

Still, cold regions are not where most of the world's future population will grow. In Africa and Southeast Asia, populations are expected to almost double by 2040 (Rosling, Rosling, and Rosling-Rönnlund 2018). Rapid financial development in these regions and new technologies such as 5G will require more data centres in the relative vicinity. This is because however we innovate data transmission networks, transmission is ultimately limited by network latency from Internet exchange points and the speed of light.

The quest for uses of data centre waste heat in warm regions became an introduction to studies of commodities in need of dehydration, as many commodities can be dried at low temperatures (see section 6.2). In fact, for some commodities – especially coffee beans – low-temperature drying is often preferred, to preserve nutrients and to avoid scorching. Recognising the promising use for waste heat for coffee bean drying, preconditions and applications in different nations were then explored, to find a suitable case for the idea. It was assumed that the country most suitable for a case study would be the one in which the ratio between current availability of data and the desire for data would be the smallest. After all, a high demand for ICT could be a potential driver for actual implementation of local data centres [AF622]. Based on data from the international telecom industry association GSMA, an index called ICRI, ICT availability/Consumer Readiness Index was developed (Terenius, Garraghan, and Harper 2021). The coffee bean producing country with the lowest ICRI

turned out to be Costa Rica. Costa Rica proved a promising case, as shown in section 7.4. In so doing, commodity drying with data centres would both bring income to coffee cooperatives and enable sustainable growth in terms of many of the UN SDGs.

The literature review also led to further uses for data centre waste heat [AF51]. One identified use, which incorporated both the material and the social, was the integration with an OTEC (ocean thermal energy conversion) plant (see section 6.4). Two papers have been produced on this topic (Terenius et al. 2021; Terenius, Garraghan, and Harper 2022).

Malaysia is a tropical country with substantial commodity production. The capital also hosts one of but a few research centres on OTEC technology. Fieldwork was therefore undertaken in Malaysia, 30 April-2 June 2022. Kuala Lumpur, the home of the OTEC Research Centre at Universiti Teknologi Malaysia (UTM), became the home base for this exploration, to view data centre waste heat in a large tropical city context [AF 61] (see section 7.3). An experimental OTEC plant was visited in Port Dickson, a small town on the Western coast of peninsular Malaysia. Finally, to view large-scale tropical commodity production in action, the Cameron Highlands was visited as it is home to several large tea plantations. As a memory aid from these travels, a diary was kept and entries added daily. In addition, a large amount of photos were taken, a few of which are shown in this thesis.

One experiment relating to the core sciences was carried out. It concerned the drying features of waterweed, in July 2021. The rationale for the experiment was to evaluate drying times from fully wet to dry weight, as an exponential decrease in drying capacity from ambient drying would increase the potential for data centre waste heat

drying [AF51]. This could, in turn, possibly enable less expensive production of biofuel. How this experiment was carried out is accounted for in section 6.3.1.

The PhD work has also meant taking active part in many industry hearings – mainly webinars – on data centre sustainability, data centre future design and OTEC technology prospects, to learn about the latest trends and events in the related industries (such as the new regulations in Germany and of the EU – see section 8.4). Such webinars have mainly been hosted by the Open Compute Project Foundation, Datacenter Dynamics, and the Uptime Institute.

### **3.5.3 Cultural factors**

In autumn 2021, work commenced on success factors for the Swedish data centre industry, with the underlying assumption that cultural factors [AF63] was a contributing, but neglected, factor for data centre establishment. The rationale for this interest was twofold:

1. On the research travels to Swedish data centres, employees repeatedly touched upon the significance of nature. From this, a hypothesis was formed that the reason why the sustainable data centre industry is strong in Sweden, Finland and Norway is not only because of geography. For instance, *allemansrätten* or “every man’s right”, is a century-old law giving the ability to roam and even to camp in every field and forest, as long as one is careful not to not disturb nature. Providing a better understanding of cultural success factors (e.g. (Porter 1990; Yeh-Yun Lin and Edvinsson 2011)) may, in other words, facilitate further deployments of sustainable data centres in the region.
2. Compared to the Nordics, tropical societies would be at disadvantage because of the much hotter climate [AF31]. Moving data centre problematics beyond

climate might reveal other potential success factors. Thus, examining what worked in Sweden and other Nordic countries may provide insights on what works in other contexts, thereby helping the establishment of sustainable data centres in warmer regions globally. This is especially important for low- and mid-income countries, which lack the financial resources of today's leading data centre nations.

The case, founded on industry reports and on Michael Porter's seminal work on nations' success factors (Porter 1990), is partially included in the thesis (see section 7.5).

A semi-structured interview (Appendix C), relating to cultural factors in Malaysia, was conducted with Dr Azzman Shariffadeen at the UTM Hotel & Residence in Kuala Lumpur on 7 May 2022. The two-hour interview was recorded and the recording stored in a designated folder on the Lancaster University servers, in line with the approval for conducting interviews from the Faculty Ethics Committee (FST21010). The interview sheet is appended to this thesis. The notes taken from the interview carried insights that are used in section 7.3.

### **3.5.4 Metrics**

As the thesis circles data centre energetics, metrics for evaluating data centre sustainability [AF 52] were investigated already in autumn 2019. It was concluded that with current data centre metrics, the claims of the thesis could not be evaluated. Hence, in November 2019, work began on replacement metrics [AF414, AF424, AF52]. Replacements for the leading metric used for data centre sustainability, *Power Usage Effectiveness* or *PUE* for short, have been researched for over a decade. The work, based on the small body of metrology literature concerning data centre energy, resulted in a



new metric, presented in chapter 8 and published in January 2023 (Terenius, Garraghan, and Harper 2023).

### 3.6 Chapter conclusion

In 1966, historian Lewis Mumford coined the term *megamachine* to depict how a system, no matter how complex, was a result of cooperative work between humans (Mumford 1966). As examples, Mumford used monarchies and the Manhattan Project. According to Crawford, artificial intelligence

*“...is another kind of megamachine, a set of technological approaches that depend on industrial infrastructures, supply chains, and human labor that stretch around the globe but are kept opaque. We have seen how AI is much more than databases and algorithms, machine learning models and linear algebra. It is metamorphic: relying on manufacturing, transportation, and physical work; data centers and the undersea cables that trace lines between the continents; personal devices and their raw components; transmission signals passing through the air; datasets produced by scraping the internet; and continual computational cycles.”* (Crawford 2021)

Crawford adds: *“These all come at a cost.”* This cost is something the thesis tries to mitigate.

This chapter has discussed how some topics of enquiry, such as societal views of data centres, can be approached from two directions: the material and the social. As it happens, there actually exists a field with the ambition to do the same. Vonderau frames her research within *infrastructure studies* (Vonderau 2019), where mainly social science researchers examine material subjects of concern. There are a few reasons for leaving their path here, its non-equal relationship (as infrastructure studies is rooted in the social sciences), and its redefined and limited scope (infrastructure).

The material social perspective solves these issues. First, it is designed for equal contribution. Second, the scope is not limited to infrastructure; a general view, relevant to many topics, can be envisioned. With such a view, the material social perspective should be applicable to anything from microplastics to lithium mining, matters only loosely coupled with infrastructure. And incidentally, one of the few existing papers within the material social paradigm uses current social and ecological issues of lithium mining as a starting point for a discussion on new compounds for battery packs (Murdock, Toghil, and Tapia-Ruiz 2021).

Infrastructure studies have a counterpoint in the *industrial symbiosis* field (discussed in section 4.2), whose aim is to be “*literally squeezing more value from the same initial inputs through co-located manufacturing processes*” (Gregson et al. 2015). Here, people are not so visible; the anthropological aspect is less prominent, and a typical task for the field is to investigate “*exchanges of by-products and wastes in planned complexes of co-located manufacturing plants*” (Gregson et al. 2015). At first glance, this thesis is not that far off, especially in that the industrial symbiosis field is quite solution-oriented. However, the field is too narrow. Material social perspective studies seems to constitute middle-ground between the infrastructure studies and industrial symbiosis literatures. Further, through expanding on the limited scopes, its degree of usefulness would be higher.

Recently, a Google researcher lamented the passive voice of *ethnographic* research, where one tries to “*wholeheartedly strive to perfect the semiotics of representation, forgetting that the forms we produce are not an end in themselves*” (Mendonca 2022). In his paper, he hoped that infrastructure studies could help propel ethnography forward, since it “*makes an explicit connection between the material,*

*historical and semiotic dimensions of contextual investigations, thereby broadening the scope of ethnography from developing insights to driving systematic change”* (Mendonca 2022).

Still, people- and culture-centred research such as ethnography is indeed helpful, not least to glue disparate findings together in order to confirm suspicions and prevent biases and presuppositions to linger. Therefore, its research methods – not least user observation and visual observation – can support the development of new findings also in material science research. Blackwell means that “*concepts such as attention, expectation, likelihood, evidence, and observation*” can indeed aid “*the imagination for engineers, entrepreneurs, managers, policymakers, and other actors involved in the social construction of information processing systems*” (Blackwell 2021). Blackwell is correct: an axiom in the human-computer interaction and cognition fields is that understanding the end-user is a key to successful design. Feeding this understanding to the developers makes for better designed artefacts – be they computer software, everyday things or infrastructural projects of great magnitude. To add to Blackwell’s statement, an “*expectation*” is something that ethnography can confirm, disprove or prove uncertain. The power of doing so should not be underestimated: this is one of the strengths of studies of geography, anthropology and similar fields – to study a person, a group or a society and draw conclusions based on intangibles that would not have been considered through reading alone.

The final point to make in this chapter concerns the two views of infrastructure. To Star, infrastructure is the glue that ties other objects – a railroad connecting two cities, or a telephone line connecting two persons speaking. The same is true for Barroso et al.: a data centre is the building that hosts servers. However, for the anthropologists,

the distinction between servers and the surrounding building is less clear-cut. For instance, Figure 25 shows how a data centre has two separate inputs – electricity and data – and two outputs – heat energy and requested data. Hence, there is an intrinsic relationship between the building and the data it contains: they only exist in a symbiotic relationship. Thus, from a material social perspective, they are chimeras, sometimes seen as two entities, sometimes as one. This double view is founded in the idea of objects-within-objects; the fact that a data centre can contain servers is not any stranger than that a car can contain (or “has”) wheels. This is just one instance where systems science principles come to aid, and necessary to view the data centre as a piece of the societal puzzle.

## 4 Data centre energetics

*So what is a battery anyway?* Ultimately, a battery is a promise. It promises that somehow, and sometime over the course of the existence of the universe, it will release whatever energy it contains. In this perspective, our sun is the largest and most important battery in our corner of the universe. A serendipitous gravitational twist in space let matter coalesce under increased pressure, to finally generate pressure and heat high enough to enable fusion and become the sun as we know it, just like the hundred billion other stars in our galaxy, in turn, one of hundreds of billions of galaxies in the known universe.

Every second, the sun transforms some 600 million tons of hydrogen to helium, and in the process, it releases roughly 1% as energy (Miramonti 2009). After five billion years, this giant battery – our sun – is now half-way through its reserve. A fraction of the electromagnetic radiation it emits is received by Earth. And a tiny fraction of that energy is collected by solar panels and indirectly hydro power and wind power and other renewables. And of *that* energy, a few thousandths help powering today's data centres.

It is not only the proportions that boggle the mind, but also, to the layman at least, energy itself, this simultaneously obvious and mysterious thing, measured in terms of joules, kilowatt-hours, British thermal units and even BOE, Barrel of Oil Equivalent. Putting either of these units after a digit means to us very little, especially since there is always the confusion between energy and power.

The reason for indulging in this rhetorical play is that the thesis does not view energy as a unit so much as a *quality*. We may speak of kilowatt-hours, or think of how much energy we need to heat our homes, but from the perspective of the thesis, energy is not really the sum of work performed. Rather, it is *the ability to perform work* (U.S. Energy Information Administration 2021). And so, when a data centre supplies heat energy to a dehydration facility, what is in fact transferred is the ability to perform work. An understanding of this quality makes thinking about energy transformation much easier: when the sun burns hydrogen, it throws out one percent of that hydrogen's stored capacity to do something. This capacity – or ability – is then sent through space, and some of it reaches us.

So, the ability to perform work – or the “promise” – is paid forward. In a manner, energy can thus be compared with money. As Smil (2017) states, energy is truly “*the only universal currency*”. To follow on this analogy, paper money cannot do much (more than produce heat when burned), but the value paper money carry can do a lot. Moving forward, energy within a society is viewed as a resource – again, similar to money – that can be used in many contexts.

## **4.1 Defining clean energy**

As stated in the first law of thermodynamics, energy can neither be created nor destroyed; it can only be more or less concentrated. It can come in several forms, such

as nuclear, heat or stored electricity, and whenever a transition from one form to another occurs, heat is released. Specifically, the thesis investigates ways to keep energy focused and useful rather than dispersed into the atmosphere through discarded heat energy. But from a sustainability perspective, even more important is the source of the energy used – the matter of “clean” energy.

“Clean energy” is a term that has received substantial focus in current years, due to climate discussions. Regrettably, but also understandably, climate policy is subjected to stark influence from other policy areas, economics not the least, but also policy regarding national security (Busby and Busby 2007), segregation (Kortetmäki and Järvelä 2021), gender (Huyer et al. 2020) and many other areas. With a society aligned to the Keynesian economics vision of everlasting financial growth, and with the problems that would face nations with collapsing economies, it is not strange that institutions such as the European Parliament (Alderman and Pronczuk 2022) debate what energy can be classified as *clean*. This question is of importance to the data centre industry, as its energy demand is high and its sustainability grading will, logically, be based on such policy.

Every energy source known to humanity can be subjected to questions regarding sustainability. In the following, the energy sources of main interest to the data centre industry are discussed.

### *Coal and oil*

These are both fossil fuels and without a doubt huge threats to humanity due to their detrimental effects on climate: when burned, they result in considerable carbon dioxide (CO<sub>2</sub>) emissions (Ritchie 2020). That said, these are inexpensive energy sources. With carbon capture and improved filtering, it is not impossible (though far from likely) that

carbon and coal could play a role in tomorrow's sustainable energy landscape. For the time being though, it is crucial that fossil fuel "stays in the ground".

### *Fossil (natural) gas*

This is also a fossil fuel, but it does not produce as much greenhouse gas (GHG) emissions when burned. Thus, fossil gas has been promoted as clean energy, or at least as a stepping stone from coal and oil toward cleaner energy sources. The main disadvantage of fossil gas is that during a dozen years after its release to the atmosphere, the main constituent of the gas itself – methane – is about eighty times (IPCC 2013) more severe a greenhouse gas than CO<sub>2</sub>. During extraction, there is typically some leakage of the gas into the atmosphere. As a result, though touted as clean energy by some of the European Parliament (Alderman and Pronczuk 2022), fossil gas is not much better for the environment than coal and oil (Ritchie 2020).

### *Nuclear energy*

During operation, nuclear energy is almost without GHG emissions (Ritchie 2020). Thus, in that sense, it is better than coal, oil and fossil gas, these unquestionably dirty energy sources. However, it is an expensive technology with obvious safety concerns, and so, often under scrutiny. With its own problem space incorporating security, financial and environmental discussions, nuclear energy is viewed differently even in otherwise similar countries, such as those of the Nordics. While Denmark does not have any nuclear energy, Sweden was dismantling its plants up until the invasion of Ukraine and Finland is building new ones. All of these nations are high users of energy. Sweden and Finland both have energy-intensive industries such as paper mills and mining, so a reliance on nuclear energy (and hydropower), and energy policy overall, have many implications for these societies. The last few years have seen renewed interest in nuclear



energy (Värri and Seppälä 2019). Research funding is granted based on the insight that available energy resources cannot sustain high-income countries on their current consumption levels, and the objective is to make it long-term viable, that is, as one report phrases it, to “*save nuclear from itself*” (Värri and Seppälä 2019).

#### *Wind power and solar power*

In terms of CO<sub>2</sub> emissions, wind power and solar power fare much, much better (Ritchie 2020). Through research and scaling up, the price of wind power and solar power has come down considerably during the last decade, and finally made them commercially viable technologies. It has been shown that also with the embodied GHG footprint of fabrication included, wind power is one of the least GHG-polluting of today’s available options (Smoucha et al. 2016). However, no matter how inexpensive these technologies could be, and no matter how environmentally benign wind turbines and solar panels could be manufactured and disposed of, they still face a natural limitation: wind power is dependent on wind, and solar power only works during the daytime. Thus, a society cannot depend on wind power and solar power alone, at least not without considerable energy backup available – and the backup solution may provide its own sustainability challenges.

#### *Hydropower*

But for gases leaking up from submerged lands (Deemer et al. 2016), hydropower is generally regarded an efficient source of power with low emissions of greenhouse gases (Ritchie 2020). Moreover, since the water masses of the hydropower dam stores energy until needed, this is “baseload” technology, meaning it works without interruptions, regardless of weather or time of day. Whilst constant and flexible energy access is a very attractive quality, hydropower is not suitable everywhere. For example,

hydropower cannot be efficiently used in the UK, due to the relatively few lakes and low altitudes of Great Britain (Berners-Lee 2021) and the great need for electricity because of a large population. Moreover, though less than ten percent of available water sources have yet been exploited worldwide (Hoes et al. 2017), centuries of construction mean many obvious locations for hydropower are already used. Therefore, besides likely having less value than the hydropower installations already put to use, further exploitation may come with increased environmental or societal problems.

### *Biomass*

As discussed in section 4.2, biomass contains large amounts of energy, and reclaiming it can be one of the keys to a circular economy. A downside of biomass is that (depending on the source and combustion technique) it too may be a rather large greenhouse gas emitter (Ritchie 2020). Thus, environmentally concerned biomass users would have to limit their choices of produce.

### *Marine energy*

Other sources of energy are being researched. Some of the more promising are marine energy, such as wave, tidal and ocean thermal energy. Of these, *ocean thermal energy conversion* (OTEC) technology would be an exciting option for the data centre industry and for sustainable data centre integration within communities, because of its multitude of outputs (see section 6.4), because it feeds on low-temperature heat energy and because it is baseload (as opposed to wave power and sea-based wind power). But for construction and dismantling, marine energy technologies all come with low or zero greenhouse gas emissions. At this moment, however, save for sea-based wind power, none of these technologies has been deployed on a large scale.

As noted above, several of the energy sources discussed – renewables especially – require backup power stored in one way or another. Backup achieved through huge stacks of lithium batteries comes with environmental problems: the lithium may have been extracted in places where mining has had negative societal or environmental consequences (Bonds and Downey 2012; Agusdinata et al. 2018; Murdock, Toghil, and Tapia-Ruiz 2021). Lithium batteries also need cobalt, leading to suffering and exploitation in, mainly, the Democratic Republic of Congo (Zeuner 2018; Nkulu et al. 2018). Further, lithium batteries are expensive and may be better needed elsewhere, not least in electric vehicles, to help the transition away from fossil fuel.

Thus, for society as a whole, investing in large lithium battery packs as a power backup for a data centre is not an optimal solution. Sodium-based batteries fare better, as sodium is an abundant element on land and in oceans. However, the time for such batteries is yet to come, as the chemistry and electrochemistry of electrode materials differ between lithium-ion and sodium-ion batteries (Abraham 2020). So, though it is believed that sodium-ion batteries *“have great potential to be the most effective and sustainable energy storage solution in the future”* (Bai and Song 2023), they are just coming out of the development phase. Within data centres, lead-acid batteries are still common (e.g. (Zabeu 2023)), but as a whole, it seems the future of stationary batteries will be sodium-rich.

There are also other options for energy storage. One is pumped hydro, meaning water is pumped to reservoirs with excessive power in the daytime, and released through a hydropower dam at night. Similarly, large blocks of stone can be lifted or lowered depending on whether energy should be stored or released (Reynolds 2022), again using gravitation and turbines to store and release energy. Producing hydrogen gas from

water, and then converting the hydrogen gas to electricity, is another option, and this technology has promise for the future energy mix (Hosseini and Wahid 2020). However, regardless of storage technology used, they all come with transportation and conversion losses, aspects that should be considered when drawing plans for community infrastructure.

Evidently, it is not easy to draw the line between sustainable and unsustainable energy. As it seems, it is even difficult to put energy sources on a “benevolence” scale, since they all come with a variety of problems: whereas some are indeed high-carbon emitters (coal, oil, fossil gas), others may not be yet common enough (marine energy), perceived as safe enough (nuclear), efficient enough for electricity production (biomass) or scalable enough (hydropower). To complicate matters further, many of the technologies require energy storage solutions. As mentioned in section 5.6.1, “clean energy” is not only a scientifically but also a politically bound concept. What route to take for the data centre industry or for societies at large is, at this point, uncertain.

## **4.2 Use and reuse within the circular economy**

Waste heat use is not new. It is often used locally, to heat nearby offices or the local swimming pool through a heat exchanger, it but can also have further ties to the society infrastructure.

The use of waste heat and other industrial waste products (such as water) is fundamental to the *circular economy*, “a cyclical ecological system which is capable of continuous reproduction of material form”, to cite Boulding’s influential essay from the late sixties on the economics of “Spaceship Earth” (Boulding 1966). Here, the aim is to reuse or recycle material to the greatest extent. In Boulding’s vision,

*“throughput is by no means a desideratum, and is indeed to be regarded as something to be minimized rather than maximized. The essential measure of the success of the economy is not production and consumption at all, but the nature, extent, quality, and complexity of the total capital stock, including in this the state of the human bodies and minds included in the system.”* (Boulding 1966)

Boulding is quite explicit in his financial view on resources, which makes sense, given the strong institutional ties between sustainability science and management, and that both economics and sustainability relate to conserving resources or spending them wisely. For this vision to come true, Boulding states that

*“what we are primarily concerned with is stock maintenance, and any technological change which results in the maintenance of a given total stock with a lessened throughput (that is, less production and consumption) is clearly a gain.”* (Boulding 1966)

Living “circularly” would have obvious environmental benefits... could it work. Perpetuum mobile systems are impossible, since they violate the second law of thermodynamics. The law states that *entropy* – in physics seen as the degree of disorder – increases for spontaneous processes. In other words, ordering of matter requires added energy, indicating that a closed, circular system cannot exist. Further, the law implies that there is no such thing as a loss-free energy conversion. Boulding correctly notes that as all material transformation requires *work* (i.e. the very thing energy – and only energy – can produce): the system simply *“cannot escape having inputs of energy”* (Boulding 1966). His suggestion to *“minimize”* throughput is also an indicator of that the circular economy is an ambition rather than a realistic goal.

As a true circular economy is thus unattainable, a wider definition of the term is more useful, in general and in this thesis. In their definition, Murray et al. (2017) emphasise benevolent and careful planning and execution of infrastructure and daily lives, and implicitly recognise that a circular economy is, in reality, not circular:

*“The Circular Economy is an economic model wherein planning, resourcing, procurement, production and reprocessing are designed and managed, as both process and output, to maximize ecosystem functioning and human well-being.”* (Murray, Skene, and Haynes 2017)

This wider definition makes sense, and spans several key concepts. Though the slightly awkward reference to a “circular” economy persists, the definition is more useful than Boulding’s, and relates well to ideas of waste heat use. Therefore, Murray’s definition is what underpins the ideal of this thesis. Figure 9 is based on the constituents of Murray’s definition.

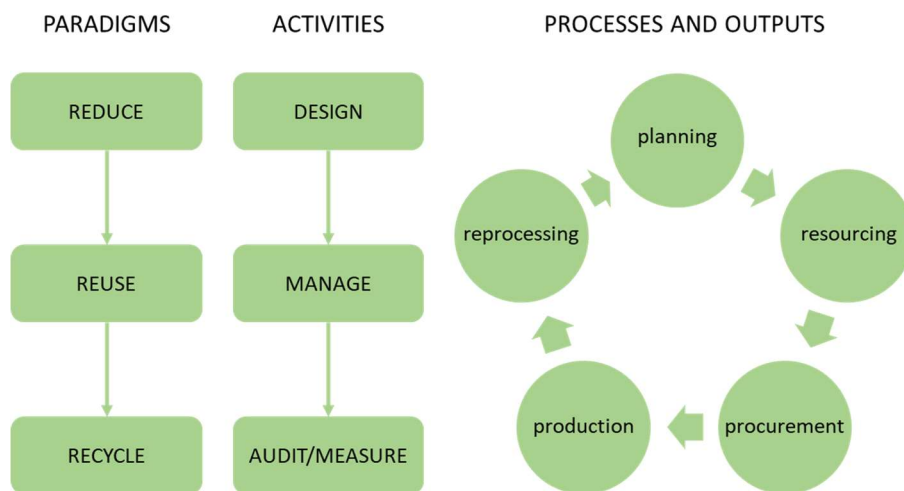


Figure 9. A systematic view of the circular economy. Here, *Activities* and *Processes and outputs* are based on Murray’s definition above. The “three Rs”, commonly used to illustrate the essence of the circular economy (Murray, Skene, and Haynes 2017), are included as foundational *Paradigms*. Boulding’s argument is expanded by adding “Audit/Measure” to Activities, as it is difficult to manage (and keep up) what is not measured (e.g. (Waas et al. 2014)).

Industrial collaboration is at the heart of the circular economy (Gregson et al. 2015). Of particular interest when discussing industrial collaboration is the sub-concept *industrial symbiosis*, where undesired output from one industry serves as input to another (Murray, Skene, and Haynes 2017).

In the particular case of waste heat, one industry's undesired output becomes another industry's resource, for dehydration, heating or other purpose. In cold countries, reusing power-intensive industries' waste heat for warmth or warm water is not a new idea (Larsson 2012). For instance, industry flue gas condensation is one of the main contributors of energy to district heating used in Swedish homes (Hwit 2019).

A “*poster child*” (Lombardi and Laybourn 2012) and a “*must-cite case*” (Gregson et al. 2015) of waste heat use within an industrial symbiosis context is the Danish multi-industrial site Kalundborg, connecting a dozen private and public enterprises, including a biomass (formerly coal) power plant, an oil refinery, a pharmaceutical plant and a plasterboard manufacturing plant (Gregson et al. 2015; Transition ApS 2021). Here, residues from one facility serves as input to another, and in total there are currently around thirty exchanges between the industries involved (Transition ApS 2021). In addition to water and ash residues being circulated throughout the site, waste heat from the power plant serves the other industries as well as the local district heating system (Deschenes and Chertow 2004). Today, one may object to labelling Kalundborg a poster child of industrial symbiosis, given the climate-threatening industries involved. Still, what has been achieved is a system with potential environmental and monetary savings for the actors involved.

In the Swedish town Borås (and others like it), a resource recovery system much in line with the idea of a circular economy has been implemented (Taherzadeh and

Richards 2019). Household waste is sorted by the locals: plastic, metal, glass, paper, batteries etc. are all disposed of in various containers and then recycled; sending waste to landfill is just about non-existent in Sweden (European Environment Agency 2021). White goods, electronics, and so on are left at one of the municipality's "recycling centrals" (återvinningscentraler). White and black plastic bags are freely provided by the municipality, white for burnable household waste and black for food stuff. Both are thrown in the garbage bin, and at the waste handling facility, a robot sorts black bags from white. Black bags are transformed into biogas and white are incinerated. The biogas is used for the municipality's buses and other vehicles, and the heat from incineration heats warm water and buildings through the district heating system. This straightforward solution saves both considerable amounts of money and the environment (Taherzadeh and Richards 2019). In so doing, the municipality of Borås takes advantage of the exchange of energy qualities: waste products replace fossil fuel and renewable energy is saved for other purposes. Thus, potato peel becomes batteries, storing energy from the soil and the sun, and releasing it as biogas. Thoughts on harvesting seaweed for energy (see section 6.2) are not that far off.

Not discrediting the merits of the Swedish waste reclamation case, it must be said that the energy derived from potato peel is likely less than the energy used when planting, harvesting and transporting the potato, so also this system "*cannot escape inputs of energy*", to quote Boulding. Further, even with the constantly added promise of work graciously supplied by the sun, the circular economy in Sweden remains a vision. This is because a high-consumption society needs an inflow also of plastics and many other materials, as it is not practically possible to fully or even partially recycle all types of items or materials.



### 4.3 Data centre energy use and waste heat

As discussed in this thesis, a data centre can come in different configurations. Servers placed on-premise in one cabinet may serve the needs of an organisation such as a hospital. Servers in the thousands reside in the hyper-scale data centres of Big Tech. Experimental data centres may be put on the ocean floor. While most servers are cooled with air, some are immersed in liquid.

The traditional data centre buildings contain one or a few floors of big halls filled with rows of server racks. The server racks are positioned facing each other, to form “*cold aisles*” and “*hot aisles*” (ASHRAE Technical Committee 2016); that is, rows where cold air is pushed into the servers, and rows where the heat is removed from the server heat exhausts.

As Barroso, Hölzle and Ranganathan (2019) argue, even though a data centre may consist of thousands of servers (ordinary or specialised computers), it can still be viewed as a single computer. Like any desktop or laptop computer, air-cooled servers need fans to cool their processors, and so does the server room as a whole. Here, computer room air conditioning (CRAC) units keep the temperatures of the servers and of the server rooms themselves in check, with the objective to remove as much heat as possible from the processors, and to keep temperature (and humidity) levels consistent in the server rooms. The CRAC units are part of a system which also includes heat exchangers (to remove the heat from the server rooms) and fans (to move the air throughout the data centre). Power distribution units provide electrical power to the servers. In an air-cooled CRAC, a refrigerant carries the heat energy outside the premises, with the help of a condenser. A water-cooled CRAC uses water to transport

the heat energy. Some water is then evaporated in a cooling tower, and the rest is returned to the data centre.

Server cooling, usually in the form of filtered air, is a crucial task for data centre managers. Cooling facilities are expensive, and they use massive amounts of energy. As this thesis will show, a modern data centre uses two thirds of the incoming power to run the processors, and one third to power the rest of the facility, mainly cooling equipment. Some power is lost in power conversion, and the remaining few percent are used for lights and such within the facility. In essence: in a 3 MW (megawatt) large-scale data centre, 2 MW are used by the processors, and 1 MW is used to shoo the 2 MW just delivered.

Pertinent to the argument of this thesis is that regardless of how it has been used in the data centre, incoming energy eventually becomes heat within the data centre structure. Today, most data centres exhaust the waste heat into the surrounding air. In colder climates, the so-called waste heat is – rightly – seen as heat energy, and sometimes used to heat buildings. Some data centre operators – particularly Finnish and Swedish – plug their waste heat into district heating networks (Koronen, Åhman, and Nilsson 2020; Davies, Maidment, and Tozer 2016) and this way heat tens of thousands of homes (Wahlroos et al. 2017). So, though Finns or Swedes would never intentionally heat their homes using electrical power – it is much too expensive so they prefer geothermal heating or district heating – the heat energy, which stems from the data centre’s acquisition of electricity, has second-hand value, to the data centre operators as well as to the environment.

In that particular case, using heat as a resource may seem easy, but under many circumstances, using heat – aptly labelled “low-grade” energy – is actually not

straightforward at all, whether to contain, use or transform back into “high-grade” energy such as electrical energy. Whilst waste heat can generally be an asset used in heat pumps, data centre waste heat is not hot enough to have financial value for companies generating electricity from industrial heat (see the introduction to chapter 8).

The energy concerns of a data centre involve increased incoming renewable power, improved strategies for minimising cooling needs, and possibly reclaimed waste heat. In addition to electricity, these processes involve use and reuse of water, which is used for cooling. Water and energy, in turn, can often be reused a few times by different industries – until the water is too dirty or the waste heat too cold, that is, until their levels of entropy have grown too high. Thus, with a systems-based view, the data centre has its place in industrial ecosystems and within the circular economy as a whole. The trick in this case is to make use of the waste heat, as it is much cooler than typical industrial waste heat. In the literature, industrial waste heat below 100°C is considered “*low temperature*” (Brückner et al. 2015). In contrast, data centre waste heat is generally not hotter than 35°C. How to reuse this lukewarm heat is discussed in chapter 6, and one way to solve the issues of water use is presented in the next section.

#### **4.4 A dive into liquid immersion cooling**

Earlier in history – that is, until 2015 or so – server heat was not a major issue for the data centre industry. Surely, it was a costly affair, with often a hundred percent increase in energy costs due to server cooling, and cooling equipment was expensive as well. But until today, server heat transfer has been manageable using air alone. With rapid miniaturisation, physics has eventually brought further concentration of computational power to a halt: the CPUs become so hot that they start to behave unreliably and finally break down. Put in data centre industry-typical jargon, “*thermal management of power*

*electronic systems is a key bottleneck to power densification”* (Birbarah et al. 2020).

The situation has rapidly become worsened by the advent of high-performance computing such as artificial intelligence computation and crypto mining, since these processes require much work from processors, and the more work they perform, the more heat they release. Therefore, new technologies are badly needed, both to increase stability and performance during intense computation, and to avoid even more unsustainable future computing.

In sum, as the temperatures of processors reach higher and higher temperatures due to growth of microprocessors’ circuit complexity (Levin et al. 2016), heat removal using air is often inadequate today. In response to the problem of accelerating heat, there is now commercial interest in *liquid immersion cooling*, a technology for server heat removal introduced in the 1980s (Simons 1994). Indeed, with increased heat, liquid immersion cooling may finally have become *necessary enough* to reach over the industry inertia threshold. And if it has, it will have major impact on data centre configuration. Save for quantum computing, liquid immersion cooling seems to be the only disruptive technology in the data centre industry today. Hence, it is presented in the following.

Server waste heat is generally transported as air. In contrast, in liquid immersion cooling, servers are submerged in what is essentially mineral oil, and so, heat energy is removed from the servers as the temperature of the oil increases. With a heat capacity more than 1100 times higher than air (Eiland et al. 2014), and a correspondingly immensely improved thermal conductivity (Ruch et al. 2017), the operating temperature remains nearly constant, regardless of workload (Bansode et al. 2018). Liquid immersion cooling thus helps to keep the server “rack” temperature under control,

which is especially advantageous to high-performance activities such as deep learning or crypto mining (Lei and Masanet 2020). In addition, because of its high heat removal capacity, it also allows for much more densely populated server racks (or vastly increases the “rack density”, to use the industry term) – no space is needed for in-between server airflow.

The ability to withdraw heat may also prolong the life of server components as submersion removes the risk for corrosion and unclean server parts, adding to the benefits of this cooling technique (Bansode et al. 2018), though it should be said that the constant friction of the oil may instead shorten the lifetime of server parts (Shah et al. 2016). In any case, the servers are not exposed to dust, which under normal circumstances build up under the server chassis and sometimes result in mechanical failure. In addition, submersion fire-proofs the servers.

A major issue of traditional data centres is water use, since water is needed for the cooling towers. With liquid immersion cooling, there are no cooling towers, and consequently no water use. Further, as servers are submerged, the usually noisy server halls turn silent, to the benefit of data centre employees (Ramamoorthy and Krishnan 2018). This point is not to be glanced over. As Ramamoorthy and Krishnan observe, the high noise level

*“is very distracting to data center workers, disturbing their ability to work in the near term and has the potential to cause serious hearing problems in the longer term. Taking these practical issues into consideration, the U.S. National Institute for Occupational Safety and Health has stipulated guidelines that restrict maximum noise levels in data centers to be less than 70 dBA.”* (Ramamoorthy and Krishnan 2018)

One can only imagine what effect large halls of silent oil baths must have on a visitor. Already today, a data centre is a factory where work is carried out despite the fact that nothing moves, a mystery on its own. However, one can *hear* it, loud and clear. With servers submerged in oil in small basins, and with minimal lighting on, a large data hall will be a futuristic crypt, with what should resemble dozens of sarcophaguses, and the silence would give the impression of a sci-fi version of the Egyptian tomb exploration of the 1930s. The fact that a tremendous amount of work is taken place in this silent world is staggering.

As mentioned above, liquid immersion cooling has been available for several decades. A paper from the early nineties cites several cases, including IBM's testing of computer chips and cooling of the supercomputer CRAY-2 (Simons 1994). So with all its promises, why has it failed to gain acceptance?

The answer may have to do with two central concerns of the data centre industry: “uptime”, that is, the percentage of time the data centre is operational, and robustness (security). The concern for uptime is in turn tied to SLAs drawn up with clients, stating, for example, access to services 99.99% of the time (Shuja et al. 2016). As a consequence, the industry has quite some inertia built-in (Klemick, Kopits, and Wolverton 2019), and it is therefore not strange if novel approaches are avoided until truly needed. Whilst industry leaders such as GRC, Izotope or Submer spend considerable effort to prove robustness (Moore 2021), data centres are still built the traditional way. *Or are they?* A major data centre visited for this work has piping for liquid immersion cooling preinstalled in case the client would desire it, to future-proof the data centre without the need of costly or time-wasting modification later. And Microsoft has experimented with running its Teams platform in a liquid immersion

cooling facility, to help the tech giant deal with the daily energy spikes at 1 PM, that is, when afternoon online meetings begin (Weston 2021).

Price is also a hindering factor for liquid immersion cooling solutions. Large-scale data centres are ideally placed away from cities, that is, they are built where land is cheap. If so, there is not much to gain from increased rack density. The oil, the basins and the construction that lifts servers up from the oil result in a system which is costlier than simply placing racks on a raised floor above cooling pipes. *Or is it?* Again, the truth is not so black and white. Especially if scaled up, it is not certain the added infrastructure for liquid immersion cooling will be that more expensive, as normal air-cooling infrastructure is costly in itself, and the added building and land costs are certainly not negligible.

For quite some time now, liquid immersion cooling has been “the next thing”. Today, perhaps the moment is right: it is no longer easy to increase the power density in server racks with air cooling alone. Moreover, its capability to increase server density, to save energy and water, and to provide a quiet environment are appealing features. Consequently, environmental goals such as low land use, zero water use, low electricity use, removal of cooling units and improved working environment for data centre personnel all speak for this disruptive technology.

## **4.5 The data centre as an energy grid regulating force**

In a society built on renewable energy, there is bound to be fluctuations in the energy flow. Heavy industry, a formidable user of electrical power, sometimes lowers its production when energy is scarce, to allow for society at large (households not least) to obtain adequate amounts of power (Energimarknadsinspektionen 2016; Statens

energimyndighet 2004). In the Nordics, this typically implies wintertime: the low outside temperatures mean commercial and residential heat pumps need to work harder.

Using service level agreements, SLAs, a data centre enterprise can be assured of power access, no matter the energy needs in the country. However, analogous to industries cutting back on certain activities during times of low power access, one may wonder if not data centres generally should do the same. Such an idea is easy enough to implement, even as an automatic process. On a small scale, this is done today through “scheduling” data computation between data centres (representative for DeepSquare, interview Basel 3 March 2023). But the strategy is likely to work also on a grander scale. For example, Meta in Luleå could run its AI algorithms when available power is plenty, and fall back to more rudimentary functionality when access is lower – or during very cold days, when the municipality may have higher electricity needs. In the process, the company would save money due to fluctuating energy costs (if so stated in the agreement with the power company).

The simple examples above are not only practical, but also provide a fresh perspective on the regulating forces of a society’s energy system. The flexibility of the power grid is dependent on energy storage, such as battery packs (for temporary storage of energy already retrieved) or hydropower dams (for energy not yet retrieved, but where the promise of power can be paid forward whenever need arises). Traditionally, the regulating force and the elasticity of the grid is the energy store and the lowered or heightened withdrawal from heavy industries. Instead, the flexibility of the energy system can be pushed to the consuming part, the data centre. Doing so opens up possibilities where a society is dimensioned for large-scale data centres: at power peak times, the data centre can lower its energy use, thus serving as a regulator of demand



within the energy grid. This strategy would be particularly beneficial in societies where renewable baseload power is not available, that is, where energy is mostly derived from intermittent power sources such as wind or solar power.

Klingert and Szilvas think along these lines too, and they note that grid elasticity can be made possible since data centres “*have highly automated management processes and a fine-grained power load*” (Klingert and Szilvas 2020). Agrawal et al. argue similarly, pushing for what they call a “*virtual battery*” (Agarwal et al. 2021). In their vision, smaller data centres are placed near sources of intermittent renewable energy such as wind or solar farms, and their workload depends on available energy. The authors presuppose that computation takes place in a distributed network. In other words: other data centres take over when there is too little or too much wind, or, for a solar-powered data centre, at night or during precipitation. Relating to this issue, Agrawal et al. note the problem with high-efficiency computation now having to be distributed between several systems, a process which would likely slow down the process. That said, perhaps not all computation today needs to be so immediate after all? It seems the world’s ICT possibilities have grown so fast that we are, in all honesty, rather spoiled today with bandwidth, download speed and immediate access to information.

And so, a final reflection on this multi-faceted problem. Based on the understanding of what energy is, and how the data centre industry relates to it, it can be argued that heavy users of the energy grid should take *responsibility*. Ultimately, the grid needs to be shared, to avoid expensive energy storage solutions and to avoid conversion losses that energy storage entails. Thus, like many heavy industries, it would only be fair to ask that data centre owners partake in the collective grid elasticity

regulation. This is not an unrealistic view. Other industries are already sharing responsibilities and lowering production on, say, exceptionally cold winter days, and it would be easy enough for the authorities to legislate capping when deemed necessary. Consumers, in turn, would then take their share of this responsibility by accepting that TikTok or NetFlix run less smoothly on especially cold winter days.

Paul Davies shows how even ants, like humans, develop their ideas collectively (Davies 2019), in an interaction he calls “*distributed computation*”. Similarly, we are since long past the reign of the single, or “monolithic”, super computer. Instead, much computation workloads and data storage are distributed among data centre servers, in the interplay of large data centres and edge networks, as well as in peer-to-peer networks, where individuals’ computers form a powerful collaborative force. Especially if rewarded with lowered energy bills, there is no reason why the data centre industry should not expand on its self-image as a distributed network of data to also include societal responsibility. This way, data centres would become grid regulators that keep society going also during winter.

## **4.6 Metrics, used and abused**

It is impossible to ascertain how much the data centre energy consumption is going to increase in the coming decades. From an environmental perspective, the exact figures do not matter: they are large enough to be considered, even on a national level. To keep this high energy use in check, the industry – or rather, governmental stakeholders and sustainability-conscious clients – need metrics that they can rely on, to be able to rank data centre solutions against each other.

Despite large-scale data centres being substantial and heavily energy-demanding buildings, data centre energy matters are today seldom – if ever – viewed in a societal

context but rather as compartmentalised resource problems. This view is much due to what has been labelled a “*very high density of energy consumption and ease of measurement*” (Mujawar et al. 2018). The “*ease of measurement*”, in turn, relates to the key metric for data centre energy efficiency presented earlier, PUE (Power Usage Effectiveness). PUE is the result of all energy going into the data centre divided by the amount of energy used for IT equipment. The lower PUE, the better: when all incoming energy is used solely by servers, that is, when no energy is needed for (mainly) cooling, PUE equals one.

By only relating to cooling, PUE cannot be used as an all-encompassing metric for data centre energy efficiency. Yet, that is exactly how it has been used, almost since its inception in 2006 (see chapter 8). For example, an Icelandic data centre with no waste heat use dedicated to bitcoin mining – with its incredibly high energy demands – will fare reasonably well in terms of PUE, as the outside air would be cool enough for the servers, and a huge proportion of the total energy therefore go to mining activities. As chapter 8 shows, there are many other problems associated with a high reliance on PUE.

There are metrics with greater ambition regarding the data centre itself. One of those, Power Density Efficiency (PDE) (Lajevardi, Haapala, and Junker 2014), aims to bridge calculations on server performance and server rack airflow, or in effect, to bridge computer science and mechanical engineering. Several metrics introduced around 2010 look beyond the data centre, and focus water use (WUE – water usage effectiveness), origin of electricity (GEC – Green Energy Coefficient), GHG footprint (CUE – Carbon Usage Effectiveness) and energy reuse (ERE – Energy Reuse Effectiveness). For a holistic perspective on the greenness of the data centre industry, such metrics are

important, but their wider scopes make them incomparable to PUE, and none of these metrics have reached the level of recognition as PUE enjoys. Most importantly, PUE, this “*imperfect*” (Morgan 2022) metric, has become the metric used for decision-making processes, thus acting as the ultimate incentive for data centre sustainability – or at least it would have been, had it not been so inadequately equipped for the purpose.

This thesis is not primarily concerned with the inner workings of the data centre, but rather views it as an important part of local, regional or global society infrastructure. Hence, rather than strictly investigating desirable data centre configurations from a PUE perspective, the thesis recognises energy and water use, what energy source the power originates from, what sort of work the servers are doing, where waste heat ends up, and what end-users actually do with the computing power.

Some of the actors of the data centre industry, Big Tech not the least, are highly profitable enterprises. Some others may not be so concerned with funding for running data centre servers (such as defence departments). In other words, money may not be the most efficient driver for energy saving measures of the industry. However, sustainability is also a key aspect, not least to local politicians. Clearly, the metrics they use have to be both accurate and, above all, purposeful. The questionable purposefulness of PUE for external stakeholders’ decision-making has already been discussed. Hence, new metrics are needed in order to facilitate the move to a more sustainable industry. Chapter 8 discusses current metrics and proposes a new one.

## **4.7 Chapter conclusion**

The very essence of energy has been the focus of this chapter, together with the imagined view of small and big promises floating around us. And it should be remembered that whenever we use any sort of device – mechanical or electrical – the

energy content is not in any way diminished; the promise is only displaced. Retrieving the energy displaced certainly has its advantages, for the wallet and for the planet. This fundamental truth constitutes the founding premise of the thesis.

Energy cannot be seen, as there is nothing really to *be* seen. However, depicting energy flows is easy. Figure 10 shows European electricity flows for an October day 2022. It is fascinating to view these flows, when one understands that what is actually sent is a number of big and small promises. It is also fascinating that once the electricity has been used, the promise is still there, ready to be sent off (typically as heat) to the next recipient. Had it not been for its unfortunate term, “waste heat” would perhaps have been more highly regarded and more widely known than it is. If so, stakeholders would maybe have looked closer at all those promises already floating around, rather than focusing one-sidedly on how to limit energy use of data centres. Thus, a holistic view on data centre energetics would serve both the governmental stakeholders and the industry better – the authorities as there would be less strain on power grids, and industry as when the authorities are happy, it is easier to build new facilities.

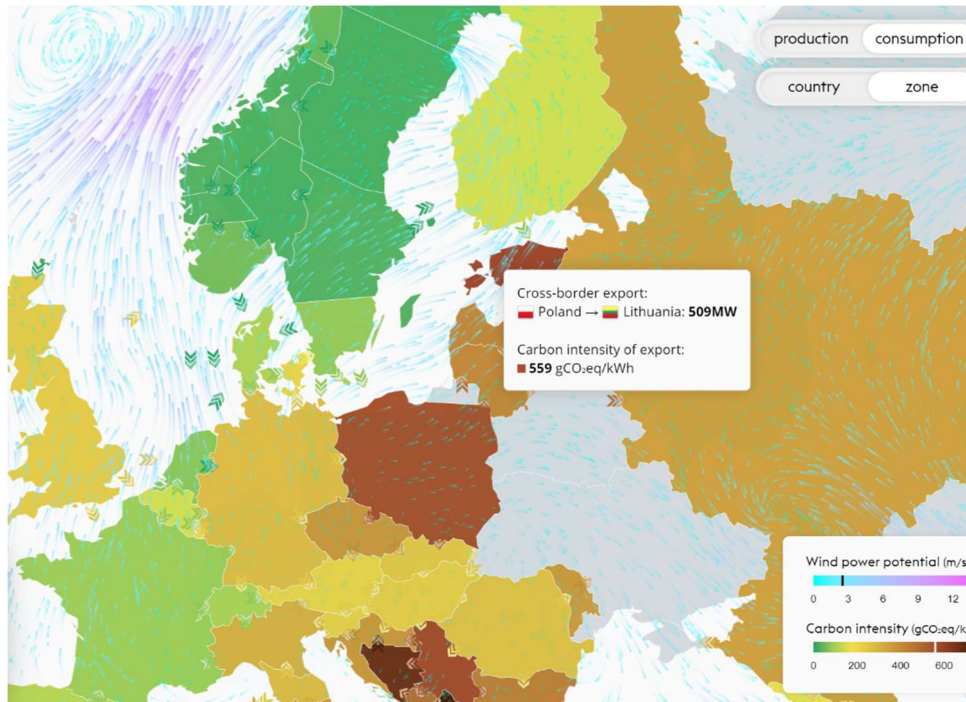


Figure 10. Energy flows in Europe on 27 October 2022. Each major arrow shows one energy flow between zones, and the colour of the arrow shows the “greenness” of that energy. The thinner arrows indicate wind direction, stressing the role of renewables in the European energyscape. Screenshot from Electricity Maps: <https://app.electricitymaps.com/map>

To conclude, it seems plausible that when combining systems engineering with the holistic visions of Boulding and Murray, we can move data centre energetics forward drastically: As mentioned above, today’s data centres use up somewhere in the region of one percent of world electricity. If their waste heat could pay forward the energy promise, we could therefore – profitably and at least in technical terms relatively easily – ideally offset the electricity consumption of approximately one hundred million people. So great is the combined power of systems perspectives, holistic approaches and the idea of the circular economy.

# 5 A framework for data centre placement

Data centres enable computation and communication around the world. Like GHG emissions, they go about their business without much respect to national boundaries. Commenting on the cloud’s “*relationship between place and placelessness*”, Hu (introduced in chapter 2) writes:

*“The cloud thus offers a vision of globalization that follows the dictates of multinational corporation – a coalition of geographic areas that move capital and resources through the most efficient path. Just as it is cheaper for Apple to use Ireland as its tax domicile to avoid paying US taxes on its French operations, for example, it is more efficient for Facebook to serve some of its Japanese customers from a Singapore data center.”* (Hu 2015)

In other words, where to place a data centre geographically is far from obvious, and something that needs investigation from many angles. This chapter therefore offers a plenitude of dimensions of concern for data centre placement. Some of the concerns directly relate to energy, but others’ connecting points do not. Indeed, viewing data

centre energetics holistically necessitates examination of both primary and secondary concerns: that is the only way to adequately fit data centres in societal and environmental contexts.

In the literature (e.g. (Barroso, Hölzle, and Ranganathan 2019)), as well as in industry briefings, it is generally supposed that a large data centre is a traditionally built entity, that is, a warehouse with a footprint of one or a few football fields. That is also what a typical large- or hyper-scale data centre looks like today. However, this forward-looking thesis does not concern itself with today's industry landscape, but with tomorrow's. In that imaginary place, a data centre can come in many shapes uncommon today: immersion-cooled, distributed among many households, placed under water or greatly energy-enhanced through waste heat recovery. These and other possible data centre configurations bring many possibilities that are not taken advantage of today, and also generally overlooked in the discourse.

Data centres face a number of everyday problems. This implies that data centre *placement* faces problems as well, and needs careful consideration. Energy was dealt with extensively in the last chapter. Some other issues are shown below.

### *Security*

Among the data centre industry's top concerns is to provide a secure environment for its customer data. As a consequence, retired military bunkers are sometimes used to establish highly secure data centres, as this chapter will show. It should be remembered though that every part of a networked system – be it a laptop computer, a smartphone, a WiFi connection or a server in a data centre – risks intrusion, hijacking or being tampered with. As pointed out during a data centre visit for this PhD work, the guards and mantraps securing the facility are of little help here, and despite often rigid physical



security measures, the main asset of a data centre is the collection of data, not the hardware.

#### *Water usage*

Water for cooling, and in some cases for hydropower, is becoming an increasingly bigger environmental and reputational problem for the data centre industry. At least this is the case for data centres placed close to populations, since water scarcity is often high in the world's more populated regions (Kummu and Varis 2011). Accordingly, excessive exploitation of fresh water sources can yield major problems for the local infrastructure (Obringer et al. 2021; McCammon 2017; Sattiraju 2020).

#### *Land usage*

If powered with solar power or wind power, data centres need much more land for energy production than for hosting servers at the site itself. And since wind power and solar power may, in some places, be the most viable sustainable energy solution, that is an issue one should account for (i.e. when finding new locations for data centres, it is presupposed one does not build for a dependency on fossil fuel). A way around the problem of land use is to build in the Arctic or the desert, where population density is low and land less expensive (Hu 2015). A more radical solution is to place data centres on the sea (Clidas, Stiver, and Hamburg 2009; Terenius et al. 2021) or under it (Microsoft 2019; Judge 2022b).

#### *Trust and legal matters*

Trust in the data centre owner, SLAs, intellectual property infringements, national data protection laws and so on may be among matters of concern for a data centre owner or manager. Indeed, the information stored within a data centre is “one of the most

*financially concentrated assets of any organization” (Koomey et al. 2008), and must be dealt with accordingly.*

*The risk of establishing a surveillance society*

The more information we as individuals hand over to corporations such as Amazon, Meta or Google (also through their subsidiaries such as Amazon Prime, Facebook or YouTube), the more they know about us and the more power these organisations have. Not only corporations but also governments can use data centres to keep control, for example to track their citizens and to enforce censoring.

To summarise, physical and “cyber” security, water use, land use, waste management and personal integrity concerns are all parts of the problematics of the daily life of a data centre owner. Some of the members of the casual list above are sub-parts of larger focus areas such as energy use, client satisfaction, sustainability endeavours and adherence to legal requirements.

As the problems facing the data centre industry are always entangled, the issues above cannot be ignored here, despite the thesis’s primary focus on energy. For example, with higher security measures in financial transactions follows higher energy use (due to enhanced software security protocols) and increased embodied energy (due to increased physical defensive measures). Also, in some instances, there is a trade-off between water use and GHG emissions (Obringer et al. 2021), and when a large-scale data centre uses hydropower, it may be at the expense of the local community (see section 5.3.2).

In other words, an understanding of different matters pertaining to data centres is needed for the energy-related investigation to have any value. Therefore, based on

the previous chapter's more general discussion regarding energy, data centres and societies, this chapter presents many additional topics for enquiry. As this is a cross-disciplinary undertaking, the topics span societal, technical, financial and sustainability-related matters. In total, there are about twenty hierarchically structured subtopics, which have initially been mapped from chapter 4 and then explored through further consultation of the literature. When explored, these subtopics should give an insight in today's and tomorrow's concerns for data centre placement. In addition, the topics will constitute the analytical framework used in the Analysis chapter (chapter 9). Figure 11 reintroduces the analytical framework, though consulting the figure – whose indexing and naming mirror this chapter's heading structure – is not necessary for the understanding of this chapter.

The topics are presented and discussed below. Some of them may seem to point to obvious and generalised solutions. However, complexities pertaining to the specific topic are sometimes more particular than at first glance anticipated: local circumstances often pose different challenges for the data centre industry. Consequently, and axiomatic to this thesis, it is assumed that location and local needs do matter.

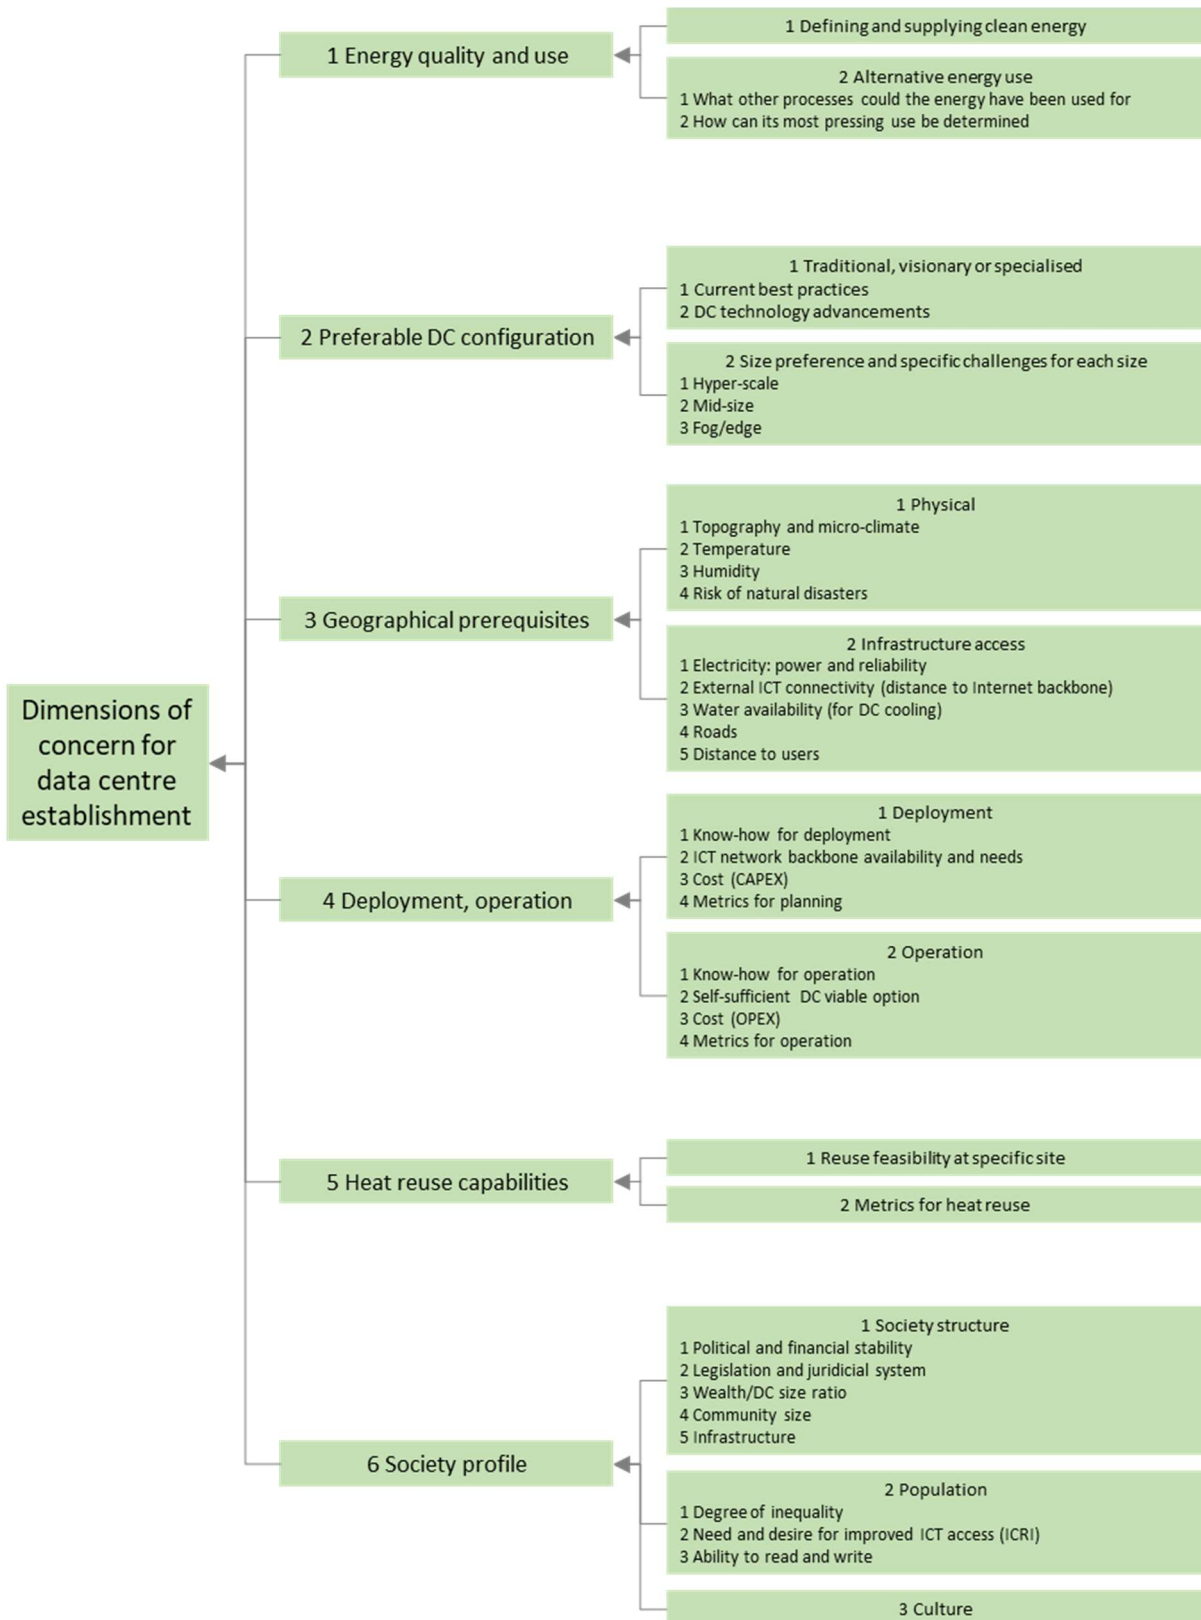


Figure 11. The analytical framework of the thesis (reintroduced).

## 5.1 Energy quality and use

Data centres not only power and empower the world, but also use up a substantial amount of available electrical power: as previously mentioned, about one percent of today's consumption (Masanet et al. 2020). Of course, this use, especially if growing uncontrollably, is a major sustainability issue, in a world to a large extent powered by fossil fuel and with steadily increasing energy needs. ...and it has been debated whether or not this will happen. On the one hand, data centre energy use has been static for the last decade (Masanet et al. 2020). On the other hand, cooling has been a culprit in power use, but it is not any longer: with a cooling overhead decrease from a hundred to, in large data centres, circa twenty percent, there is not that much more room for improvement. Moreover, because of rapidly growing economies in Southeast Asia and Africa and higher focus on latency, it is reasonable to believe future data centres will be placed nearer the tropics. This means that power use (again, for cooling, but also for increased data computing volumes) will go up. Discussions relating to this topic are found in chapters 6 and 8.

As explained in chapter 4, energy can be seen as a large set of promises, and as the first law of thermodynamics states, these promises can only be handed over, never vanish. However, their value to us strongly depends on the energy *carrier* – what it is, where it is and how much is retrievable from it. Thus, though seemingly straightforward in nature, *energy quality* relates to views on energy and energy use. Such views are not entirely objective, but framed within the context of social complexity. In other words, here is where generalisability and uniqueness meet in a chance encounter, and the point at which we should stop thinking about energy as a physical concept only.

### 5.1.1 Local definitions of clean energy

As stated in the opening of this chapter, GHG emissions know no national boundaries. Thus, it may seem that clean energy is, well, clean: basically not fossil fuel. However, as with so many other topics concerned with here, the matter of energy cleanness is not so binary after all. A very recent example (2022), touched upon in chapter 4, is the debate within the EU whether fossil gas and nuclear energy can be viewed as sustainable energy sources. It is well-known that fossil gas has serious consequences for the environment – it may even be as potent a GHG as coal if methane gas leakage during extraction (Pandey et al. 2019) and transportation through leaky pipes (Weller, Hamburg, and von Fischer 2020) are included. Nuclear energy, in turn, is a low-carbon emitter, but its associated risks should not be overlooked – and among policy-makers, they certainly are not. For the data centre builder relying on sustainability-bound contracts, the very definitions can therefore potentially pose both ethical and financial problems.

As shown in chapter 4, even hydropower, solar power and wind power have their environmental footprints (Obringer et al. 2021). Moreover, local governments view them differently. For example, solar power is used in Great Britain, but the land use of solar arrays may be debatable, as at least England is quite densely populated. So, what would a UK-based data centre operator think about using solar power? And – perhaps a bigger question – what would its clients feel about it, bound as they may be by GHG emission targets? Would they bet on solar power rather than nuclear power to be classified as sustainable in Britain, also ten or fifteen years from now?

It may well be that the question whether or not fossil gas and nuclear energy are classified as sustainable will largely redraw the clean energy landscape of Europe, and it is probable that similar discussions will take place on other continents.

### **5.1.2 Alternatives for energy use**

Some of the Big Tech companies but also other data centre operators are rapidly increasing their shares of renewable energy (Chalise et al. 2015; Google 2020). It is undeniably positive that these sustainability measures are being taken. However, wherever there is a shortage of available sustainable energy, other services may have had better use of it. For example, if a major hospital and a planned Amazon data centre both would be in need of additional power, it would make sense from an ethical perspective that the hospital has precedence, and that the data centre in question is denied building permission (cf. section 4.5). This is not only a hypothetical scenario: Today, both cities and nations are capping data centre establishment, because the energy is needed for other purposes (Laan 2008; Department of the Environment Climate and Communications 2021) and even more so as both fossil fuel and nuclear energy are phased out.

What complicates such a discussion is that with electricity use of mid-sized countries (Bryce 2020), Google (Google 2020) and similar companies build their own electricity production units, such as new solar power parks. This leads to questions that are simultaneously ethical, philosophical and utterly practical: In many parts of the world, there is an obvious shortage of not only available renewable energy, but also of *obtainable* renewable energy (Berners-Lee 2021). In such a case, should a new data centre – let alone, a collection of data centres – be allowed? If the decision is affirmative, what effects would this have for regional infrastructure and energy security? If negative,

then what precedence does this have for other industries? Answering “it depends on what the data centres are used for – bitcoin mining may be out, but genome research accepted” would likely be impossible for a politician, due to borderline cases and the risk for being trapped in conflicts of interest.

Moreover, even if it were possible to obtain enough renewable energy at the desired site – say, if there is enough room for solar panel arrays – is this the best use of the materials used in the process? For example, should there be a worldwide shortage of solar panels, one may object to many of them being swooped up by Big Tech.

## **5.2 Preferable data centre configuration**

If one makes a true commitment, solving many of the problems relating to the data centre’s power consumption is doable – as seen by yearly PUE efficiency improvements by Google (Google 2020). In fact, there are many well-known strategies for this work, including server virtualisation and keeping correct airflow in the server room (Barroso, Hölzle, and Ranganathan 2019; Zakarya 2018; De Napoli et al. 2016). However, as discussed in sections 4.6 and 8.1, PUE has clear – and limiting – system boundaries. To improve data centre sustainability, one also needs to address water usage, GHG footprint, the degree of energy reuse, construction materials of the data centre, the life cycle of on-site servers and other matters.

With this expanded view of data centre sustainability in mind, this particular section looks at preferable data centre technologies and configurations, today – and also tomorrow.



### **5.2.1 Traditional, visionary and specialised data centres**

A typical large data centre is a warehouse building made of concrete and steel. This means that even before servers are installed, the building has had a significant GHG footprint, from cement and concrete production as well as from iron extraction and steel manufacturing. Moreover, both products need transportation to the site. Building with timber is beneficial in that it lowers GHG emissions and traps CO<sub>2</sub>, and a few data centres have in fact been built in timber: EcoDataCenter (see section 7.5) is one of these. It should be noted that building in timber instead of in steel has a trifold environmental gain: lower emission during production, lower transportation emissions (if forests are nearby) and CO<sub>2</sub> entrapment for the longevity of the building. Against this background, when designing data centres and when making predictions or recommendations for future establishment, one should consider that future data centre solutions may look very differently from today's.

For a long time, data centres have become increasingly bigger. To allow for even more computation and storage, future directions may involve stronger emphasis on distributed systems (see section 4.5), where hyper-scale data centres run background operations, and much smaller edge data centres run communication with nearby users in order to decrease latency and minimise energy use. Here, it should be noted that as mentioned in section 5.3, data transmission networks use approximately as much electricity as data centres do. Putting compute closer to end-users therefore decreases the need for some of the power-draw of transmission networks. For the smaller edge data centres, the relative energy use would likely be higher, since cooling infrastructure needed has scalability advantages. On the other hand: should an edge data centre be placed in an office, hotel or residential building (see section 6.1) in a reasonably cold country, then its airflow could easily be put to good use.

There is no shortage of newly built and planned traditional data centres. Consequently, we will live with conventional server racks for quite some time. Still, liquid immersion cooling (discussed in section 4.4) has many advantages. For one thing, it makes the argument for waste heat use stronger, since it is easier to transport heat as a liquid, and since higher temperatures are reachable through this method. Liquid immersion cooling also substantially decreases data centre square footage, as servers can be more densely packed because of its improved heat removal capabilities.

It is safe to say that a global rollout of liquid immersion cooling would have quite an impact on data centre configuration. On top of what has been mentioned above, it would improve server reliability in humid areas, since servers are not in contact with air. Thus, waste heat use aside, a good use case for liquid immersion cooling could be Hong Kong, where space is extremely limited, ambient temperatures high and the air humid.

Data centres need backup power in the event of a power failure. For sustainability reasons, there is a trend among big data centre owners to replace existing backup diesel generators with enormous battery packs. Just accounting for the disposal of the diesel generators, acquisition of the battery packs and nightly recharging of the batteries, it is doubtful whether this exchange is beneficial to the climate. Further, piling up battery packs withdraws them from other uses, such as the automotive industry, which is concerning: Apple's latest lithium-ion based battery backup system alone (Tesla 2021; O'Kane 2021) stores as much energy as the batteries of a quarter of a million of Tesla S luxury cars (Mohamed, Halim, and Zakaria 2021). Thus, the lithium used for Apple's backup system for just one gigantic data centre could have powered hundreds of thousands of smaller cars or likely millions of scooters. Therefore, even

though the data centre's sustainability rating may increase, it is questionable whether the move to battery backups truly is a sustainability gain. And an increased sustainability rating from switching to lithium-ion battery packs is questionable in itself, due to the environmental and social issues of lithium and cobalt mining (e.g. (Crawford 2021)). The change is especially questionable – and somewhat ironic – since *as opposed to scooters or cars, backup systems are meant never to be used.*

The reason why lithium has trumped lead and sodium as the alkali metal of choice for battery making is its higher energy density capacity, or in other words: a lithium-ion battery weighs less than its lead or sodium counterparts. Therefore, it is well-suited for cordless tools and vehicles. Whereas lithium sources are quite limited, and lead can be environmentally concerning, sodium is abundant and non-toxic. Particularly with new research and commercial breakthroughs on sodium batteries (King 2023), it may now be more reasonable to use sodium for stationary battery packs – such as data centre backup batteries (Abraham 2020). Doing so would leave more lithium to batteries that need to be light-weight, serving uses from household items to cars.

This section has covered new technologies with the possibility to change the future of data centre configuration. It should be said that there are other gains from current engineering advancements, not least in the continuous improvement of cooling systems and processor architecture. “Greener” cement, which can substantially lower the GHG footprint of cement production, will also help. However, incremental winnings do not fundamentally change the prospects for data centre configuration, and are therefore not dwelled upon further in this thesis.

### 5.2.2 Size preference

From local racks to large warehouses containing thousands of servers, data centres supply all sorts of ICT services, from storing and delivering emails to running large-scale scientific calculations. Therefore, Big Tech relies heavily on data centres. In fact, the ability to store and compute customer data – in the cloud rather than on the end-user’s desktop – is at the very core of their businesses today. That said, many modern organisations, from universities and hospitals to small and big enterprises, put their data in large-scale data centres. There are at least three main reasons why:

1. Security – in case of burglary or fire, it is advisable to keep (copies of) valuable data outside of the premises, and using a professional data centre nearly eliminates the risk of failed backup routines. Moreover, it is possible (though not certain) that immaterial security measures are given higher priority at the external data centre, since security is part of the data centre owner’s core business.
2. Cost – especially for small organisations, hiring IT personnel to establish and maintain a local (“on-premise”) data centre is often inefficient use of human and business resources. Through consolidation, a large-scale data centre has the advantage of scale here.
3. Focus – to focus on the core business, it makes sense not to maintain data on-premise, in the same way an organisation does not own a power facility.

The service can be “cloud-based”, that is, using servers also utilised by others. Amazon’s AWS and Microsoft’s Azure are two of these, though it should be said that web-based email is hosted similarly, say by Google (Gmail). Computation capability and storage can also be done on own, or “dedicated”, servers, put in so-called “co-

location” data centres. Here, one pays for space, power, security and, not least, high-speed and high-volume Internet connectivity; the client is in charge of the actual servers.

A large data centre – though consuming many megawatts of electrical power – is much more energy efficient than a small. This is for several reasons. First, cooling can be made more efficient. Second, server performance can be dimensioned more freely. Third, due to their size and power demands, the largest data centres are positioned where environmentally and financially sound.

Because of the significant investment costs involved, “hyperscalers” are often built in modules, completed as demand rises. This is the case for EcoDataCenter in Falun, Sweden, visited in August 2021. On paper it is an 80 MW facility (EcoDataCenter 2021). In reality, only one of the planned four modules was completed at the time of the visit, though a second module has been built since then (EcoDataCenter 2022). Thus, capital costs can be minimised, and investor or owners can be assured that there is ample room for expansion. Meta’s main European data centre, based in Luleå, Sweden, has been built the same way: one building at a time. The first of the social media giant’s buildings was inaugurated in 2013 (Vonderau 2019).

Server virtualisation, a major trend around 2010 and today a given, is the process of combining the services of several servers onto one physical computer. At the dawn of this millennium, a small organisation may have kept one server for email handling, one for the organisation’s web site and one for the employees’ files (appropriately called “mail server”, “web server” and “file server”). In such a setup, each server would have run with a fraction of its capacity. In contrast, with server virtualisation, these would be consolidated onto just one machine, saving just about two thirds of the electrical power.

The larger the data centre, the easier to optimise server virtualisation, or other sorts of portioning of data storage and computation, depending on server capacity and use case. This is especially true for dedicated server environments, as some types of software applications are designed for niche servers. For instance, AI applications often run on processors used for graphics computation (GPUs), and so does crypto mining, as shown in Figure 12.



Figure 12. Crypto mining in October 2020 using graphical processing units (GPUs) in the vicinity to Luleå, Sweden. On this latitude, and with unpolluted air, servers can take advantage of free cooling (note the snow in the background). The photo shows parts of a system of 500 kilowatts of bitcoin mining in progress.

Building small data centres that can compete with large data centres' energy use efficiency is more problematic. In truth, the data centre industry is an economy of scale. That being said, there are situations where smaller data centres can make more sense from a sustainability point of view. One example is the Costa Rican case (see section 7.4), where waste heat can possibly be reused almost fully. Alas, few hyper-scale data centres take advantage of waste heat. And why should they? Compared to the revenue stream from selling computation or storage, the return on investment (ROI) from waste heat use may not be perceived as high enough to bother. Moreover, as waste heat use

unfortunately does not improve PUE, there is little sustainability competition incentive for these actors to do so.

### **5.3 Geographical prerequisites**

The main issue concerned with in this thesis is reclamation of data centre waste heat. As shown in 7.5, provided the municipality infrastructure is sound, reclaiming substantial amounts of the waste heat in cold regions is not technically difficult. However, that is not where most people on the planet live.

Data to and from data centres are sent as photons, and so, in theory, a data centre's connectivity is limited by the speed of light. In practice, the fibre optic cables transmitting these data reduce the speed to about two-thirds of this figure (Stockholm Data Parks 2019). The Internet is a giant network, connected through several high transmission (subsea and on-land) cables. To send data between two devices such as smartphones or laptops, the information also needs to go through gateways, like train stations connecting rail networks. (In fact, this is more than an analogy: fibre-optic cables tend to follow railroad tracks (Hu 2015).) "Changing tracks" increases latency. Thus, the fewer the gateways, the faster the transmission. On a related note, the combined energy use of the gateways and related infrastructure is equivalent to the data centres' (IEA 2022), so again, the closer the data centre is to the end-user, the better for the environment.

New technologies such as 5G rely on extremely short transmission times – down to the millisecond-level for a roundtrip between, say, an autonomous vehicle and the local data centre. Some of today's leading sustainable data centres are built in the Arctic regions of the Nordics. Here, the populations are disproportionally small compared to the data centres they host, meaning that they need to serve people further away. Today

a relatively minor problem, the proximity advantage may prove more significant with future technologies' possibly higher demand for fast connectivity. None of the Nordic countries' cold regions is ideal in this respect; since most people in the world live closer to the equator, there is a trade-off between inexpensive or even free cooling and the vicinity to people.

As mentioned above, this distance is not a paramount problem at this point (after all, information is carried by light), and the projected lifetime of a typical data centre does not exceed two decades (Barroso, Hölzle, and Ranganathan 2019), so future-proofing is not a major concern. Like Edgevana's CEO stated at an industry conference in 2020, "*not everyone will need sub-ten millisecond latency*".<sup>2</sup> However, what is concerning is that with the quest for low latencies, more data centres need to be built closer to the majority of the world's populations, that is, in tropical and subtropical nations. This insight demonstrates why the need for sustainable data centres in the tropics is such an important and pressing matter.

A data centre can be placed anywhere on Earth – in a city, in a forest, in a desert, in snow, even under water. However, every location comes with its own issues, as discussed below.

### *City*

Locating a data centre in a city is ideal in terms of proximity to users (and of special importance to services relying on important microsecond decisions such as financial

---

<sup>2</sup> DCD, "Building the Edge. A global discussion on edge infrastructure & innovation", 27-28 May 2020. Talks are still available, at <https://www.datacenterdynamics.com/en/broadcasts/edge/2020/>.



transactions or self-driving vehicles), and to recruit operating personnel. However, land cost would typically be high, and keeping a secure perimeter around the buildings may be more difficult than in rural parts of the country (see section 7.5). On top of this, since a large-scale data centre may be an unwelcome sight in a city due to its size, since it is a significant user of both electricity and water, and since cooling towers may be noisy (Barakat 2022), such a facility adds strain to the local infrastructure. As a consequence, it should not come as a surprise that cities with an abundance of data centres limit further development. Amsterdam is a well-known example (Laan 2008), but others plan to follow suit, such as Dublin (Department of the Environment Climate and Communications 2021).

### *Desert*

Placing a data centre in a desert has one obvious advantage: low land cost. A desert in close vicinity to more populated country is therefore an attractive place to build on. On the other hand, there are plenty of drawbacks. First, heat poses an obvious problem, as data centres need cooling to operate. (This problem is smaller than one may think though. Removing heat from dry air is more efficient than from humid air, and the outside air is not the primary driver of inside air temperatures; computer work is (Sperling 2009).)

Second, though a desert is benevolent for solar panel installations, sand dust and high operating temperatures lower the efficiency of the solar panel arrays (Kennedy et al. 2021; Rao et al. 2014). Third, providing fresh water for the cooling towers is an expensive – and unsustainable – affair, threatening to consume much of local water resources (Mytton 2021; Data Centre Magazine 2021; Sattiraju 2020; McCammon 2017). Fourth, this is typically not where people live, so attracting operating personnel

may not be so easy. Fifth, depending on the location of the desert, proximity to the end-user may be an issue. For the large scale data centres in the USA, proximity would not be a major problem (provided the data centres serve populations nearest to them) as the distance to major populations is not that great: even the desert states New Mexico, Nevada and Arizona are homes to over two million people each, and border on populous states such as Texas and California. For a data centre placed in Northern Sahara, proximity may well be an issue, though.

### *The Arctic regions*

Erecting data centres in the Arctic regions works well. In fact, the Nordics are currently facing substantial investments, partly for the cold climate and access to reliable hydropower (see section 7.5). The practically limitless fresh water supplies also help. The drawback here is, again, proximity to users, as most of the world's population lives close to the equator.

### *Under water*

Finally, submerging container-sized data centres under sea water has been tried by Microsoft (Microsoft 2019; Periola, Osanaiye, and Olusesi 2021; Brown 2017), with good results. Doing this on a grand scale in the tropics may actually be feasible, not least with new subsea cables built around Africa (AUC/OECD 2021; Adame 2021; Cotterill 2021), allowing for high volume, fast and reliable Internet connectivity. Ben Cutler, Microsoft's manager for their experimental "Natick" underwater data centres, has also pointed to the physical stability of the ocean floor: temperatures are low, and storms are confined to the ocean's surface (Judge 2016). However, the idea comes with a few obstacles. First, obtaining access to the data centres is not easy for the data centre owner (Periola, Osanaiye, and Olusesi 2021): it requires transportation on the open sea

and bringing up the data centre in question. Next, there may be security issues involved, since there is no security parameter for the data centres. Nonetheless, underwater data centres should be tempting around tropical coastlines, and the first semi-commercial underwater data centres are currently being built in China (Judge 2022b). Still, the fact that these cannot easily reclaim the waste heat raises concern. In a nutshell, they achieve great PUEs, but compete less well with holistic energy efficiency metrics (see section 9.2). More studies are also needed to address sustainability concerns: how do the submerged data centres alter the local environment (the micro-climate)? What are possible side-effects of corroding data centre structures? Interestingly, Microsoft has filed a patent (Cutler et al. 2016) for creating artificial coral reefs out of their collections of Natick data centres, and is thus designed to “*actively promote reef life and sustain a surrounding ecosystem*” (Cutler et al. 2016). The patent outlines how components are added to minimise vibrations and sound waves, not to disturb the ecosystem, and how a plurality of sensors will report on the health of the coral reef. Here, if successful, the bio-fouling that may become an obstacle for submerged data centres (Judge 2016), thereby turns into an asset for the local marine environment.

#### *Distributed workloads*

Another option altogether is to utilise end-users’ unused computing resources, thus replacing data centre computing capability with massively distributed workloads (the web comes to mind). This is not a new idea, and especially useful as a free alternative to high-performance computing (HPC) clusters. One famous example is SETI@Home, established by Berkley researchers in 1999 (Korpela et al. 2001). Here, thousands of idle personal computers in people’s workplaces or homes worked together on the search for extra-terrestrial intelligence (Agrawal and Kulkarni 2019). Following SETI@Home, similar projects have been established, such as Folding@home, which aims to find

folded protein structures for use in medicine (Agliazanov, Sit, and Demir 2020). In all these cases, end-users are involved in a crowd-sourcing activity, but rather than involving people, the entire process is automatic and only uses computer resources. Some such projects, taking advantage of *citizen science* (in which individuals help researchers address research questions, usually through gathering or filtering data) also involve research relating to environmental issues (Agliazanov, Sit, and Demir 2020). However, one can question the rationale behind using computer resources for such calculations. After all, also the client computers use power, and their all-purpose CPUs are not optimised for these sorts of calculations. Thus, moving the workload from the data centre to end-users is a two-edged blade: on the one hand, it saves computer hardware and does not affect research funds much, on the other, the accumulated power might become higher due to the use of non-specialised hardware and software configurations. From an environmental perspective, the question is whether the servers' GHG emissions over their lifetime are higher or lower than the partaking personal computers' power use, and if heat is likelier to be reused in people's homes than in a data centre's context. Still, citizen science raises the public's interest in research, and that educational element of this strategy has value in itself.

As shown in the above, many issues pertain to data centre placement. In terms of geography specifically, these further relate to physical matters such as the *proneness to natural disasters* and to *infrastructure access*. Sections 5.3.1 and 5.3.2 comment on these aspects in more detail.

### **5.3.1 Physical**

“Earth” has two meanings. It is our home in the universe, but it also the earth on which we stand, and on which we build our infrastructure. In our everyday lives, we take earth

for granted – we do not wake up in the morning expecting to suddenly fall into a hole, pass through the core and ultimately get ejected into outer space. For an engineer or a geophysicist, however, Earth is not a solid object. Indeed, the firm ground we live on is not only limited in terms of area, but also in terms of *thickness*. If we imagine the world the size of an apple, the crust is as thin as the apple skin. Beneath the crust – the mostly solid, but sometimes half-molten mantle and occasionally huge quantities of magma.

This fact is not without its implications. Though we do not need worrying about digging too deep to construct foundations for data centres, the earth becomes unstable where the tectonic plates grind against each other or go apart. Entire nations such as Indonesia or Japan are built on volcanic islands that are the result of tectonic plate movements. This means that many people live under the constant threat of earthquakes or tsunamis (themselves often a direct effect of earthquakes). For example, large cities such as Tokyo, Santiago and San Francisco are exposed to these risks. As a consequence, these may not be the most appropriate places to build large-scale data centres. As seen in Figure 13, earthquake-prone regions, and regions haunted by volcanic eruptions, follow plate tectonic boundaries well.

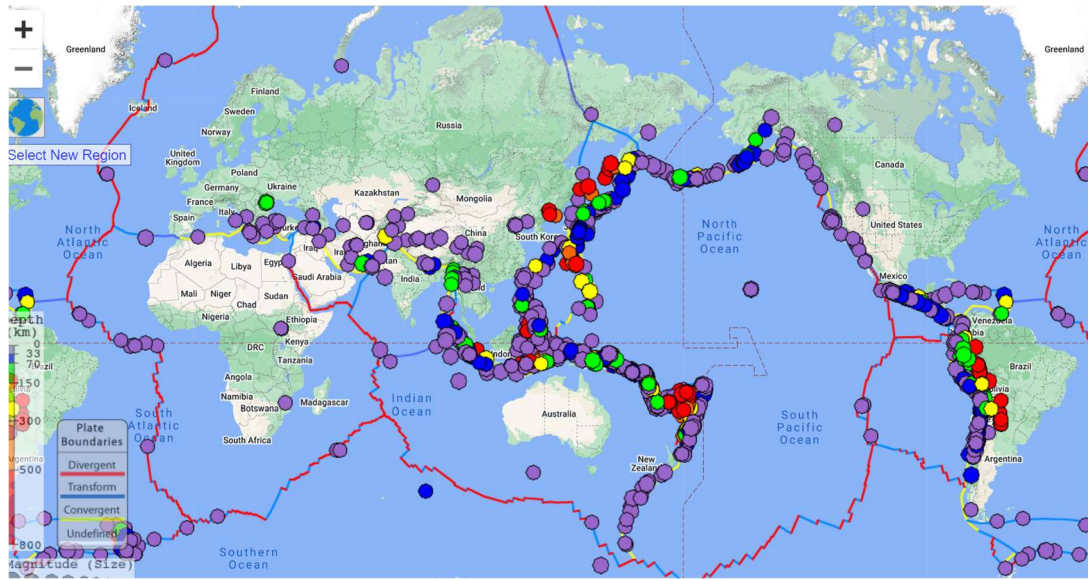


Figure 13. The thousand largest seismic events (earthquakes) since 1970, as shown by the IRIS browser. The map is centred so that the Pacific “Ring of Fire” – connecting Japan in the west with California in the east – is clearly visible. Software reproduced by permission from the Incorporated Research Institutions for Seismology, USA. Map used by permission from Google. Link to visualisation: <https://tinyurl.com/tereniusearthquakes>

Unfortunately, many of today’s major cities are situated in areas on the edges of the plates. Since uptime and security are top priorities of data centre operators, optimal places for data centre placement must be found elsewhere. This is also one of the reasons why American banks have built their data centres in Arizona rather than California: it is close enough to serve large portions of the US population, but less exposed to unpredictable weather conditions, earthquakes included (Crawford 2021).

Relief from extreme weather conditions and exposure to natural disasters is something that data centre owners take in consideration when establishing large-scale data centres, and is used by industry organisations in the Nordics to attract clients (Dybdal Christensen et al. 2018). Whereas the Nordics are well positioned in terms of plate tectonics and natural disasters, California, Japan and Indonesia are definitely not. And locating data centres in those areas is not just a problem for the data centre owner:

one can only imagine what would happen in our personal lives and in industry if large-scale data centres of governmental institutions and Big Tech would become damaged.

The ability to cope with disasters as they happen is also an important ability, perhaps even more important than exposure in the first place. This is discussed in the next section.

### **5.3.2 Infrastructure access**

A data centre needs access to the “Internet backbone”, a network run by key Internet traffic actors. For reasons of speed and stability, the data centre needs to access this network with as few Internet exchange points (IXPs) – routers and switches that is – as possible. For a data centre located in Europe, USA or Japan, the distance to the Internet backbone would be short. In some other parts of the world, say, land-locked countries in sub-Saharan Africa, the distance may be longer, with increased latency and stability worries as a consequence.

A single fibre-optic cable can only carry so much data. In other words, to supply the world with high-data activities, such as streamed video, a number of cables are needed. The aforementioned Europe, the USA and Japan therefore feature a number of powerful and costly Internet backbone (“Tier 1”) networks, but until recently, sub-Saharan Africa has lacked those. Recently built (as well as additional planned) sub-sea fibre-optic cables connecting Africa should decrease lag and increase possible data volumes. Figure 14 shows major existing and planned Internet cables of large parts of Europe, Africa and Asia.

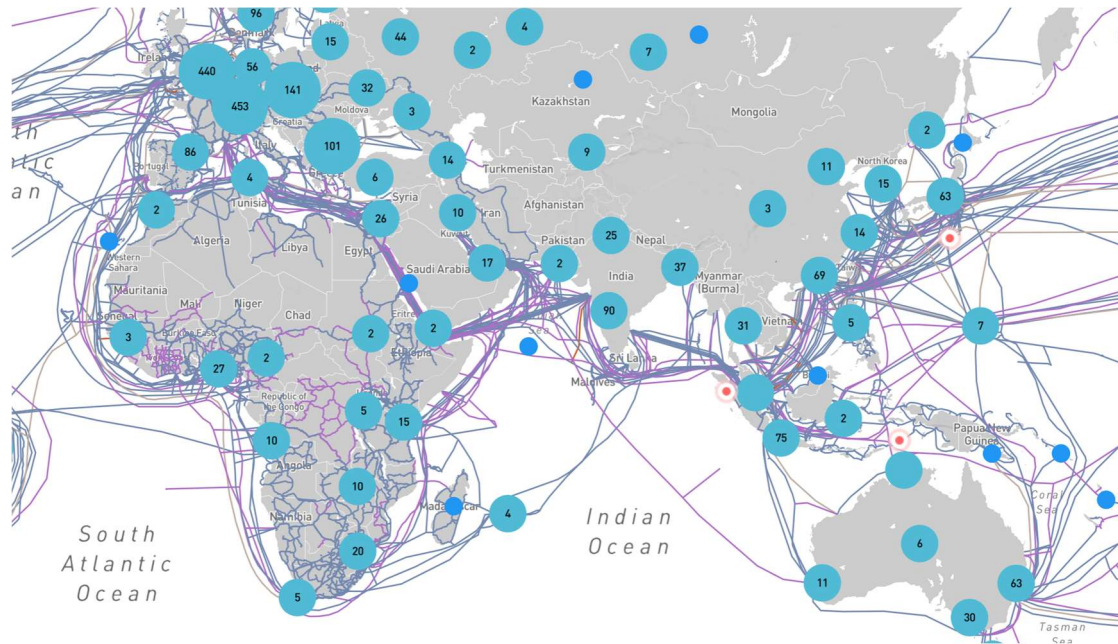


Figure 14. Screenshot from Infrapedia (<https://www.infrapedia.com>) over terrestrial and subsea Internet cables, as well as groups of data centres (the numbers in blue circles). Blue lines indicate active cables; purple indicate planned. The map also shows a few major earthquakes (cf. 5.3.1). Used with permission.

Whilst connectivity is a top priority for a data centre, a prerequisite for its existence is access to electricity. In addition, its cooling systems often need water. In other words, a traditionally built data centre should ideally be located where access to the Internet backbone, electricity and water overlap. Finally, adequate road infrastructure is desired – not least for commodity dehydration of reasonable scale, which this thesis argues can become a solution to data centre energy problematics (see section 6.2).

In other words, where to place a data centre is severely limited because of infrastructure access. As a result, Big Tech and other major data centre owners are forced to make trade-offs to build new units in attractive locations. For example, large-scale data centre development in the USA and Europe has been followed by objections from the local population relating to electricity and water issues (McCammon 2017; Judge 2022a), and Meta recently stopped the construction of the Netherlands' largest



data centre due to extended controversy with the local population over electricity (Moss 2022).

The establishment of an IXP can provide new business activities – also in high-income countries. For example, a transatlantic cable landing station currently being set up in Blackpool on England’s west coast will greatly increase transatlantic data transfer capability and eliminate the need for a roundtrip via London’s cables (Live Blackpool 2021). This establishment is seeking to “*attract more cutting-edge, tech-based industries to the North West coast*” (Live Blackpool 2021).

For an industry so concerned with service reliability (uptime in this case), it is not strange risk avoidance is critical for data centre owners. As a data centre manager said in an research interview: “*these are the most risk averse human beings you’ll ever meet*” (Klemick, Kopits, and Wolverson 2019). Risk awareness binds physical prerequisites and infrastructure access. Acknowledging this bond, the World Risk Index defines risk according to four categories:

- + “**Exposure** to earthquakes, cyclones, floods, drought, and sea-level rise
- + **Susceptibility** depending on infrastructure, food supply, and economic framework conditions
- + **Coping capacities** depending on governance, health care, social and material security
- + **Adaptive capacities** related to upcoming natural events, climate change, and other challenges.” (Aleksandrova et al. 2021)

The World Risk Index’s focus on infrastructure rather than on exposure to natural hazards means the wealthy and geologically stable Europe fares well, whereas particularly small island developing states (SIDS) rank as those facing the most severe

problems. Of the highly populous nations, some low-income countries in Southeast Asia are among the most vulnerable, especially the Philippines (8 in the index 2021), Bangladesh (13) and Cambodia (15). At the other end of the table, one finds wealthy Arabic states, the Nordics and the Baltics (Aleksandrova et al. 2021). One may reflect on that Iceland is among these countries, even though Iceland is built on the plate borders and, from time to time, experiences volcanic activity. Thus, in the modern world, adaption and resilience trump exposure.

Of course, infrastructure such as electrical networks, water networks and Internet cables must also be reliable, preferably for the lifespan of the data centre. It is not always obvious this will be the case, even for the infrastructure-mature European market.

## **5.4 Deployment and operation**

Energy may seem very real. Smil (2017) notes that *“everything in the observable universe can be seen, analyzed, and explained in energy terms”*. However, he also notes that energy, despite this ubiquity, is an *“intellectual construct”*. This is true: energy does not, in a sense, exist, but is a capacity, or, as expressed in chapter 4, a promise. So is money. When it comes to deployment and operation, both activities are tightly coupled to finances, to capital expenses (CapEx) and to operating expenses (OpEx). As elsewhere in this chapter, the challenge for a data centre owner is to find a location with favourable geographical and societal conditions, but also where concerns can be navigated at a reasonable price.

### **5.4.1 Deployment**

As a result of the many data centre types possible and the great variation between them, data centre deployment can be carried out in different ways. A hyper-scale data centre can be regarded a typical industrial building. Security measures are high, and there is a

focus on *connectivity*, to the power grid as well as to the Internet. Erecting the building itself is, however, not a challenging task. Putting a data centre under water is an entirely different experience, especially if the data centre is not close to shore. Whenever the data centre is submerged or brought to the surface, a small seafarer expedition is needed, and so is good weather. As a Microsoft in-house reporter lyrically noted: “*The deployment and retrieval of the Northern Isles underwater datacenter required atypically calm seas and a choreographed dance of robots and winches that played out between the pontoons of a gantry barge.*” (Roach 2020)

In the hypothetical study of coffee-bean drying using data centre waste heat in rural parts of Costa Rica presented in section 7.4, large data centres are not deemed suitable. Instead, what is proposed is to deploy fifty container-sized data centres, which would be enough to dehydrate Costa Rica’s entire coffee bean production, whilst simultaneously increase and democratise ICT access. This scenario presents unique deployment barriers and possibilities. First, for reliable Internet access beyond the data centre, the idea relies on available fibre-optic cable. A less exotic but equally important requirement is access to an electrical power grid, with the possibility to supply an additional and continuous 50 kW per container. Finally, there must be available roads, on which the container can be transported. A reasonable assumption would be to search for towns with decent power grids. Here, the roads would likely be sufficient for the trucks needed to build the concrete foundation, as well as for the transportation of the data centre container itself. Some parts of Costa Rica, but not all, have close access to optical fibre networks. Thus, deciding on where to deploy the fifty Costa Rican data centres would involve the study of maps: Where is high-quality optic fibre access reasonably close? Where are the coffee plantations? In this mountainous landscape,

what does the terrain look like? And not least important: where would the data centres benefit the population the most?

### 5.4.2 Operation

The operation of a data centre requires knowledgeable staff, but despite the vastness of a large-scale data centre, the number of people working within them is not high. In fact, municipalities that have welcomed data centre operators such as Meta or Google have later been criticized for not being able to keep promises of employment opportunities (Vonderau 2021; Sattiraju 2020). Tellingly, despite 24/7 operation, Google's largest data centre employs not more than 400 employees (Johansson and Kriström 2021). There is a reason for this. Once the data centre has been built and its infrastructure installed, actual operation is limited to server maintenance, security and general management roles. Co-location data centres, that is, data centres that let storage and computing on dedicated servers to clients, often have the clients come in and do server maintenance themselves. In such a case, there is little more for the data centre operator to do than to check the client personnel's credentials.

Data centre security can be thought of in at least three ways: a) physical countermeasures to threats, b) prevention of hacking and other attacks through networks and c) guaranteed uptime. Of these three, the first has gained substantial attention from data centre managers and from clients. In fact, a data centre manager met with during this work said that his clients were too concerned about the physical security of the building; software intrusions, ransomware attacks and similar threats are likelier to appear.

Hu shares this opinion, pointing to a data preservation project for a "*time capsule*" filled with digitised information (rather than today's newspaper and the other

usual suspects), called Planets (Preservation and Long-term Access through Networked Services). Though, Hu states, the project was

*“assembled by a sober-minded consortium keenly concerned about the future of preservation – including the British, Austrian, Dutch, and Danish library systems, the state archives of Switzerland, and a number of European universities – they chose to bury their box inside a former Swiss Air Force data bunker guarded by bulletproof checkpoints and twenty-four-hour surveillance systems, electromagnetic pulse protection, negative pressure systems to flush out chemical weapon attacks, and ‘hermetically sealed’ air gaps.”* (Hu 2015)

The time capsule itself is placed in an extra-secure room inside a massive safe, sealed in wax (Daily 2010). Ironically, the data stored in the time capsule are somewhat pointless:

*“The time capsule, they decided, would consist of five core ‘source objects’: the Planets brochure in PDF, a few lines of Java Code [sic], parts of the Planets homepage in HTML, a JPEG picture of the Planets all-staff meeting, and a short MOV video. It also contains a full menu of storage media old and new: punch cards, microfilm, floppy discs, audio tapes, CDs, DVDs, USB and Blu-ray”* (Daily 2010).

One cannot help but wonder if any of these media can be read a hundred years from today. The story goes to show the lengths clients go to regarding the physical security when storing their data. However, the excesses of investments in physical security are not always proportional to the risks of software attacks or theft through networks. Or for that matter, from tampering with by disgruntled employees. As data security researcher Steven J Murdoch puts it,

*“the threat of a rogue system administrator deleting all the data because they know they are about to be fired is orders of magnitude more likely than an EMP [electromagnetic pulse] or invasion. The fact that the hosting centre is under a mountain doesn’t stop the system administrator corrupting the backup files stored there.”* (Daily 2010)

Indeed, the data centre world deals with software, which is ultimately made by and for humans. And even largely content people working within or outside a data centre introduce errors; they’re “only human”. Therefore, as a consequence, data centre risks are more closely related to human error than to materialities. Hu recognises both, but where the emphasis lies is clear:

*“Because of its reliability and ubiquity, the cloud is a particularly mute piece of infrastructure. It is just there, atmospheric and part of the environment.*

*Until something goes wrong, that is. Until a dictator throws the Internet ‘kill switch’, or, more likely, a farmer’s backhoe accidentally hits fiber-optic cable. Until state-sponsored hackers launch a wave of attacks, or, more likely, an unanticipated leap year throws off the servers, as it did on February 29, 2011. Until a small business in Virginia makes a mistake, and accidentally directs the entire Internet – yes, all of the Internet – to send its data via Virginia ... A multi-billion-dollar industry that claims 99.999 percent reliability breaks far more often than you’d think, because it sits on top of a few brittle fibers the width of a few hairs. The cloud is both an idea and a physical and material object, and the more one learns about it, the more one realizes just how fragile it is.”* (Hu 2015)

Physical security is taken seriously also elsewhere. As shown in section 7.5, it can be imperative to actors dealing in co-location. Since security issues are likely to be ranked high on the agenda of CIOs (chief information officers) at client organisations,

a co-location facility must be able to prove itself regarding security measures, to be able to compete.

For other types of data centres, operation, as deployment, poses vastly different challenges. Bringing underwater containers to the surface to replace a faulty server is not a realistic option. Instead, Microsoft decided to leave their Natick data centre be for the duration of the experiment. As it turned out, the strategy was successful: Less than one percent of the servers failed during the five-year operation (Roach 2020). This meant a failure-rate one fourth of what Microsoft experiences in a traditional data centre. Of course, it is dangerous to draw conclusions from just one isolated experiment, but if the low failure-rates can actually be regarded indicative, the company speculates there are three reasons why:

1. The cold environment: the subsea waters outside Orkney were much colder than a typical data centre building.
2. The container was nitrogen-filled. This means corrosion from oxygen was low (or possibly even non-existent).
3. Interestingly, Microsoft suggests the “lights-out” environment might have been helpful as well. As no operators physically interfered with the system during operation, human errors were non-existent. (Roach 2020)

It is still too early to say what made the Natick experiment successful, but it was evident the reliability and longevity of servers lived up to – and possibly exceeded – the company’s expectations.

## **5.5 Waste heat use capabilities**

Heat energy can be recovered in several ways. Doing so requires ingenuity and a clear goal, as it is not always so easily done. In addition, to be able to measure the positive

effect on the environment (and financials), new metrics need to be designed. Heat use options and associated metrics constitute parts of this thesis, and are explored in other chapters, 6 and 8. For the purpose of the completion of the analytical framework, they are also briefly presented in the following.

### **5.5.1 Reuse feasibility at specific site**

Since this thesis circles waste heat recovery, a major matter of concern is the potential of heat energy reclamation at the specific site. The Danish industrial site Kalundborg was discussed above (see section 4.2). For data centres, there must be uses for the low-temperature waste heat that has been produced. Chapter 6 covers some possible applications for this heat, such as direct heating of buildings, greenhouses and swimming pools (section 6.1) and commodity dehydration (section 6.2). With possible use cases for this heat in mind, it is easier to find suitable locations for data centre placement. For example, a beach and spa resort such as Blackpool may use data centre waste heat for its many swimming pools. For Blackpool's ambitions to become a new data centre hub, this should be good news: First, reusing heat makes financial sense. Second, it puts less strain on the environment and that way presumably makes the prospects of major data centre establishment easier for the local population to accept.

In both low- and high-income regions, a little more than thirty percent of all produce goes to waste, adding up to millions of tons of food each year. The difference is that whereas in high-income countries a vast majority of the waste comes from the end-user (the consumer), in many low-income countries about thirty percent of the crops never reach a distribution network. Instead, they are either never harvested, or harvested but eaten by rodents or insects due to poor storing capabilities (Berners-Lee 2021). With a data centre's warm and dry airflow, a substantial amount of this produce can be



dehydrated. Once dried, it is easier to store and transport the produce because of the substantial decrease in size and weight. Obviously, long term storage is also made possible this way, which is beneficial to homes lacking access to a refrigerator. Though this procedure may seem easy in theory, it may not be so straight-forward in practice. A walkthrough of steps needed are presented by Terenius, Garraghan and Harper (2021). Chapter 6 discusses this particular case and some others.

### **5.5.2 Metrics for waste heat use**

Ever since its inception, many researchers and industry practitioners have argued against PUE as a general ballpark metric for data centre energy efficiency – and even more so, data centre overall greenness. The large number of proposed metrics (see section 8.1.2) are clear indicators something is not right. This interest is not strange, given today's huge investments in data centres. Further, morally conscious organisations would look to metrics when deciding on what data centre to choose for their business (at least where they are in control, such as at co-location facilities), and these organisations need more accurate metrics too. It is now overwhelmingly evident (see sections 8.1.1 and 8.1.2) that PUE has played out its role as a general metric for data centre energy efficiency.

Against this background, a new metric is introduced in chapter 8. Due to its qualitative elements, it is bound to lack in precision: under a real and severe climate change threat, it should be recognised what work is in fact carried out with the world's limited computing resources at hand.

The data centre is a factor of high impact in today's societies, and will be an even bigger factor in tomorrow's. Steered by the proper metrics, it is easier to make future

data centres responsible industrial actors, to benefit the environment as well as future societal prosperity.

## 5.6 Society profile

Societies are complex, and made of people. As this thesis is rooted in systems engineering, it mainly discusses communities on the societal level rather than the individual, and focuses on infrastructural solutions. Still, data centres are cornerstones of a modern society. Consequently, this section discusses *society structure, population and culture*.

### 5.6.1 Society structure

A large-scale data centre may be portrayed as a faceless factory, freed from geographical or cultural constraints, and this is not necessarily a wrongful impression: data centres belonging to Big Tech might be built based on a company-designed template, with little concern of where they are erected. Still, data centres have ties to the societies in which they are placed, and must comply with the way the society works.

First, the society needs *political stability*. The data centre industry relies on stability and reputation, and the site at which it is hosted should share these ideals, preferably for the lifespan of the data centre.

There are also matters of *governmental surveillance* and *censorship*. For example, Google has been required to meet government's demands in order to run business in China (Chandel et al. 2019), a situation that is of course undesirable for a company whose business evolves around information sharing. The aforementioned case thus shows the difficulty of running operations in other societal structures. Likewise, to avoid the risk of a political nightmare, Microsoft would likely not put many data centres

in Russia, despite the free cooling available in parts of the large country and the short route to populous regions of China, Russia, Japan, Turkey, Ukraine and India.

Indeed, it is impossible to shy away from policy when trying to set technical solutions in action, since the powers of a nation, of a local authority, of the end-user and of the facility owner itself intersect (Kaijser and Kronsell 2014). As Velkova shows, “*thermopolitical*” tensions may form when data centres scoop up available electricity (Velkova 2021). Conversely, too low electricity prices may hinder the quest for sustainable data centres (Chilukuri, Dahlan, and Hwye 2018). Many other concerns could be listed here. For a data centre owner, governmental policy must therefore be one of the first matters to investigate.

The size of the society may or may not be an issue, depending on the nature of the data centre. If it is to serve the near population (through low-latency applications and networks such as 5G), then it should be dimensioned accordingly. However, for a supercomputer built for weather predictions, where it is placed matters less. Then again, if its waste heat is planned to be used, it must be positioned quite close to an actor in need of the waste heat.

A question of major importance is ICT infrastructure. To decrease network latency, data centre hotspots are usually positioned where the Internet backbone cables connect (Hu 2015). It is further crucial that energy and water capacity is adequate: with uptime as the main concern, an unstable power grid or draught during hot summers are red flags to the data centre owner.

### **5.6.2 Population**

It can be argued whether the population in the vicinity of a large industry building is of any relevance to the successful deployment or operation of it. After all, the huge particle

accelerator in the French town Grenoble, visible from several parts of the city, has likely little to do with the majority of *les Grenoblois*. Still, the presence of a major data centre can be reflected in the local sense of being, and even become a part of the local culture (Vonderau 2019, 2021).

However, there are bigger issues concerning population, at least when the data centre is supposed to support the local society (be it a city, region or country). When scaling the data centre for a population, one should look at its ICT maturity, or as the mobile telephony industry association GSMA puts it, its degree of “*consumer readiness*”. (Note that consumer readiness is the denominator in ICRI, an index for decision-making concerning what markets to invest in. Developed for this PhD work, ICRI is described in section 7.2 and used in section 7.4.)

Perhaps surprising, the low- and mid-income countries of the world have quite high mobile phone penetration. In fact, over five billion people have mobile phone subscriptions today (GSM Association 2020b). Even the world’s ten financially poorest (Somalia omitted due to lack of reliable data) nations have a mobile phone ownership of 37% (GSM Association 2020a), which is a stark contrast to the one-tv-per-village situation not so long ago. In Mozambique, one of these nations, the ratio between SIM cards and population is 52/100 (GSM Association 2020b). Thus, it can be concluded that in every country in the world, data access *matters* and has the ability to play a part in society.

The sustainability aspects of this thesis concern not only the environment but also the social dimension. Vital to social progress, in turn hopefully leading to environmental sustainability, is access to data. Seen in a societal context, data centres are not really about the data – what matters is the service they provide. In a low- or mid-

income country, a data centre can help children attend school online, let a local industry connect easier with distributors, or aid sick persons through online meetings with the doctor where road infrastructure is limited.

Finally, a data centre owner of a major company must today look to population on a global scale. With many of Big Tech’s data centres built in the USA, there should be an unsatisfied demand in other parts of the world. After all, less than one human in twenty lives in the USA. Therefore, but for new data centres positioned strategically around the Globe, data access would only have become more unequal over time. Indeed, the world expects another billion people in Africa and in Southeast Asia by 2040 (Rosling, Rosling, and Rosling-Rönnlund 2018). This growth is visualised in Figure 15.

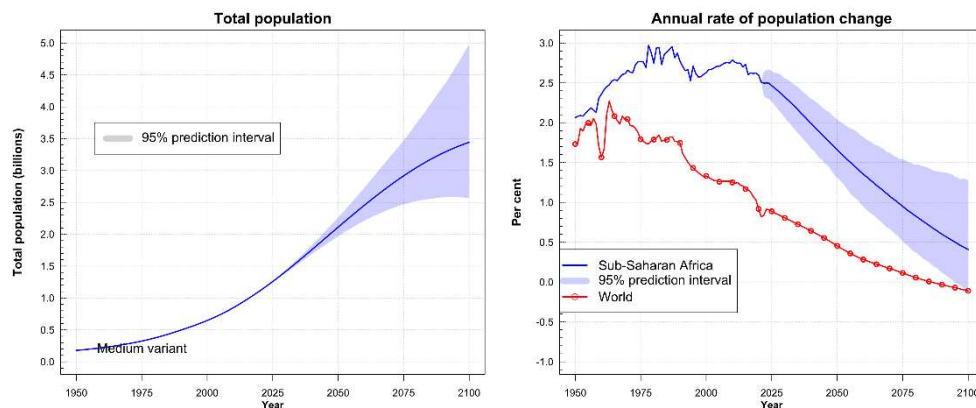


Figure 15. Population growth in sub-Saharan Africa (United Nations - Department of Economic and Social Affairs - Population Division 2022). The left graph shows total population. As GDP rises, people have fewer children (Rosling, Rosling, and Rosling-Rönnlund 2018), but the generational lag and the fact that people will live longer means it will take time before the curve flattens out. The right graph shows population change, proving what was written above: fewer children will be born, and as a consequence, population stabilises. It is notable that by 2100, world population growth is projected to reach zero, despite longer lifespans. Graphs by United Nations, Department of Economic and Social Affairs, Population Division. Published with permission.

Global wealth is incredibly unevenly distributed: *one percent of the global population owns half, and half of the global population one percent* (Shorrocks, Davies,

and Lluberas 2021). With a projected annual GDP (gross domestic product) growth in Africa of more than 4% until 2040 (IEA 2019), this relationship will change. As the African population will be growing fast during the coming two decades, and GDP will increase as well, many more people will get connected. And as the International Monetary Fund (IMF) notes, integrating ICT in Sub-Saharan Africa particularly would bring with it many positive features:

*“The global diffusion of digital technologies promises to create new opportunities for progress and inclusion, so digital reforms and infrastructure will help boost sub-Saharan Africa’s resilience and efficiency, expanding access to global markets, improving public service delivery, increasing transparency and accountability, and fostering the creation of new jobs.”* (International Monetary Fund 2021)

Here it should be observed that since many web services are free (that is, we pay through revealing details of our work and personal lives), low income is not much of a factor for Internet access, and consequently not much of an issue for the data centre industry. That being said, income is a factor for acquisition of the phone or computer needed to access the Internet, for payment of telecom services, and for the ability to purchase goods and services advertised at web sites and apps hosted by the Internet ecosystem, in which data centres play a central part.

Population growth today and tomorrow does indeed mean the labour force is growing stronger. The IMF notes that the working age group of Sub-Sahara (25-64 years) grows faster than any other age group,

*“providing a valuable opportunity for accelerated growth. More than 1½ million people enter the labor force every month ... making the region potentially one of the world’s*

*most dynamic economies and one of its most important markets...*” (International Monetary Fund 2021)

During this time, the population in Nigeria is expected to double, and Nigeria to become the world’s third most populous country. Nigeria already has a strong ICT sector, and is well connected through access to subsea cables. With four hundred million people, its population will outnumber the USA’s. The appearance of local social media platforms is not hard to imagine. So what are the odds that today’s Big Tech continues its dominance fifteen years from now? Probably high, provided the companies manage to tie people in through their massive and deeply entangled ecosystems (such as Apple’s ecosystem made of the iPhone, the iPad, Apple Watch, Apple Pay and other products and services) – which explains Big Tech’s current interest in the African continent (e.g. (Cotterill 2021)).

The current interest is also reasonable given how relatively uncharted the African continent is to the data centre industry. Figure 16 provides a coarse overview of opportunities: here, on a continent with a population already double the size of Europe’s, the number of data centres is very low. Whilst the groups of African data centres fade in comparison with their European counterparts, the difference would likely be magnitudes bigger if the size of each continent’s data centres were to be added up.

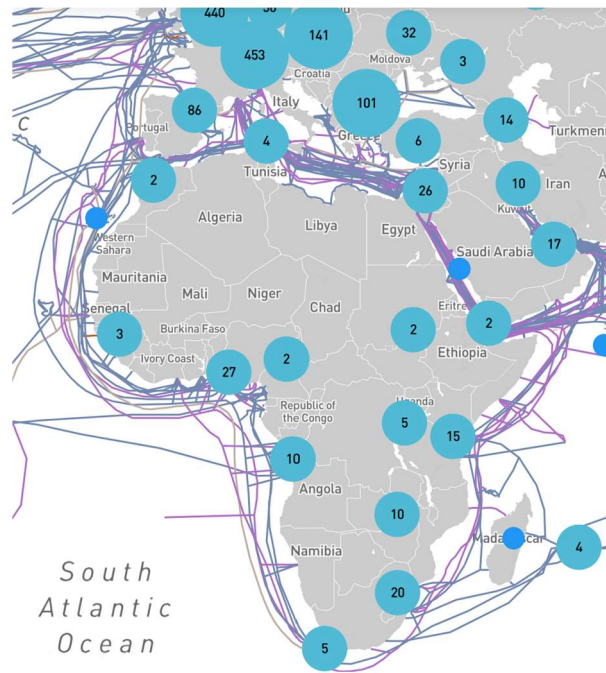


Figure 16. Map of Africa, showing subsea Internet cables and groups of data centres (the numbers in blue circles). Blue lines indicate active subsea cables; purple indicate planned. Used with permission from Infrapedia (<https://www.infrapedia.com>).

In closing, future population growth should be an important aspect of data centre placement. The time is long gone since ICT equalled North America and Europe. With populations that eventually double and with relatively high projected GDP growth, Africa and Southeast Asia are of high interest to the data centre industry, as global players opt for further expansion. The IMF sums up the situation well:

*“The global population is set to grow about 2 billion people over the next three decades. Half of that growth will take place in sub-Saharan Africa ... This trend represents the region’s single greatest challenge, but it also presents perhaps its greatest opportunity: it embodies a vastly growing pool of human talent and ingenuity...”* (International Monetary Fund 2021)

For the industry, these are not empty words: One can only imagine how attractive a billion young people must be to Big Tech and other e-commerce actors, whether they pay with dollars or with giving away personal information.



### 5.6.3 Culture

Culture is a multi-faceted term that is difficult to pin down. Both the Oxford and Cambridge encyclopedias note however, that it is *a way of life*: “*the way of life of a people, including their attitudes, values, beliefs, arts, sciences, modes of perception, and habits of thought and activity*” (Oxford University Press n.d.) (Oxford) and “*the way of life, especially the general customs and beliefs, of a particular group of people at a particular time*” (Cambridge University Press n.d.) (Cambridge). One may assume that “the way of life of a people” is of little concern to data centre placement; after all, the materiality of a data centre is often universal (Dourish 2022). However, as shown in the Swedish case study (see section 7.5), there may be reasons to consider also culture, both as a deciding factor for placement of a data centre, and when figuring out how to run it wisely. In section 7.5, it is argued that on top of an efficient interchange between industry and local, regional and national government initiatives, a genuine interest in nature and a long history of collaboration has aided when establishing sustainability-leading data centres in the Nordics. And as discussed earlier, other regions on Earth may have other cultural traits to take advantage of when aiming to establish a group of sustainable data centres.

On a larger scale, the question of how peaceful a nation is, is of special concern to the data centre builder, as well as local business culture traditions of corruption and similar unethical business practices. Wealthy and infrastructure-advanced nations are among the leaders of the Global Peace Index, and nine of the ten highest rated countries are found in Europe (Institute for Economics & Peace 2021). High levels of perceived corruption haunt many large nations outside Europe and North America: Russia, Mexico and also Brazil come to mind (World Economic Forum 2020). From the index,

it is clear that establishing a data centre hub may be risky business in many parts of the world.

There are other cultural traits to consider when discussing the establishment of data centres. These are due to the fact that culture is an incredibly complex phenomenon, impossible to capture or encircle with today's models offered by systems science or other fields. Instead, here is where visual observation is immensely helpful. The photo gallery below (Figure 17) from a research visit in Kuala Lumpur shows some aspects of Malaysian culture – Ramadan gatherings by the fountains, the consequences of smoking at the underground station, use of motorcycles for transport in rural areas and Malays waiting for the monorail in Kuala Lumpur. These photos all carry meaning, and help an understanding of possibilities and barriers for implementation of ideas proposed in this thesis.




Figure 17. A few facets of culture in Malaysia.

When planning new locations for data centre establishment, one must consider the national and local political climate. This is because for the data centre industry, uptime is crucial, and so is availability of its clients. With policy resistance and new

legislation relating to energy, censorship or foreign information-sharing, there can be threats to both uptime and general availability. Therewith, the data centre industry has political ties. But it has so also in a more profound way. Through social media, media repositories hosted by Big Tech and other actors, and news shared at an unprecedented speed, transmission networks have the power to shake the foundation of political stability. Transmission networks, in turn, utilise and increasingly rely on data centres.

The importance of transmission networks and data centres received particular attention during the Arab Spring in 2011 (Howard et al. 2011). Theodor Tudoroiu, a scholar of international policy, writes:

*“This was a complex transformative process with individual, local, national, and pan-Arab dimensions. Social media served as an instrument of local and national mobilization, communication, and coordination; helped propagate international revolutionary contagion; and contributed to the enhancement of a pan-Arab consciousness which facilitated the contagion process.”* (Tudoroiu 2014)

We who remember the Arab Spring also recall that it was reported as a series of seemingly unrelated events, where *“disparate ‘Facebook uprisings’ [would] take the center of the scene, turning the Arab Spring into a patchwork of erratic and analytically incoherent ‘cascade protests’”* (Tudoroiu 2014). Tudoroiu denies this image stating that in fact,

*“the Arab world witnessed an extremely coherent process of revolutionary contagion whose liberal and democratic ideology was disseminated transnationally mainly by social media. The impressive speed, scale, and effectiveness of this contagion would have not been possible without the effect of the Arab public sphere – itself partially enabled by the social media – on the increasingly cohesive pan-Arab consciousness.*

*Fundamentally, the Arab Spring was the first revolutionary wave ever to reflect the change in power relations originating in the rise of new communication networks.”*

(Tudoroiu 2014)

It can be debated whether other “*new communication networks*”, such as the telegraph or the telephone, have not been able to cause and coordinate uprisings in the past (when *they* constituted the new networks), but if limiting the discussion to social media, it is clear a change has come... and it has come to stay. Through virtual private networks, VPNs, democracy and unwanted media find their way around the world. This massive change to the political landscape ultimately relies on data centres, which is why one cannot discuss benefits and pitfalls of data centres without mentioning their impact on democratisation and, through Wikipedia, reputable news channels etc., also of enlightenment (policy and data centre establishment is discussed further in section 6.6.1, and education through the advent of data centres in sections 9.2 and 9.3). With all this power, one cannot help but wonder if not the anonymity of data centres helps limiting the responsibility of data centre owners: when Lockheed-Martin builds fighter airplanes, or when Germany starts new coal mine in the 2020s, there is bound to be some controversy. But the intangible nature of the principal product of data centres – structured data – makes it easier for data centre owners to hide in the shadows. Based on Facebook’s establishment in Sweden, Vonderau suspects the same, discussing “*the practices of packaging, foregrounding and obscuring infrastructural materialities*”:

*“As I have shown, Facebook tries to limit its own local visibility and to disentangle its infrastructures from local social contexts by framing the social network as merely virtual and global, and by downplaying the role of specific infrastructural sites, in order to ensure the most effective and frictionless flows of data and profit.”* (Vonderau 2019)

It must be said that Facebook was established ten years ago. Data centres are much more well-known today, so there might be less reason to hide. A search at Google Trends (Figure 18) supports this notion: the large data centre expansion has not meant that people search more for datacenters as a topic than what they did in 2004. Indeed, the percentual search volume for the term datacenter is just about static. Assumedly, this has to do with a higher acceptance of these tangible and intangible structures; had many opposed them today, Internet searches for the datacenter topic should have skyrocketed.




Figure 18. A Google Trends search for “datacenter” as a topic. The graph shows search interest from 2004 to today. More precisely, what is shown here is the percentage of searches for the given topic (language-agnostic) on a monthly basis, compared to the number of searches in total. <https://tinyurl.com/tereniusGoogleTrends> (Screenshot from 7 February 2023.)

Thus, the responsibility of data centre owners is not only linked to power but also an important research topic in itself. As Hu notes, cloud computing produces “*a layer of abstraction that masks the physical infrastructure of data storage*” (Hu 2015). This abstraction layer pushes the question of responsibility away from the data centre owner, and the anonymity of the data centre building does so even more.

## 5.7 Chapter conclusion

This chapter has spanned many topics, illustrating the complexity of dimensions of concern for data centre placement. These dimensions, in turn, are evidenced in the

multitude of factors, perspectives and solutions presented above, which together would be under advisement of a decision-maker.

By now, it should also be evident that for a (future) data centre, there is no “one size fits all” solution. Microsoft’s CEO, Satya Nadella, believes in underwater data centres (Ars Technica 2018). The World Community Grid put their hope in continued use of consumers’ unused computation resources (World Community Grid 2021). GRC envisions liquid-immersed servers in existing data centre facilities (Moore 2021). A realistic vision for the coming years might be large-scale data centres for AI work selling heat in cold regions, and smaller, preferably liquid-immersion-cooled, data centres of various configurations for latency-sensitive work in areas closer to populations. Still, the industry continues to invest in large-scale data centres all around the world, following geography-oblivious best-practice solutions to minimise cooling, but ignoring waste heat use and remaining traditional in data centre architecture and placement. In fact, despite the many revolutionary services the IT industry provides, it seems it has a very traditional understanding of what a data centre, actually, *is*.

Section 5.3 provided a few examples of possible locations for large-scale data centres. Connecting to that section, the chapter ends with a summary (Table 2) of advantages and disadvantages of the different locations. Again, there are more possible configurations than shown here. For one thing, liquid immersion removes some of the disadvantages for city-based data centres, and for Big Tech and other data centre-heavy industries, there is room for customised solutions that may affect the estimated value defined in the table.

Still, the table not only summarises the contents in section 5.3 but also much of what has been covered in this fact-filled chapter. In so doing, it lays the foundation for

the Analysis chapter. More importantly, it has hopefully helped to convey both some of the problems and possibilities pertaining to this exciting area of research, and an insight in the complex world they are situated within.

Table 2. Some of the themes and cases presented in this chapter. The Estimate column is ranked, from least to most desirable, as Bad, Problematic, Doubtable, Good, Very good and Excellent. A few other terms are used outside of this ranking: Not applicable, Unknown and Debatable.

| <b>Themes and cases</b> |                 |  |
|-------------------------|-----------------|--|
| <b>Energy needs</b>     |                 |  |
| <b>Location type</b>    | <b>Estimate</b> | <b>Comment</b>   |
| City                    | Problematic     | Also puts strain on the local grids, as evidenced in Dublin and Singapore. However, thanks to the central location, waste heat can often be reused relatively easily.                        |
| Desert                  | Problematic     | Cooling needs are significant, but cooling dry air is less energy-demanding than humid.  |
| Arctic                  | Excellent       | There is little need for cooling. Energy, ideally supplied from renewables, can often be reused to district heating networks nearby, or in their absence to greenhouses, swimming pools etc. |
| Subsea                  | Excellent       | There are almost no energy needs but for server use.   |
| Widely distributed      | Problematic     | All-purpose consumer electronics are not meant for the sort of computing carried out in data centres. Background processes that run on the computers add to the energy use.                  |

| <b>Stability (uptime)</b> |                 |   |
|---------------------------|-----------------|---|
| <b>Location type</b>      | <b>Estimate</b> | <b>Comment</b>  |
| City                      | Excellent       | Access to data centre personnel and likely closeness to Internet exchange points result in high availability.   |
| Desert                    | Excellent       | The deserts in question would be in the relative vicinity of populations, such as Arizona. Thus, the same applies as in cities. Moreover, the weather in a desert may be more stable, meaning weather-related outages may be fewer.   |
| Arctic                    | Excellent       | Stability is high, not least to the cold climate.   |
| Subsea                    | Very good       | According to tests run so far, the stability of underwater data centres is quite high, higher than data centres on land. A downside is that underwater data centres are easy to sabotage and difficult to guard. Sabotage would greatly diminish uptime.  |
| Widely distributed        | Depends         | The stability of a distributed network depends on its configuration. A network that depends on no parts failing would be quite error-prone. However, systems like these (such as the Internet itself) typically rely on nodes to supply certain bits of information (say, a web site in Internet's case). In other words, the fact that one node fails should not have much impact on the other parts of the network. |
| <b>Latency</b>            |                 |   |
| <b>Location type</b>      | <b>Estimate</b> | <b>Comment</b>  |
| City                      | Excellent       | A point of building a data centre within a city is to minimise latency.   |
| Desert                    | Doubtable       | Unless close to major desert cities, latency could become an issue.   |
| Arctic                    | Doubtable       | Similar to deserts. Ultimately, it is a matter of distance to end-users as well as to the Internet backbone with as few jumps as possible.  |
| Subsea                    | Excellent       | Latency should be low, as these data centres should be placed in the close vicinity to large cities such as Seoul, Shanghai or Lagos.   |
| Widely distributed        | Bad             | Latency would typically be high, at least with a global distribution of computation workloads.  |



| <b>Security</b>  |                 |  |
|--|-----------------|--|
| <b>Location type</b>                                       | <b>Estimate</b> | <b>Comment</b>   |
| City   | Excellent       | A building in a city is relatively easy to guard. Law enforcement is not far away.   |
| Desert   | Excellent       | Provided there is ample security personnel in place, and construction security measures taken.   |
| Arctic   | Excellent       | Similar to deserts. Also, the remote and cold location makes raids less attractive.  |
| Subsea   | Bad             | Security concerns would be a major issue. In fact, even though they did not contain sensitive data, Microsoft never disclosed where their two Natick data centres were positioned. Guarding subsea data centres would be difficult and expensive.  |
| Widely distributed   | Problematic     | Security concerns are only software-related, since the whole system is distributed among many actors. On the other hand, these systems would be more prone to hacking since the software is already spread across the globe and accessible to tampering. Moreover, the systems are more well-known than internal systems in data centres, which makes them easier targets for hackers, and more appetising for those who want to hack the systems for the purpose of reputation. |
| <b>Capital and operating expenditures (CapEx and OpEx)</b> |                 |  |
| <b>Location type</b>                                       | <b>Estimate</b> | <b>Comment</b>   |
| City   | Bad             | Data centres are large entities (unless liquid-immersed), and land cost in a city is high. In major cities, salaries are also typically higher than in rural areas.  |
| Desert   | Good            | A major point of building in the desert is inexpensive land use.   |
| Arctic   | Excellent       | Provided needed infrastructure (roads, Internet backbone access) is in place. Strategically placed arctic data centres can take advantage of inexpensive sustainable power, free cooling and waste heat use.   |

|  |                 |   |
|--|-----------------|---|
| Subsea   | Unknown         | CapEx for the Naticks have not been disclosed by Microsoft, but OpEx should be relatively low. It is true it must be costlier to bring up an underwater container than work with land-based data centres. On the other hand, the cold and nitrogen-filled environment, as well as the lack of operators interfering with the system, means reliability and longevity of servers are good.                     |
| Widely distributed                               | Not applicable  | The costs are carried by the consumer. In the case of SETI@home, this was without any financial compensation; people did it altruistically. For bitcoin miners, the point of participation is to be part of the mining lottery. For some “peer-to-peer” (P2P) operations, such as Napster, sharing one’s computer’s hard drive resources was a precondition for gaining access to other members’ hard drives. |
| <b>Sustainability</b>                            |                 |   |
| <b>Location type</b>                             | <b>Estimate</b> | <b>Comment</b>  |
| City   | Depends         | On the one hand, waste heat use is easily achieved in many cities. On the other hand, energy may not as easily come from sustainable sources as remote data centres, which can be built close to renewable power.   |
| Desert   | Problematic     | Cooling needs are obvious. Also, water needs for cooling towers may cause societal and environmental issues related to scarce local water resources.  |
| Arctic   | Excellent       | Besides there being little need for cooling, waste heat can also be repurposed. In addition, the data centre would logically be built where it is sustainably reasonable, for example, taking advantage of clean power from a large river nearby.   |
| Subsea   | Good            | With low (if any) need for cooling, PUE approaches 1. However, waste heat is perhaps not so easily used.  |
| Widely distributed                               | Debatable       | On the negative side, personal computers are not optimised for data centre-typical computation. On the positive side, this idea lessens the need for additional servers, so what is lost in energy may be gained in hardware. Also, some of the computers would likely heat the home or workplace.  |
| <b>Government helpfulness with establishment</b> |                 |   |
| <b>Location type</b>                             | <b>Estimate</b> | <b>Comment</b>  |

|                    |                              |   |
|--------------------|------------------------------|---|
| City               | Bad                          | Erecting a large data centre is often problematic within city walls, due to strain on the local power grid.   |
| Desert             | Good                         | Getting land permissions in the desert may be easy, and large data centre establishments may bring in work opportunities to the local population. Still, there may be questions regarding electricity and water use.  |
| Arctic             | Excellent                    | Governments are quite involved and keen on data centre establishments. Major data centres in the Nordics have brought recognition, money and skilled people to these typically depopulating regions (see section 7.5).  |
| Subsea             | Likely excellent when mature | Provided the underwater data centres are meant to serve the local population, governments should applaud these. After all, they provide improved ICT access, and their land use is close to none. However, underwater data centres will probably need more proof-of-concept installations to be accepted. |
| Widely distributed | Bad                          | It is not deemed probable that local governments would welcome increased energy use for computation mainly serving global purposes. After all, many local authorities have ambitions to lower GHG emissions, and higher use of personal computers does not align with these ambitions.                    |

## 6 Applications for data centre waste heat

Heat is usually considered the least valuable form of energy, as it is challenging to distribute and use. In addition, heat is regarded difficult to store. That being said, as long as the end-use is heat and the distance between source and end-use is adequate, heat energy is useful as an energy carrier. Also, energy storage – regardless of form – is never truly straightforward.

Recognising the challenges above, this chapter investigates existing and possible future uses for data centre waste heat. *But what is it specifically?* Well, the electrical power supplied to a data centre ultimately transforms into waste heat, which is typically vented out. Before reaching the servers, the airflow has been filtered and often dehumidified, to prolong the lifespan of the electrical components.

Industrial waste heat temperatures often exceed 100°C, but a data centre's airflow is much lower, generally 30-40°C. This makes turning heat into electricity a non-viable option, as energy conversion losses would be massive. Hence, in order to reuse the waste heat – that is, 1% of world electricity – this chapter focuses on secondary

uses for dry, clean, lukewarm air. Below, a non-exhaustive list of suggestions for waste heat use is presented. The recommendations of this thesis fall into four categories: heating of other facilities, commodity dehydration, energy storage and increasing work efficiency of renewable energy technologies based on thermal principles.

## **6.1 Heating buildings, greenhouses and swimming pools**

Industrial waste heat has many uses, and is sometimes employed to supply district heating networks with heat. Waste heat use is attractive, as it may offset an equal portion of energy gathered from other sources. The major obstacle in reclaiming data centre waste heat specifically is its relatively low temperature, which is why data centre waste heat use is fairly uncharted territory. Still, the high amount of heat energy – today mostly wasted – makes it worth trying to use it.

Already today, heat reclamation is sometimes employed (Davies, Maidment, and Tozer 2016; Velkova 2016; Huang et al. 2020; Koronen, Åhman, and Nilsson 2020; Wahlroos et al. 2018; Wahlroos et al. 2017). In the Nordics, district heating warms over 50 percent of people’s homes (Davies, Maidment, and Tozer 2016), and accordingly, it is becoming rather common for the Nordic data centre industry to work with municipalities to supply their district heating networks with heat (e.g. (Wahlroos et al. 2017; Vonderau 2021; Koronen, Åhman, and Nilsson 2020)). In fact, many Finnish, Swedish and Danish households are partially heated with data centre waste heat today (Velkova 2016), and the Swedish capital Stockholm alone heats tens of thousands of homes using data centre waste heat by simply leading the waste heat airflow to the district heating network (Koronen, Åhman, and Nilsson 2020; Oltmanns et al. 2020; Wahlroos et al. 2018). To quote Vonderau: especially in the Arctic region, *“ice and cold – symbols of Northern Sweden’s remoteness and harsh climate – turn into assets,*

*necessary to save the data of the world and to make the region profitable”* (Vonderau 2019). Velkova arrives at the same conclusion:

*“The convergence between the data centre industry and urban heating infrastructures is dependent on a well developed and broadly used district heating system. In Europe it is the Nordic, Baltic and some Eastern European countries that have more than half of the population heated in this way and serviced by fibre optic internet, making them potentially attractive locations for such approaches.”* (Velkova 2016)

For cold and temperate regions, there are some other known uses for data centre waste heat besides district heating. Site-specific uses include the heating of water, for fish and lobster farming and for swimming pools (Miller 2008). Amazingly, fifteen years after the first instance of swimming pool heating using data centre heat, the concept still gains recognition, even by the BBC (Kleinman 2023). This goes to show just how little data centre waste heat use is exploited.

Further, heating greenhouses to prolong the growing season in cold regions is an option (for data centres located within city limits, urban farming comes to mind). For small-scale data centres, heating a nearby facility directly may be feasible. A merit of these uses is that heat pumps can be avoided; to avoid conversion losses, need for supplementary power and expensive investment, it is always beneficial to use energy as-is. The next sections present some novel ideas on suggested applications for data centres’ outgoing heat.

## **6.2 Commodity dehydration**

As discussed above, hitherto data centre waste heat has – when at all utilised – primarily been used to heat buildings, which can be seen as unproblematic but also as a little “lazy”. Sometimes, lazy may be fine. That being said, a limited view of what the value

of energy actually *is* – the ability (or, the promise) to perform work – does not help when the surrounding buildings are either not in need of warm air or there is no infrastructure in place (such as district heating) to supply other buildings with waste heat. Thus, to extend the uses for data centre waste heat worldwide, other applications must be imagined, tested and considered.

In both low- and high-income regions, a little more than thirty percent of all produce goes to waste, adding up to millions of tons of food losses each year (Bradford et al. 2018). The difference is that whereas in high-income countries a vast majority of the waste comes from the end-user (the consumer), in many low-income countries about thirty percent of the crops never reach a distribution network. Instead, they are either never harvested, or harvested but eaten by rodents or insects due to poor storing capabilities (Berners-Lee 2021; Bradford et al. 2018). These causes need further investigation, but it is evident the inability to store produce over long periods is an issue here. The underlying idea presented in the following is to use the warm and dry airflow to dry a substantial amount of this produce. Once dried, it is easier to store and transport the produce because of the substantial decrease in size. Obviously, long term storage is also made possible this way, also in homes having no access to a refrigerator.

Besides the possibility to conserve a large part of otherwise spoiled produce, the access to large quantities of dehydrated and safely stored food would help mitigate the effects of periods with risk for starvation. In several places in Sub-Saharan Africa, famine is partly due to lack of planning for the future: if it rains, there is food on the table; if there is no rain, people starve. Compared to undertaking international aid campaigns, it would be a relatively simple task to provide communities with conserved, dehydrated (and then re-hydrated) food when rain fails. Accomplishing this task not

only keeps people and perhaps also cattle from starving, but also lets individuals avoid the health risks associated with lack of nutrients. For very poor communities, recurring periods of drought and locus infestation destroy so much of what they try to build in the years in-between. A steady flow of food is vital for these communities to progress and get out of this problematic cycle. Dehydrated and then re-hydrated food can help these societies maintain a steady flow of healthy food, for comparatively little money.

On the other end of the value chain, much of today's produce is either never harvested, lost during storage or in transit, or rotting at the market, because of the lack of cool and rodent-free long-term storage facilities (Bradford et al. 2018). For instance, Parmar et al. show that Ethiopian sweet potatoes must be sold within approximately one week from harvest, or be discarded of (Parmar, Hensel, and Sturm 2016). If the value chain includes a food dehydration facility as a separate sales channel for farmers, more of the harvest can be used. In turn, this means prices farmers get for produce entering the grocery market would be higher, due to decreased competition. Moreover, since the dehydration facility would introduce flexibility in the system, the gain from farming can be more predictable.

Indeed, with a holistic approach to data centre waste heat and data access, it may to some extent, and under some circumstances, be possible to address problems relating to lack of education, food security and financially and environmentally sustainable communities, as shown in Figure 19. In the figure, a data centre provides a local community with communication abilities and access to data. In addition, using waste heat for dehydration, the community saves up on commodities to build a stock of food for dry periods, and can distribute dried commodities such as coffee, tea, fish or crops, for export and to give a financially viable basis for the community. Finally, energy reuse



increases environmental sustainability, and so does a community's choices, made by citizens more informed through data access. This topic is elaborated upon and exemplified in section 7.3.

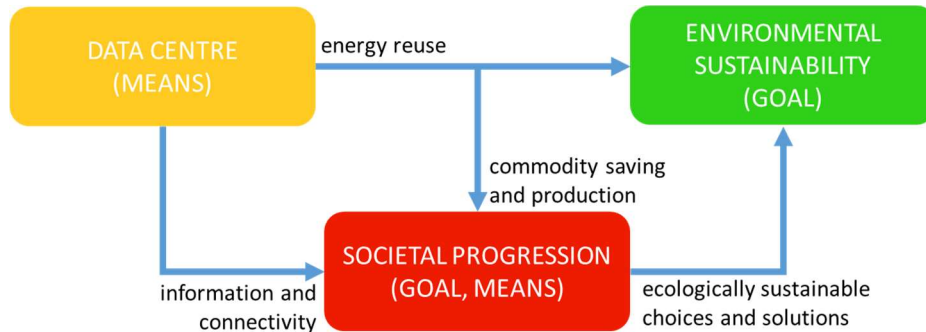


Figure 19. Means and goals to enable sustainable societies. The data centre brings heat and ICT access to the society.

The above mainly referred to sub-Saharan contexts. However, commodity dehydration can become a powerful method to enable societal growth and to combat climate change worldwide through lowered energy use. First presented to the research community in 2020 (Terenius, Garraghan, and Harper 2020), commodity dehydration using data centre waste heat would not decrease energy use, but provide an ability to easily repurpose heat to the benefit of high-, mid- and low-income countries around the Globe. Commodities can be dehydrated on large and small scales, as explained below.

### 6.2.1 Large-scale commodity dehydration

One suitable commodity for large-scale dehydration is wooden pellets from forestry industry leftovers. Once the powdered wood has been dehydrated by the waste heat, the powder is pressed into pellets. The pellets are then used for heating during winter. In Sweden, heat for wood powder dehydration is currently supplied to the Falun municipality by EcoDataCenter to help them become *“the world’s first climate-positive data center”* (EcoDataCenter n.d.), and further north, ambitious research on this topic has been carried out (Vesterlund et al. 2021).

Many other commodities can potentially be dehydrated using data centre waste heat, such as fruit, fish (see Figure 20) or tea leaves for human consumption, fodder for cattle, and seaweed for biofuel production. In so doing, data centres can substitute massive amounts of electricity and fossil fuel used for dehydration today. It is true that in the tropics, many of these commodities are sun-dried. However, industrial dehydration in warehouses increases food security since rodents, insects and birds cannot access the commodity and since the air is clean. Moreover, the produce can be dehydrated in a controllable fashion and more evenly than through sun-drying, and also at night, when sun is absent and ambient temperature decreases. The relatively cool waste heat maps particularly well to commodities risking scorching (fruit, coffee beans) as well as to those where hotter air causes fuel volatilisation (seaweed is one example of such commodities (Skoglund et al. 2017)).



Figure 20. Dried fish sold in a supermarket in Kuala Lumpur. Dried fish (anchovies) is also part of the national dish, nasi lemak.

What to dry is very site-specific. Ultimately, a plenitude of factors decides: available commodities, heat and humidity, road infrastructure, government policy and several others. The principles of large-scale commodity dehydration are found in Figure 21. Here, through a network of actors and activities, farmers become both producers (of data and the commodity at hand) and consumers (again, of data, but also of conserves).

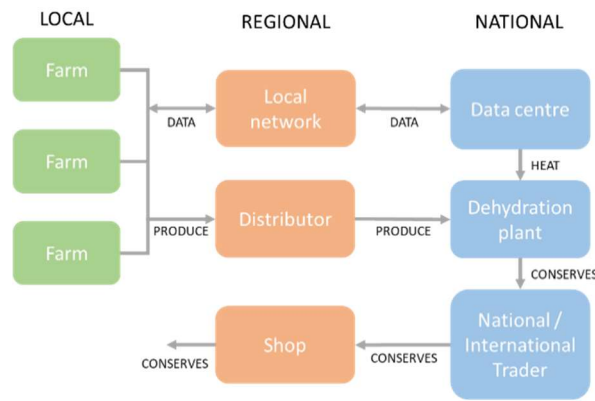


Figure 21. Large-scale commodity dehydration. Typical proposed supply chain structure of an industrialised country.

### 6.2.2 Small-scale commodity dehydration

In addition to addressing the world's electricity use, data access and commodity dehydration are enablers for building local sustainable communities. One particularly suitable commodity for dehydration is coffee, one of the world's most traded agricultural commodities. Coffee beans need drying to achieve moisture levels low enough for extended storage, and to avoid taste deficiencies, the airflow should preferably not exceed 45°C (Phillips 1963). Moreover, coffee is often produced in humid areas, where sun-drying for partial dehydration is not possible. Therefore, drying is commonly (partly or wholly) carried out using diesel-, wood- or electricity-powered machinery.

For an agricultural community away from major cities, the prospects of using data centre waste heat to dry coffee beans or other commodities are appealing. Conversely, suppose an existing drying facility in a community may be powered by data centre waste heat. That may attract international data centre builders for both financial and CSR (corporate social responsibility) reasons, in turn increasing ICT availability locally or regionally.

The concept is illustrated in Figure 22. Here, many farms (in this case, coffee plantations) need dehydration, so they already need to transport their undried coffee to a drying facility. In this vision, a small data centre, consisting of servers in a standard freight container, supplies the dryer with heat. As shown in Figure 19, the dryer and the data access are thus enablers of socially and environmentally sustainable community development. The data centre supports knowledge-creation and increases profits through supplying the dehydration plant with heat. An educated society may also make more sustainable choices.

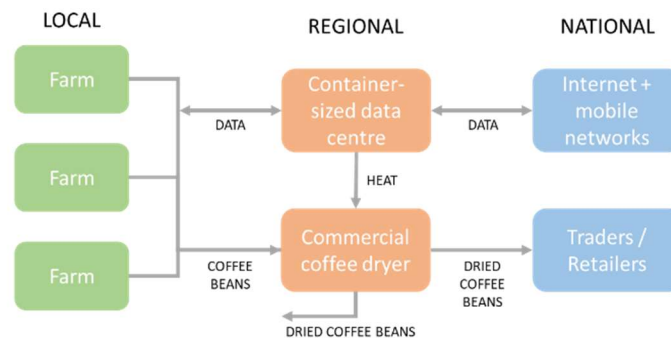


Figure 22. Small-scale coffee beans dehydration. Typical proposed supply chain for a local coffee cooperative.

This approach – elaborated upon by Terenius, Garraghan and Harper (2021) – gives more power to local coffee producers: drying coffee beans close to source and then playing a more active role in supply chains may substantially increase profits for local farmers or collective efforts. Bridging political ecology and engineering, this strategy can be used in many countries around the world, including the leading coffee-producing countries (Brazil, Vietnam, Colombia, Indonesia, Ethiopia, Honduras, India, Uganda, Mexico, Guatemala, Peru, Nicaragua, China, Côte d’Ivoire and Costa Rica).

### 6.3 Energy storage solutions

In an energy supply chain, heat is usually considered the least desirable form of energy, as it is challenging to distribute and use. In addition, heat is regarded difficult to store.

However, there are ways to store waste heat energy economically, safely and reliably, and where the materials in question act as batteries. The dehydrated wooden pellets mentioned above serve as one example of such energy storage, biofuel production from seaweed another.

### **6.3.1 Seaweed for biofuel production**

Seaweed is abundant worldwide. It is used in Asian cooking and for some other purposes, but has potential as biofuel as well. Since seaweed is entirely water-born (thus not taking up space on land), it may come to play an important role for future biofuel production. While biomass – the source of biofuels – does not decrease carbon dioxide in the air, it is carbon neutral: CO<sub>2</sub> emitted during digestion is equivalent to the CO<sub>2</sub> taken from nature during growth. Thus, the more biomass used in the world, the less the need for fossil fuel.

The research and literature on seaweed as a source for biofuel is relatively small and scattered, which is why this thesis involves an experiment related to seaweed dehydration. The underlying assumption was that in a humid setting, seaweed could be initially sundried but to get closer to dry mass would be more challenging. The working hypothesis was that sun-dried seaweed is not dehydrated linearly, but with gradual efficiency decrease (as calculated in percentages of the original weight). Hence, the assumption was that reaching semi-dry commodities would be a much quicker process than the time for semi-dry commodities to reach fully dry state, especially on the coast, where humidity is always high. Provided that seaweed would indeed need increasingly longer times to dry to the point it could be used for biofuel, the dry airflow from a data centre could potentially be of use: initial drying would be done through sun-drying and final drying using data centres.

Commodity dehydration is vital in many of today's equatorial nations. Since dehydrated commodities (such as fish on Zanzibar) provide income to many low-income families, and this without heavy investment, the experiment was designed to rely solely on free and commonly available resources: the sun, the slightly fluctuating wind on a hot summer day, and some wooden frames. The experiment was carried out by the Swedish lake Sommen (58.03° N, 15.27° E), 6-8 July 2021, using available waterweed (*Elodea canadensis*). Wind speeds fluctuated between 0/s and 3 m/s. At the time of waterweed collection, lake surface temperature was 24°C. Clouds appeared and disappeared as they day passed, and darker clouds set in during early evening. The experiment indeed showed that the majority of the moisture content was removed rapidly by sun and wind (see Table 3 and Figure 23). Experiments made by Spanish and French researchers have yielded similar results (López-Hortas et al. 2019).

Table 3. Duration, time of day, temperature, humidity (Rh, %) and batch weight in grams after subtracting dry mass.

| Duration | Time  | Temp (°C) | Humidity | Batch A | Batch B | Batch C |
|----------|-------|-----------|----------|---------|---------|---------|
| 0h       | 11.15 | 24        | 63       | 450     | 453     | 457     |
| 1h       | 12.15 | 26.5      | 59       | 215     | 207     | 259     |
| 2h       | 13.15 | 25.5      | 59       | 154     | 141     | 170     |
| 3h       | 14.15 | 24        | 60       | 122     | 105     | 122     |
| 4h       | 15.15 | 25.5      | 60       | 96      | 80      | 88      |
| 5h       | 16.15 | 27.5      | 48       | 70      | 51      | 60      |
| 6h       | 17.15 | 28        | 53       | 48      | 31      | 39      |
| 7h       | 18.15 | 25        | 47       | 33      | 20      | 26      |
| 8h       | 19.15 | 24        | 58       | 29      | 18      | 23      |
| 9h       | 20.15 | 22        | 68       | 24      | 15      | 19      |




Figure 23. The dissipation of moisture in three batches, seen over time. The y axis shows moisture (water) weight, after subtracting dry mass (circa 50 grams per batch). The x axis shows hours passed. A repeated experiment indicated the same results, but was eventually abandoned due to changing weather conditions.

Thus, though further experimentation is needed, indicative results point to that industrial waste heat should be applied not in a first, but in a second stage. Consequently, when performing calculations for waste heat dehydration of any commodity that can be sun-dried, one needs to understand at what point it is most advantageous to subject the given commodity to the heat, rather than simply calculate dehydration/time to judge the feasibility of dehydrating a specified commodity.

Since biofuel has the ability to power the automotive and aviation industries, energy-efficient dehydration of seaweed might become hugely beneficial. As the initial testing indicated (see Figure 23), certain types of waterweed do not dry linearly. Consequently, in humid environments, seaweed in need of dehydration can hopefully initially be sundried, and when the produce reaches a certain humidity content, the remaining water can be removed with data centre waste heat.

Whether or not this is a viable option is yet to be seen. Even if the idea itself would work effortlessly, there are underlying questions that still needs answering regarding how to most efficiently and profitably convert seaweed into biofuel. What can be concluded already is that seaweed harvesting could generate incomes to shore-based countries around the world. It seems clear that provided that seaweed used for biofuel will need very low moisture contents for biofuel production, data centre-dehydrated seaweed has the ability to generate massive quantities of biofuel, to lower the GHG footprint of the automotive sector and to enable relatively sustainable flying.

### **6.3.2 Salt hydrate charging**

Another option worth further investigation is dehydration of salt hydrates. This technique – suggested for data centre waste heat purposes in 2008 (David Harman V 2008) but not further explored – involves running a warm airflow over the salt hydrate, to remove the water and leave the salt behind. The salt is then packed and transported to wherever there is use for heat. Here, humid air is introduced to the salt and rehydrates it. Doing so releases energy, and a building can be heated. “Charging” thermo-chemical materials such as salt hydrates has several benefits: salt hydrates are inexpensive, easily stored for prolonged times, and non-toxic. Importantly, some of them work at the low-temperate heat provided by data centres (Onder and Sarier 2015; Noël et al. 2022), avoiding the need for heat pumps. As expressed by Salviati et al., *“smart heat management is presently a central topic in greenhouse gas mitigation and the approach of thermal energy storage (TES) has a key role in achieving this goal”* (Salviati et al. 2019). It would be most helpful if data centres could supply this heat, especially in cold and temperate regions with a market for salt hydrates.



## 6.4 Industrial symbiosis with OTEC plants

Both large (Microsoft, Google) and small actors are currently expressing interest in placing data centres near the coastline, or even as floating (Clidas, Stiver, and Hambrun 2009; Nautilus Data Technologies 2022) or submerged (Roach 2020) facilities, to take advantage of seawater for free server cooling. However, the resulting heat – in this case, the heated seawater – goes wasted. One option put forward (Terenius et al. 2021) is running a working fluid such as ammonia or water between a data centre and a dedicated heat exchanger in an ocean thermal energy conversion (OTEC) plant.

In OTEC, warm surface water and cold deep ocean water (DOW) interplay to either evaporate and condense a working fluid (closed cycle or CC-OTEC; Figure 24), or evaporate the incoming surface water, which produces drinking water as a by-product (open cycle or OC-OTEC). In both cases, the evaporated fluid drives a turbine and generates electricity. OTEC works well in tropical and subtropical regions at temperature differences around 20-25°C, which is a lower (better) differential than other heat recovery methods. Optionally, the plant can be designed to produce drinking water as a by-product. One barrier to commercial success for OTEC plants is the low energy conversion potential, about 4% (Vega 2013). The amount of energy converted is directly corresponding to the temperature differential, so if this can be increased, OTEC plants will become more profitable.




Figure 24. The working principle of CC-OTEC (Terenius Dessne 2015). A medium such as ammonia is alternatively in liquid and gaseous forms. Warm surface water makes it boil, and cool deep ocean water turns it back into a liquid. As a gas, the ammonia is run through a turbine, which drives a generator that produces electricity. Figure reprinted with permission.

In this proposition, the heat energy from the fluid which cools the servers (air or mineral oil) is transferred via heat exchangers to the OTEC plant. The plant's working fluid then cools the multi-megawatt data centre through another heat exchanger, as seen in Figure 25. The data centre's electricity is thus reused as heat in the OTEC plant, and the retrieved heat energy helps to power the data centre as electricity. Next, the OTEC plant can then potentially supply a dehydration plant with heat, and in any case serve the community with data access (see Figure 25). A less intricate option is to provide the data centre with seawater air conditioning (SWAC) from the DOW, and then use the waste heat for the dehydration facility.




Figure 25. A tropical, shore-based large-scale data centre, and its industrial symbiosis with OTEC and dehydration plants. A heated working fluid from the data centre feeds the OTEC plant, which provides electricity and (in this version) returns the working fluid, now chilled. In the figure, warm air is sent off to a dehydration facility, though it is also possible the energy from the data centre would serve only one of the other two facilities.

Using OTEC in conjunction with industrial waste heat is not an entirely new idea. For example, research has looked into using heated water from nuclear power plants or gas plants (Soto and Vergara 2014; Kim, Ng, and Chun 2009). It has been found that with the elevated warm water temperatures, not only can the OTEC process be more efficient, but also used in waters that are typically too cold for the technology to be economically feasible. One case study shows that OTEC can be used in the Chilean waters, increasing the net efficiency of a thermal power plant by 1.3% (Soto and Vergara 2014), and another study that in South Korea, the net efficiency of a thermal power plant can be increased by 2% (Kim, Ng, and Chun 2009).

Based on the above examples, it is evident that the area in which OTEC works can be significantly expanded if treating waste heat as a resource. Of particular interest is that with this view, OC-OTEC could be used to deliver drinking water to the Arabic Peninsula, which is home to many large-scale data centres. Here, the water surface is too cool for high efficiency of OTEC technology, but with additional heat from a large

data centre, OTEC would be of greater use. It has been shown that “*there is a market for OTEC plants that produce electricity and desalinated water*” but that OTEC needs scaling up to become more cost-efficient (Vega 2013). Hopefully, the ability to serve the Arabic peninsula with drinking water could be a selling point for OTEC, and rapidly create awareness and a market space for the technology.

The Arabic Peninsula aside, much of the world’s population and financial growth in the coming three decades will take place in Africa and Southeast Asia. Hence, this is where new data centres are needed, and the areas coincide with OTEC’s most favourable locations.

OTEC plants can produce hydrogen gas through electrolysis of the sea water, and data centres can provide these plants with heat energy. Incidentally, it has been suggested that future data centres should rely on hydrogen fuel cells for backup power, to replace today’s diesel backup solutions (Enterprise Ireland 2022). Thus, not only is there a symbiotic relationship regarding electricity between data centres and OTEC plants, but also a supply- and demand relationship regarding auxiliary power. As a bonus, an ocean-based (possibly submerged) OTEC-powered data centre can take advantage of the hydrogen gas as a backup when the plant needs maintenance.

## **6.5 Challenges and chapter conclusion**

For all of the above strategies, some challenges remain. One use of data centre waste heat today is to *heat facilities*. The ideal case would be to heat the building that the (small) data centre is placed in, meaning 100% of the waste heat can be utilised – at least during the cold months. Where district heating is not possible, heating swimming pools, or (pre-)heating water for other uses, can be exciting options.

It is true that regarding *commodity dehydration*, wooden pellets are dehydrated in the Nordics today. Still, to enable commodity dehydration work in tropical or other contexts, many factors must be taken into consideration, such as ambient temperature and humidity, commodities to dry, road infrastructure, political climate and so on, as discussed in chapter 5. To accurately define the monetary and environmental benefit of this strategy, one must also consider the point at which to introduce the commodity to data centre waste heat dehydration.

The greatest challenge relating to *seaweed* is not the data centre-enabled drying process, but finding the ultimate manner in which to use the resource. From conversations with seaweed researchers in the Netherlands, the UK and Malaysia, it seems this research needs much more funding to produce viable results. Given the attractive qualities of seaweed as biofuel and the urgency of moving away from carbon fuel, more funding for this research will hopefully be granted soon.

Storing energy as *salt hydrates* has several benefits, as discussed above. However, as with most energy-saving applications, there is a question of financial viability. For the foreseeable future, the economics of salt hydrates heating seem to push the technology to niche uses. With an increasingly volatile energy market, this may change rapidly. And even if not, one should consider salt hydrates as a solution for homes in remote locations, away from district heating networks. As seen in the United States and in the Nordics, remote locations are often preferred by the data centre industry.

The combination of *OTEC* and data centres is appetising, not least as it may provide a financial injection to the still small OTEC industry. However, the strategy presents two challenges. First and foremost, there is the question of the OTEC industry

itself. The technology has been proven, but due to unfavourable renewable energy policy and a fear of putting substantial funding into sparingly used technologies, it faces difficulties scaling up to the point where it becomes financially competitive. Second, more research is needed to understand how to use the waste heat as efficiently as possible, to avoid conversion losses during energy conversions. This work would involve finding the appropriate working fluid and discovering where in the OTEC process to inject it to achieve optimal results.

Finally, a challenge relating to all suggestions discussed in this chapter is the lack of suitable energy efficiency metrics in the data centre industry. PUE does not consider waste heat use, and the metrics that do include it have been too limited in scope to gain traction. This issue will be discussed further in chapter 8.

Overcoming all these challenges will require more research, within and crossing disciplinary boundaries. With enough resources, the financial and environmental gains can likely be quite large. However, all suggested solutions in this chapter presuppose that there is a need for the waste heat at the given site. Locating the right site for a data centre build and matching it with a well-suited case for heat use is a challenge in itself. This challenge will be discussed in the coming chapter.

# 7 Data centres and society

There are circumstances when using waste heat from data centres is not that simple, nor, for that matter, an obvious solution or even idea. Intuitively, finding uses for waste heat in subtropical and tropical regions may seem especially difficult. Since the world will grow with another two billion people in the coming decades – one in Africa and one in Southeast Asia (Rosling, Rosling, and Rosling-Rönnlund 2018) – and as overall data use is projected to increase (Freitag et al. 2021), a large quantity of the data centres built in the foreseeable future will likely be located in warm regions. Here, using waste heat to heat buildings is not an exciting option. Still, with many gigawatts of data centre power available, the benefits of waste heat reclamation has the possible potential to be enormous – if only there were a use case for this heat. This chapter is about this heat reclamation, and how it can benefit the actors of a society.

## 7.1 The three actors of a society

The societal ties of data centres involve three principal *actors*, commonly referred to as individuals, organisations and states (Meyer and Jepperson 2000). Here, for clarity and relevance, these are named Industry, Consumer and the Authorities. Every actor has its

own desires, concerns and agendas. For every actor, the duality of a data centre – the material presence and the information it handles – is reflected in the relation to it.

### *Industry*

For Industry, a local data centre presents a way to handle company information locally. This may be a requirement from authorities, but working locally may also facilitate work when data operations do not go as planned, since problem areas can be communicated face to face. For some business enterprises, such as stock traders, a short distance and a reduced number of “hops” between the client computer and the data centre is a competitive advantage, and the same may be true for a research institute working with extremely large datasets. The materiality of a data centre may, on the other hand, pose a challenge: the heavy power draw (and sometimes water use for cooling towers) of a large data centre can potentially hinder new industry establishment (see section 5.1.2). Still, and vital to the underlying argument of this thesis, the material output of the data centre – the heat energy that is – can serve other industries with heat for different purposes (and possibly return a chilled fluid in the process).

### *Consumer*

For the Consumer, depending on what kind of data it hosts, a local data centre can provide faster and more reliable access to data, improving everyday life: from schooling and communication with the doctor to selling produce to enable communication with loved ones. An established data centre industry may also raise general IT maturity in a region, thereby aiding its transformation to a more knowledge-based society. The materiality of the data centre means employment opportunities, both for (a few) employees and for construction workers. In addition, heat energy from the data centre can heat homes as well as facilities allocated to consumers, such as swimming pools.



At the same time, the massive appearance of a large data centre, and the noise that its cooling systems make, can become severe problems to the public (Ortar et al. 2023; Diguet and Lopez 2019).

### *Authorities*

For Authorities, closeness to a data centre provides not only faster data access but also potentially a way to keep data under control. This means the government can uphold the law and protect its people by activities such as scanning for hate crimes or detect terrorist attack attempts, but it can also use this control for censorship (Chandel et al. 2019; Warf 2021) or for keeping citizens in check. A large data centre puts strain on the energy grid and may, under some circumstances, threaten scarce water resources (Sattiraju 2020), but behaves much like other industrial facilities in many other aspects. Generally speaking, provided the energy grid can accommodate for it, a data centre should be attractive also to the Authorities, as it can benefit Industry and Consumer.

Table 4 provides an overview of the postulations above, partially based on the SWOT (Strengths, Weaknesses, Opportunities, Threats) model commonly used in the field of economics. Though the list of opportunities and threats is not exhaustive, it brings attention to some of the concerns that should be accounted for when placing a data centre in a local context.

Whilst the goals of the actors may be somewhat similar between societies, their preconditions may not. Therefore, the thesis investigates greatly contrasting cases – high-income countries in cold regions on the one hand, low- or mid-income equatorial countries on the other.

Table 4. Actors with concerns relating to a local data centre placed in a society.

|                   | Industry  | Consumer   | Authorities   |
|-------------------|---|--|---|
| Opportunities     |   |  |   |
| <b>Immaterial</b> | Easy communication brings more control<br><br>Locally hosted data strengthens security<br><br>Extremely fast access to data | Enabling many everyday activities, such as education, telemedicine and keeping connected to family and friends | Keeping control of data   |
| <b>Material</b>   | Waste heat reclamation  | Employment opportunities<br><br>Heat for homes and facilities  | Employment opportunities  |
| Threats           |   |  |   |
| <b>Immaterial</b> | –   | –  | Improved access to data can destabilise a corrupt government                                    |
| <b>Material</b>   | The data centre increases competition over electricity and water  | A large data centre is noisy and its presence dominates the neighbourhood                                      | Needs adequate infrastructure for electrical power, and may put strain on local water resources |

## 7.2 Site selection

In addition to some exploration in the United Kingdom and Switzerland, the three cases of this thesis concern *Costa Rica* – a tropical, agricultural country with no data centres at the studied site, *Sweden* – a high-income country with an evolved sustainable data centre industry and a cold climate in the data centre-rich part of the country and *Malaysia* – an equatorial country with a well-developed economy and a flourishing data centre industry. Table 5 displays similarities and differences between the three nations.

Table 5. Comparison of some key metrics for the three countries.

|   | Malaysia   | Costa Rica  | Sweden  |
|---|--|---|---|
| Country geographies and demographics  |  |   |   |
| <b>Climate</b>  | Tropical/equatorial  | Tropical (though coffee plantations are on high altitudes, thus less hot) | Subarctic (where many of the data centres are placed) and continental |
| <b>Land area</b>  | 329,000 km <sup>2</sup><br>(peninsular Malaysia<br>132,000 km <sup>2</sup> ) | 51,000 km <sup>2</sup>  | 447,000 km <sup>2</sup>   |
| <b>Population</b>   | 32 million   | 5 million   | 10 million  |
| <b>HDI world ranking</b><br>(United Nations<br>Development<br>Programme 2022) | 62   | 58  | 7   |
| ICT landscape   |  |   |   |
| <b>Setting for case</b>   | Major city (8 million<br>in the urban area)                                  | Rural   | Small/mid-size city<br>(80,000)                                       |
| <b>Data centre industry</b>   | Mature   | Non-existent<br>(at studied site)   | Leading   |
| <b>Mobile connectivity</b><br>(GSM Association<br>2020b)                      | “Advanced”   | “Transitioner”  | “Leader”  |
| <b>ICRI (2018)</b><br>(see section 7.2.2)                                     | 0.903  | 0.735   | 0.921   |

### 7.2.1 Rationale for the selection of countries

The three countries were chosen for very specific reasons. Sweden is a leader in sustainable data centres and home to many large-scale data centres, and from what has

been gathered at industry meetings, the Nordics are paving the way for data centre waste heat use. Further, there is some leading research on these topics in northern Sweden. Factoring a cold climate and an impressive economy, the choice of Sweden as a case showed what can be possible with today's data centre solutions, where environmental and political circumstances are close to ideal. As a consequence, solutions from Sweden and other countries of the Nordics (plus Switzerland, which shares many geographical, political and industrial traits with the Nordics) seem to have a chance of being exported to other nations.

Malaysia was chosen for the perceived difficulty of waste heat use. After all, this is a quite hot country, so whereas Sweden would be the near-perfect place for data centre establishment, an equatorial country would be the most challenging, especially for reclaiming waste heat. Moreover, the northern towns in Sweden have rather small populations (less than 100,000 people, and often much smaller than that). Hence, contrasting the multi-million city Kuala Lumpur with the Swedish case could potentially bring both new possibilities and new dimensions of concern to the table.

Rural Costa Rica, finally, was the object of the first study. It was concluded that coffee-bean drying would be a promising case for data centre waste heat use (see section 6.2), but there are fifty coffee-growing countries to choose from. For the first trial case, a country where the idea could realistically be carried out was required. First, the crop yield needed to be substantial for the idea to be worthwhile. This requirement led to the identification of the fifteen largest producers of coffee. Second, the country's infrastructure and policy needed to be rather highly developed, and stable growth was desired, so that the data centre industry would dare to invest. Therefore, the GDP trends of the fifteen largest coffee producers (Figure 26) were examined. Among coffee

producers, Mexico, Costa Rica and China have the highest wealth per capita. Unlike in many other coffee-producing countries, the wealth of Mexico and Costa Rica has been growing steadily for several decades, as Figure 26 shows. Furthermore, Costa Rica is relatively politically stable (Central Intelligence Agency 2023). Hence, though coffee bean dehydration as a principle may be even more beneficial to countries with lower GDP, Costa Rica was considered a likely candidate for implementation of the ideas proposed.




Figure 26. Per capita income seen over time for the fifteen largest coffee exporters, shown in Gapminder. Costa Rica is highlighted, and its stable economic growth since 1950 is clearly visible. Chart made in June 2020. Link: <https://tinyurl.com/tereniuscoffeeexporters>

The justification for Costa Rica was strengthened by a new index. The index is described in the following.

## 7.2.2 Introducing ICRI

Recalling the discussion of actors in section 7.1, it is clear that when implementing a sustainable data centre, one must also consider the end-users. Indeed, the commodity dehydration idea proposed in this thesis is not only about industrial dehydration per se, but also about consumers' access to data. Given the high mobile data penetration in

low-and mid-income countries, data were collected from the mobile network operator organisation GSMA to deduce an index referred to as the ICRI (Infrastructure/Consumer Readiness Index):

$$ICRI = \frac{\textit{Existing infrastructure}}{\textit{Consumer readiness}}$$

(eq 1)

By visualising the gap between existing infrastructure on the one hand and the IT maturity and desire for data access of the population on the other, the point of ICRI is to show where investments in telecom – or any similar infrastructure-related investment – make the most sense. Consequently, the aim is to identify countries, regions or local communities with low ICRI ratios, indicating that in a not too distant future, populations will demand – and have a willingness to pay for – improved infrastructure. A low ICRI is desirable for an ICT systems vendor, as that shows the demand for ICT is higher than the availability. A high ICRI, on the other hand, is an indicator of a saturated market. Of the selected coffee exporters, the country with the lowest ICRI turned out to be Costa Rica (see Table 6).

An interesting feature of ICRI is that it does not discriminate between low and high wealth. For example, Mexico has a much higher level of infrastructure than Costa Rica has, and Uganda a much lower; still they both have significant higher ICRI than Costa Rica.

Table 6. Country infrastructure and consumer readiness. Data retrieved from GSMA (GSM Association 2020a).

| Country       | Cluster      | Infrastructure | Affordability | Consumer Readiness | Content and Services | Infrastructure / consumer readiness |
|---------------|--------------|----------------|---------------|--------------------|----------------------|-------------------------------------|
| Costa Rica    | Transitioner | 56.80          | 54.76         | 77.26              | 64.34                | 0.74                                |
| Colombia      | Transitioner | 58.50          | 49.45         | 76.43              | 71.36                | 0.77                                |
| Peru          | Advanced     | 59.93          | 67.95         | 74.01              | 67.13                | 0.81                                |
| Uganda        | Emerging     | 39.88          | 34.08         | 49.23              | 38.29                | 0.81                                |
| Vietnam       | Transitioner | 59.84          | 64.05         | 73.24              | 63.51                | 0.82                                |
| Honduras      | Transitioner | 52.75          | 39.89         | 62.86              | 51.11                | 0.84                                |
| Indonesia     | Transitioner | 57.39          | 61.04         | 66.10              | 63.12                | 0.87                                |
| Brazil        | Transitioner | 66.35          | 42.83         | 75.21              | 71.43                | 0.88                                |
| Ethiopia      | Emerging     | 33.70          | 41.07         | 36.49              | 33.69                | 0.92                                |
| Mexico        | Advanced     | 66.37          | 60.29         | 70.96              | 72.25                | 0.94                                |
| Nicaragua     | Transitioner | 56.41          | 37.58         | 59.89              | 53.78                | 0.94                                |
| Guatemala     | Transitioner | 55.46          | 49.51         | 58.59              | 58.56                | 0.95                                |
| Cote d'Ivoire | Emerging     | 50.88          | 51.23         | 46.70              | 37.45                | 1.09                                |
| India         | Transitioner | 53.50          | 64.56         | 48.96              | 56.44                | 1.09                                |

### 7.3 Building a sustainable future for Malaysia

As mentioned above, The Nordic countries can take advantage of the cold climate, both for cooling and for waste heat use. The tropics face more challenges. Here, at the centre of the world (more than half of the world population lives in Southeast Asia, India or China), large data centres are built without waste heat use capabilities. They also suffer from much higher cooling costs than data centres in colder climates do. With fast economic growth, growing populations and high energy use, waste heat use would be especially attractive. This section investigates the above-mentioned prospects of waste heat use in Malaysia.

Malaysia, home to 33 million people and – much thanks to oil export – on the brink to achieve high-income status (see Figure 27), has a long tradition of semiconductor manufacturing. This industry started out as a grassroots initiative to bring workers back to a de-populating region of the country (Azzman Shariffadeen, interview Kuala Lumpur 7 May 2022). A quite hot country, where day temperatures seldom go below 30°C, can be regarded a challenge for waste heat use. Still, many of the ideas proposed in the previous section can be utilised also here.



Figure 27. The Kuala Lumpur skyline, seen from the student accommodation skyrise for Universiti Teknologi Malaysia. In the middle are the Petronas towers (Petronas is the government-owned oil company). The much higher KL Tower is visible on the right.

In the late 1980s, ideas were presented on how to transform Malaysia to a knowledge-based nation and to achieve prosperity through ICT (Azzman Shariffadeen 1988). Not long after, the former president Mahathir bin Mohamad built the “Multimedia Super Corridor” in a new town called Cyberjaya, alongside the new governmental town Putrajaya.

Cyberjaya has now become home to many data centre operators. The town is large and spacious (see Figure 28). Hence, adding warehouses for commodity



dehydration would be much easier here than in, say, Singapore, London or Amsterdam, where data centres are abundant, but available land severely limited and prohibitively costly.

A major agricultural producer, Malaysia has a number of commodities that could be dehydrated with data centre waste heat:

- Forestry products – Malaysia is a substantial timber producer, even when excluding Borneo (overseas transportation makes drying forestry products from Borneo in Cyberjaya less favourable). With about seven million cubic metres of produced wooden goods per annum in Peninsular Malaysia (Miyamoto et al. 2014), Cyberjayan waste heat may well be fully utilised solely to dry wooden products such as timber, pellets and construction boards.
- Fish, fruit and vegetables – These commodities often require dehydration. A warehouse solution increases food security and provides more evenly dried commodities. Speaking for data centre waste heat dehydration specifically is the fact that the air inside a data centre is filtered, thus further increasing food security.
- Seaweed – Seaweed has attracted much interest in sustainability circles since it grows extremely fast without intervention and requires no land footprint. Seaweed is typically sundried, but drying decreases over time (see section 6.3.1) and obtaining dry mass with supplied heat is energy-intensive. Therefore, dehydration may be carried out as a two-step process, where initial sun-drying removes most of the water contents, and further dehydration is carried out with data centres, until the moisture content is below 10%, which is appropriate for the food processing industry (López-Hortas et al. 2019). Can the dried seaweed

instead be used for biofuel, it may power the automotive or aviation industries, partially solving several global energy and sustainability problems in one blow.



Figure 28. Cyberjaya’s centre. These particular plots may be reserved for future builds, but the entire area is quite spacious. In other words, it would be easy to place dehydration plants here, as well as to gain access to the sites.

The dehumidified air of a data centre makes dehydration faster and more efficient. Still, with the high outdoor heat of the region in which Cyberjaya is placed, realistically the efficiency gains will be less favourable than in, for example, Japan, Europe or North America. Things may change soon, though. With liquid immersion cooling implemented, outgoing heat can reach 50°C or more, and Qarnot, the French supplier of liquid immersion cooling systems, has managed to deliver 65°C (presentation, OCP European Summit, Prague, 20 April 2023). Increasing temperatures, in turn, increases the number of use cases for tropical data centre waste heat, especially as liquids transport heat much better than air does.

Energy storage, especially hydrogen gas production, is an important key to the future Malaysian economy (WWF Malaysia and Boston Consulting Group 2021). As mentioned above (see section 6.4), ocean-based OTEC has been planned to support this scenario, since bringing electricity to shore from a floating plant far away is problematic. Further, ocean-based data centres can use hydrogen gas when the OTEC plant needs maintenance. OTEC works well in the waters north of Malaysia, and important OTEC research is carried out in the nation. However, due to the shallow and

warm waters in the Malaccan Strait, OTEC is not financially viable close to Cyberjaya and to the majority of the Malaysian population. This fact pushes the question of the suitability of an OTEC-data centre industrial symbiosis to questions about geographical practicalities and the geopolitical worries about conflicts over Malaysian territory in the Chinese Sea. Using the heat energy to “charge” salt hydrates is also viable for large and small Malaysian data centres. Still, doing so implies the salt hydrates must be shipped to a location having a demand for the salt hydrates, such as northern Europe. In turn, Europe has a large share of the global data centre market, so provided this strategy pays off, Europeans could avoid shipping from Southeast Asia and easily supply their own salt hydrates where no better uses for the waste heat are found.

With a large portion of the future world population living in this region, it is likely more data centres will soon be built here, as argued in section 5.6.2. The question is where in this region data centre establishment is optimal. A look at the seismic events map (Figure 29) reveals that the western coast of peninsular Malaysia is well-protected from both direct and tsunami-induced seismic events. Cyberjaya is situated several kilometres from the coast, which further decreases the risk of tsunamis affecting data centre infrastructure. As the map indicates, Kuala Lumpur (with Cyberjaya) stands out among larger cities in this region: Indonesia and the Philippines are located on the Pacific Ring of Fire, and though relatively far from earthquake-prone areas, Singapore is a coastal city, exposed to tsunamis.

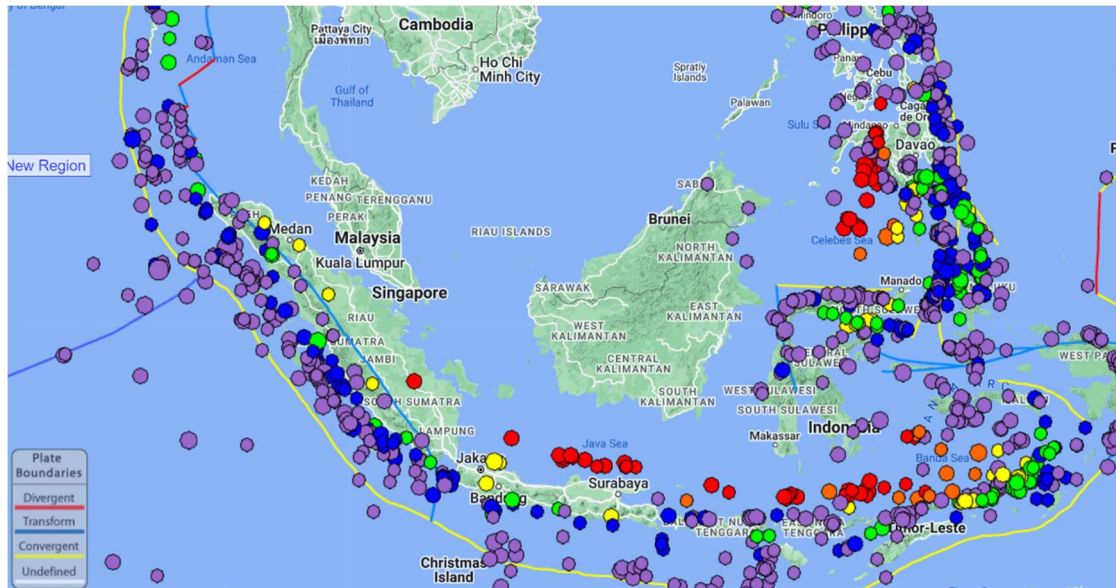


Figure 29. The thousand largest seismic events (earthquakes) since 1970, as shown by the IRIS browser, zoomed in on Malaysia (Figure 13 shows a global view). Software reproduced with permission from the Incorporated Research Institutions for Seismology, USA. Map used with permission from Google. Link to visualisation: <https://tinyurl.com/tereniusearthquakesmalaysia>

The Malaysian government is bound by sustainability pledges, as expressed in several resolutions (e.g. the Twelfth Malaysia Plan, with its emphasis on the “*prosperous, inclusive, sustainable Malaysia*” (Economic Planning Unit 2021)). One way to increase the sustainability of the local data centre industry is to use its waste heat. As expressed in the Green Technology Master Plan Malaysia 2017-2030, “*a sectoral target does not only produce a single isolated outcome within the sector but also has concurrent outcomes in other sectors*” (Ministry of Energy Green Technology and Water Malaysia (KeTTHA) 2017). If the government lives up to these words, the holistic attitude to industries’ contribution to Malaysian prosperity brings hope for waste heat use in Cyberjaya.

## 7.4 Drying coffee beans in Costa Rica

Coffee is the world’s most traded agricultural commodity in need of dehydration. An industry employing sixty million people (Sachs et al. 2019) and delivering almost ten

million tons of produce to the global market yearly, coffee production can have a substantial impact on consumption, values and climate, thus responding to the demands of many of the UN SDGs. This is particularly valid if lessons learned from coffee production are used in other parts of the commodity sector, such as tea and cocoa production, and thereby multiplying the societal and environmental impacts.

Hinted in the above, the main objective of coffee production is income generation. However, as pointed out by the International Coffee Organization (2003), beyond revenue, sustainable coffee also brings immediate results for the local society and environment. According to the organisation, these results include improved natural resource management, biodiversity conservation, improved crop resilience, reduced financial exposure due to on-farm diversification, societal development, increased use of rural labour and, finally, fewer health risks due to misuse of agrochemicals (International Coffee Organization 2003).

In other words, many of the merits identified by the International Coffee Organization mirror the societal and environmental outcomes of integration of data centres in growing societies. Though the list above was written twelve years prior to the advent of the UN SDGs, it also ties to them. When the UN SDG framework was presented in 2015, the fact that sustainable coffee connects to several of the goals did not go unnoticed by the International Coffee Organization:

*“Higher coffee prices are associated with more rural employment, higher contribution of agriculture to GDP, lower levels of poverty (SDG 1), increased food security (SDG 2), reduced inequality (SDG 10), and higher political stability (SDG 16). Hence, policies that help to increase and stabilise income levels of coffee-producing households can have a significant impact on economic and social development ...”* (International Coffee Organization 2019)

Sachs et al. (2019) connect fourteen of the SDGs to the coffee production sector, omitting only SDGs 10, 11 and 14. Terenius et al. (2021), in turn, estimate that their data centre-enhanced proposal for sustainable coffee production in low- and mid-income countries has high or very high contributions to eleven of the SDGs (all but 3, 6 and 14-17). In other words, it seems that with some incitement, the coffee production sector can become an enabler for sustainable communities. Still, that requires a catalyst. *Could – as envisioned by Terenius et al. – the data centre industry be that catalyst?*

Based on ICRI and financial standing, rural Costa Rica was selected as the main candidate. Like Malaysia, Costa Rica is an upper-middle income country (International Monetary Fund 2022). However, many of its five million inhabitants (and migrant coffee plantation workers from neighbouring countries) live under the poverty line of USD 1.90 per day. As argued by Terenius et al. (2021), it may be that for semi-rural areas in many low- and mid-income countries, the future lies in container-sized data centres, integrated in a system of growers, retailers and waste heat users, as shown in Figure 30. Hence, such a solution is used in the calculation below, taken from Terenius et al (Terenius, Garraghan, and Harper 2021). (Because of the quite crude estimates in this example, the figures presented cannot be taken at face value. However, they do show the magnitude of the possibilities for repurposing data centre waste heat in a relatively warm, mid-income nation.)

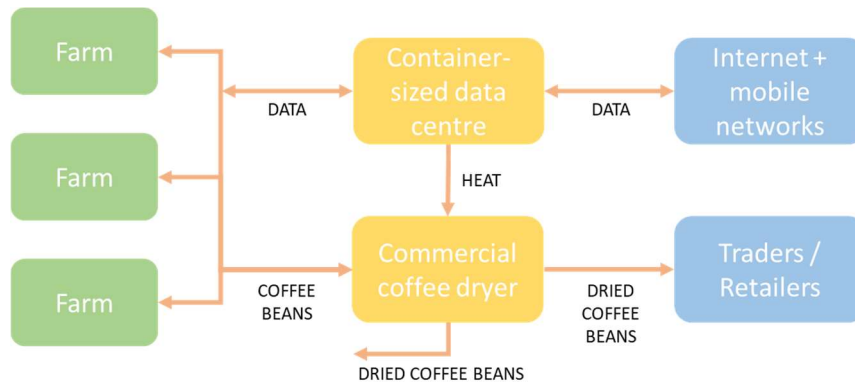


Figure 30. Coffee farms, data centres, dehydration facilities, traders/retailers and ICT actors are all coupled in the proposed strategy (reintroduced figure).

For Costa Rica, about 275 kWh of electricity is needed for every ton of dried coffee (Adams and Ghaly 2007). Firewood typically complements the electrical load, yielding around 1.5 cubic metres of burned wood per ton (Arce et al. 2009). With an annual export of 90 000 tons (International Coffee Organization 2020), Costa Rica thus needs approximately 25 000 MWh and 140 000 m<sup>3</sup> of wood to dry their coffee.

Typical 100% load of a container-sized data centre is 50 kW. Using the data centre as a heat resource, a maximum of 450 000 kWh of waste heat over a year is possible. This results in an ability to dry approximately 1 700 tons of coffee each year, that is, close to 2% of national export. Hence, in theory, 50 container-sized data centres would eliminate the need to power the coffee industry's dehydration processes – whilst providing the population and tourists with access to data.

In practice, the efficiency would be lower: a data centre would not be run at maximum power, and the calculation above assumes 24/7 drying, which is likely unreasonable, also if automatic dryers are utilised. However, as discussed in section 4.4, with new technologies such as liquid immersion cooling, the temperature of available heat would be higher, and their rack density much higher, resulting in much more concentrated heat energy.

Moreover, Costa Rica's harvest season lasts only about five months (depending on weather conditions and region), so for at least half a year, other uses for the heat must be found. However, harvest coincides with the dry season, and there should be many other uses for dehydration during the wet season, not least as coffee workers may turn to other industries at that time.

To learn of the installed data centre capacity of Costa Rica, the seven largest co-location data centres in Costa Rica, all placed in the capital San José, were contacted by email. Based on their replies, the rated power of the co-location market was estimated at 5 MW. The literature has approximated the energy use of co-location data centres to a third of all installed capacity (Barroso, Hölzle, and Ranganathan 2019), again a very crude figure. This means that 15 MW would constitute the entire available data centre power in Costa Rica. Thus:

$$5 \text{ MW} * 3 = 15 \text{ MW total power}$$

(eq 2)

$$50 \text{ kW} * 50 = 2.500 \text{ kW} = 2.5 \text{ MW additional power}$$

(eq 3)

Thereby follows, that if fully implemented, the idea would also bring an additional estimated 10-20 percent of possible data storage and computing to Costa Rica, increasing network capacity and stability. More importantly for the local community, the idea would enable computation closer to the source ("edge



computing”). That, in turn, would increase data speed locally, and lessen the amounts of data in need of shuffling to San José (see Figure 31) as well as abroad.

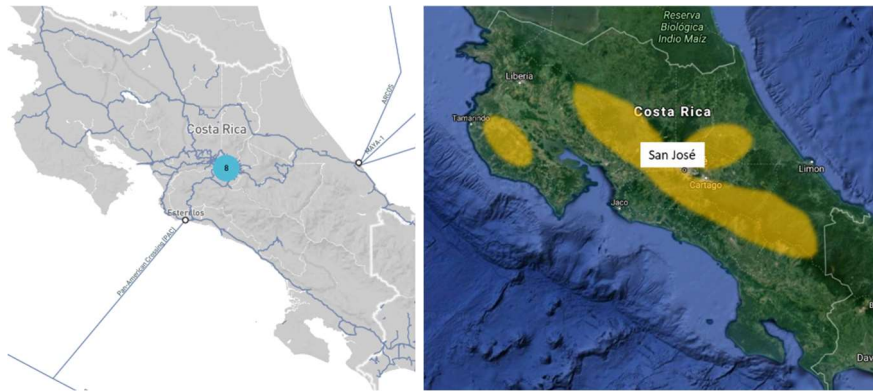


Figure 31. Maps showing various dimensions of concern for data centre placement. Left: Map of terrestrial networks; Infrapedia. Right: Coffee plantation locations (author’s approximation). As can be seen, a large portion of the country lacks powerful ICT infrastructure, and the main computing resources are concentrated to the capital. Infrapedia screenshot used with permission.

## 7.5 Sweden and the new Klondike

The use cases above involved tropical, mid-income countries. However, the material social concerns of the data centre industry also stretches to the wealthier regions of the far north.

Sweden, like its Nordic neighbours Norway and Finland, has some promising conditions for data centre establishment: a climate that enables free cooling, plenty of renewable, stable energy (mainly hydropower), populations with excellent foreign language-speaking skills and a deeply rooted concern for the environment, political stability and practically no censorship, tectonically safe land, a history of practical, innovative solutions to technical problems, well-established road, railway and Internet infrastructure, a collaborative spirit between municipalities and local industries and – like some other countries (Skatteverket 2015) – discounted energy tax for energy-

intensive industries. On top of this, Sweden is built around district heating, making waste heat recovery possible, and often profitable.

A down-to-earth mindset has resulted in many highly practical engineering innovations, from the modern crescent wrench to Bluetooth, and in truth, many internationally known Swedish ventures, such as IKEA, Volvo and Tetra Pak were built on innovative products. In telecom, these include Ericsson, an early innovator of not only the telephone itself, but also of mobile phone networks and 5G. Many big Swedish companies are based on innovations from the late eighteen hundreds – Alfa-Laval, ASEA (today ABB), Atlas Copco and Ericsson were all established between 1862 and 1883 – which might sound worrisome from a business perspective; a prosperous business needs constant re-innovation. As Peter Drucker wrote half a century ago, *“the growth has been largely along lines that had been laid down well and truly in those distant days of our grandparents and great-grandparents”* (Drucker 1969). But building on previous successes can be profitable, and makes for competitive advantages and further expansion. For example, Ericsson developed the first automatic switchboard many decades ago, and is still a key player in telecommunication networks.

A century after the first Swedish tech boom, the Swedish online industry meant a new start for innovation, and spawned enterprises such as Skype, Spotify, Klarna and the game developers King (Candy Crush), Digital Illusions (the Battlefield series) and Mojang (Minecraft). Encouraging data centre establishment is therefore logical, and innovating practical solutions to problems relating to the data centres – read: energy issues – is not far-fetched.

The reasons above speak for Sweden – again, together with Norway and Finland – as a leading future data centre nexus. Interviews among decision-makers (Dybdal

Christensen et al. 2018; Korhonen 2018) further support this argumentation, pointing to many of the factors mentioned above. These countries are also sustainability frontrunners (Burck et al. 2022), and actively researching future prospects for data centre sustainability: fish farming is carried out in Norway, Luleå Technical University and RISE research commodity dehydration in Sweden, and Finnish industry experts have established the data centre Energy Reuse Factor (ERF) as an ISO standard (Suomen standardisoimisliitto SFS RY 2021). Unsurprisingly, the Swedish Data Center Industry Association is happy to share this positive image, and has a desire for integration in societal contexts. Their consortium website uses the tagline “*For a sustainable future!*” and in one of their webinars, a board member stated that data centres “*play an important role, have an undeserved bad reputation and can successfully integrate with energy systems*” (Swedish Data Center Industry Association 2023).

Sweden is to 70% covered with forests. The main rivers used for hydropower are found in the north, a relatively low-populated region. The last half century has seen a migration from the north to the southern part of Sweden, the result of reduced employment opportunities. For the Swedish government, establishing large-scale data centres was a part of a strategy to stop this gradual depopulation (Harding 2015).

Against this background, it is not surprising that the prospects of an established data centre industry in Sweden caught the attention of the government and of foreign heavy data users. In 2011, Meta (Facebook) built their first, thereby “*paradigmatic*” (Velkova 2016), European data centre outside Luleå in northern Sweden (see Figure 32), and Google and Microsoft followed suit. Today, several operators sell their waste heat for district heating (see section 6.1), meaning that data centres transform into tax-

relieved electric heaters with the ability to store and compute data. On top of monetary savings, this method has both societal and environmental benefits. The technical manager at the city-owned energy company Stockholm Exergi explains:

*“Increased digitization gives the city’s residents new opportunities while the heat from the data centers is taken care of and contributes to a more sustainable energy system. You can say that the electricity is used twice – first in the servers, then for heating.”*

(Conapto 2022)



Figure 32. Erecting Facebook’s fourth data centre building in snow, Luleå, October 2020.

Vonderau, who made over 50 interviews with people associated with the Facebook establishment, writes:

*“With IT firms constantly looking out for cheap energy, remote and energy-rich regions around the world compete as potential sites of cloud infrastructure in the hope of solving local demographic and economic problems. The city of Luleå is a case in point. After undergoing a lengthy and highly competitive selection process among more than one hundred places in Sweden and Europe, Luleå won the bid due to its cool climate, low electricity prices and geopolitical security.*

*Equally crucial, however, were attractive deals with regional authorities and companies regarding the data centre’s energy and water supply systems, cheaply sold land, and not least state subsidies amounting to 100 million Swedish Krona [10 million Euro] in total ...” (Vonderau 2019)*

The importance of the collaborative work to get Facebook into Sweden is not to be glanced over. On the contrary, a hallmark for the Nordics is collaboration between industry and the local authorities, and the collaborative spirit also extends to the data centre industry. A glossy and excessively seductive catalogue (see Figure 33), produced in 2020 by the local university, the local regions and the EU European Regional Development Fund, makes this quite clear. The catalogue features interviews with “*the sustainability engineer*”, “*the cloud entrepreneur*”, “*the politician*”, “*the student*”, “*the investment developer*” and so forth, as well as some feature articles. This continuous effort (note that the Facebook data centre was several years old when this promotional catalogue was published) is an indicator of the importance of data centre establishment for the local and regional actors and of the extent to which Swedish industrial and governmental actors of all levels cooperate.



Figure 33. A promotional catalogue for data centres in Luleå and surroundings.

Though this tight collaboration may seem exotic to CEOs from other nations, it is neither strange nor new to the Norrbotten region, to which Luleå belongs. On the contrary, Vonderau writes,

*“...Northern Sweden’s cool climate and cheap hydro power are not the only reasons why this region is a perfect place for locating cloud infrastructure. Of equal importance is the IT industry’s integration in and dependency on already existing local infrastructures that were established in the course of the country’s industrial development over the last two centuries.”* (Vonderau 2019)

Referring to Hansson’s term for Norrbotten as a *“technological megasystem”* (Hansson 2006), she notes that Norrbotten features

*“...an extremely stable, multi-layered industrial infrastructure connecting forest, steel, mining, and other traditional industries with hydro power stations, railways and harbours, as well as large military defence facilities designed to withstand a nuclear explosion.”* (Vonderau 2019)

In truth, such a *“technological megasystem”* would not be possible without tightly integrated affairs and shared goals of the societal actors.

A decade after Meta/Facebook’s establishment, the data centre industry has brought employment opportunities to Sweden, and the once depopulating northern counties are now repopulated, turning northern Sweden into what’s been referred to as a new *“Klondike”* (Myrén 2021). Clearly, despite a per capita wealth four-five times higher than Malaysia’s and Costa Rica’s (International Monetary Fund 2022), Sweden was subjected to the same migratory patterns that haunted Malaysia. And in both cases, the IT industry helped to turn the tide.

## 7.6 Three research approaches to society

The methods used in the Swedish, Malaysian and Costa Rican cases differ, partly because of how the research plan was set up, partly because of the restrictions imposed by covid and partly because of how the understanding of the topic developed. The Swedish case is based on quantitative facts, though visual observation has been considered when trying to pin down success factors for data centre establishment. Malaysia has been approached much more from an anthropologist's perspective (again, mainly through visual observation). Here, there is no history of sustainable data centres, and there is even the question whether there can be such a thing. The only way to understand the topic better was to explore the region, particularly the agricultural regions and Cyberjaya, the home of the Malaysian data centre industry. And just as importantly: to talk with people – researchers, taxi drivers, store clerks – in order to obtain an understanding of the local culture.

As international travel was impossible in 2020 and 2022, any trip to Costa Rica was prohibited. For this reason, the case was conducted much more formally and abstractly than the others, including the creation of the ICRI. Upon proper *in situ* investigation, it may become clear coffee cooperatives in Costa Rica are perhaps not ideal for this idea. For instance, the cost for extending ICT networks to smaller villages may be prohibitively high, and political ambitions may be lacking. It is also possible higher connectivity is, for some yet unknown reason, not deemed very important to the villagers. Thus, despite a low ICRI, the demand/cost ratio may be unfavourable in the specific case. Be that as it may, there are many small agricultural communities worldwide, and with varied data centre configuration other cases can be identified. For these smaller communities especially, it is logical data centres' dual outputs – ICT and

heat – can have major impacts and strengthen societies in profound ways, like newly built train stations in small villages a century ago.

## 7.7 Chapter conclusion

The Malaysian and Swedish cases show that for the local authorities, an established data centre, especially from a high-profile owner, is an attractive goal. The fact that Facebook’s Luleå data centre was placed at the site in competition with a hundred other European towns, validates this finding:

*“Regional actors, for example officials at Luleå Business Association, the municipality’s urban planner, or IT researchers at Luleå Technical University, knew very well that the concept of data centres just a few years ago had been non-existent among the region’s authorities and the local population, and that the idea of financial or moral gain from storing people’s data emerged as absurd until very recently. Now, however, politicians and business developers at regional and national levels were in full swing to push the project forward by mediating between foreign companies, government authorities and the locals, and they tried to prevent any disputes that could slow down the expansion of the cloud.” (Vonderau 2021) (author’s translation)*

In this case, it is clear that ambitious work was carried out to make the establishment happen, with the hope of strengthening the local industry sector. As shown in section 7.5, marketing efforts by the local industry, authorities and university were still ongoing in 2020.

One reason why the sustainable data centre industry is strong in Sweden, Finland and Norway may be cultural factors, in this case, a deep-rooted sense of the need to take care of the environment. (For instance, *allemansrätten* or “every man’s right”, is a century-old law every Swedish child knows about. It gives the ability to roam and even



to camp in every field and forest, as long as one is careful not to not disturb nature.) Providing a better understanding of cultural success factors may, in other words, facilitate further deployments of infrastructure in a specific region, as shown by Porter (1990) and by Yeh-Yun Lin and Edvinsson (2011).

Also for a rural setting with only small-sized industries at hand, as in the Costa Rican case, it seems likely the local government should support the parallel establishment of increased ICT possibilities and free dehydration. Indeed, what ties the three cases is the view that through regarding energy a promise of extended value, society and nature can gain. This view, previously illustrated in Figure 19, is a reminder of the Luleå collaboration, of prosperous coffee cooperatives and of Dr Azzman's ambitious plan to transform Malaysia into a knowledge-based nation.

Whilst there are similarities between the cases, there are also important differences. The substantially differing geographical traits between Sweden, Costa Rica and Malaysia provide one of the reasons the cases in this thesis are taken from the equator and from the Arctic circle. Another is GDP. The World Bank recognises the problem that low- and mid-income countries face in this regard:

*“Because of the extremely high standards of reliability required for data infrastructure, as well as concerns about the carbon footprint of data, the ideal private sector investment climate should provide for reliable, clean, low-cost electricity, natural cooling, and negligible disaster risk – conditions that are not always readily met in low- and middle-income countries.”* (The World Bank 2021)

Reliable, clean and low-cost electricity (hydropower), natural cooling and extremely low disaster risk are all trades of the northern regions of Sweden and Finland, and among reasons why the Nordics have become a popular choice for European data

centre establishment (see section 7.5). Still, as chapter 5 shows, many other factors need investigation. It is clear though that geography and financial stability play crucial roles among the dimensions of concern.

Speaking of finances, one should keep in mind that although Sweden has a much higher GDP than Malaysia and Costa Rica, these are all countries enjoying above-average HDI rankings (see Figure 34). As importantly, the rankings of these nations increase gradually and just about linearly. In other words, these are – relatively speaking at the very least – prosperous and developed countries. Therefore, though three very different environments for possible data centre establishment was aimed for, the differences between them on a country level may not be so dramatic. Still, the cases were not meant to only differ in country wealth and development, but also in prospects for waste heat use. A snowy landscape with large and uninterrupted access to energy differs vastly from mountainous, rural coffee plantations with unreliable power supply, or an equatorial major city with very high temperatures.

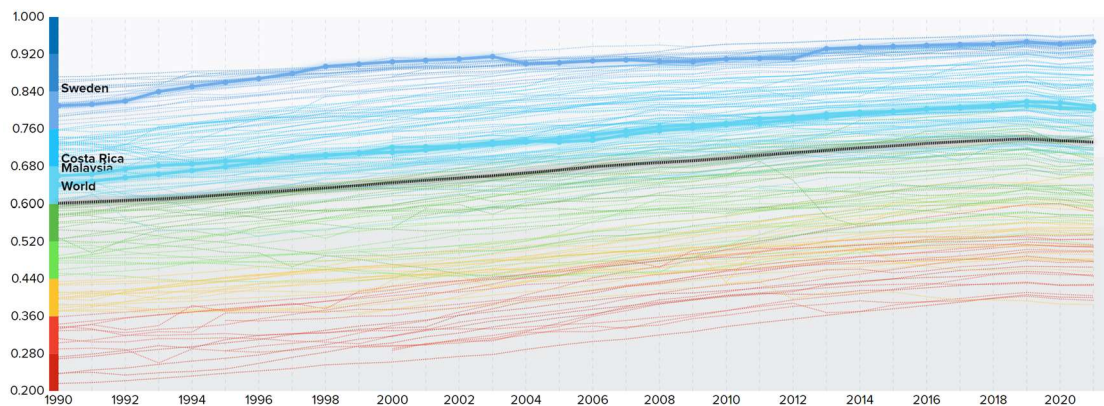


Figure 34. Human Development Index of the countries under examination. Screenshot from 31 March 2023. Link to chart: <https://hdr.undp.org/data-center/human-development-index#/indicies/HDI>

The fact that the many dimensions of concern make cases similar and dissimilar at the same time, indicates that it is quite difficult to find the ideal place for a data centre.

Including the possibility of waste heat use increases the complexity of the decision process further. Hence, there seems to be a need for tools that can make such decisions easier. Of such tools, correct, relevant and helpful metrics seem obvious – but as shown in the next chapter, there is a lack of relevant metrics today. This is why the next chapter also explores possible new metrics to remedy the situation and help stakeholders decide upon proper data centre placement.

## 8 Data centre energy efficiency metrics

One thing having remained consistent throughout this work has been the need for more appropriate data centre energy efficiency metrics. There are two main reasons for this need. First, as has been apparent in metrics-related webinars attended as well as in interviews conducted, the data centre industry itself needs more illustrative metrics to establish market advantages for sustainable data centres. As this chapter will show, a reliance on Power Usage Effectiveness (PUE) – today’s de facto key metric for data centre energy and overall sustainability ratings – is counteractive for sustainability work within these entities. Second, since PUE only deals with data centre internals, the prospects of waste heat use are lost for the society, and whether the data centre is powered by coal or hydropower becomes inconsequential to the rating.

These clear drawbacks of PUE as a holistic metric, and other metrics’ failure to replace PUE in this regard, necessitated the development of new metrics: without them, the proposals for new applications for data centre waste heat would have limited value. In light of these concerns, this chapter presents the history of data centre energy

efficiency metrics, why they fail, and what is needed. Next, a new, holistic, metric is proposed. Last, a real-world case indicates the importance of this matter.

## 8.1 Metrics, used or proposed

This section describes some of the many different metrics proposed for measuring energy efficiency. The aforementioned starting point is PUE. An introduction to the metric follows here, and then some of the other proposed metrics are presented.

### 8.1.1 PUE

Proposed by HP researchers Christopher Malone and Christian Belady in 2006 to deal with the high energy costs associated with data centre cooling (Malone and Belady 2006), PUE was soon adopted by the data centre industry and by other stakeholders (The Green Grid 2010a). Almost two decades after PUE's inception, it remains the dominant metric for data centre energy use. It formalises energy efficiency according to this formula (though this definition from 2007 (The Green Grid 2007) uses “power” rather than “energy”):

$$PUE = \frac{\textit{Total facility energy}}{\textit{IT equipment energy}}$$

(eq 4)

PUE eventually became an ISO standard (ISO/IEC 30134-2:2016). This time, it was defined in terms of energy rather than power, thus establishing the conclusive definition of PUE: “*ratio of the data centre total energy consumption to information technology equipment energy consumption, calculated, measured or assessed across the same period*” (International Organization for Standardization 2016a).

The way the formula is written means that the lower PUE, the better. When all incoming energy is used solely by servers, that is, when no energy is needed for cooling or other side-uses, PUE becomes one. Conversely, everything above 1 is overhead. This overhead is what has primarily been addressed by data centre constructors and engineers – and for some years, this focus made sense as the overhead often exceeded 100%.

One striking drawback of the formula is that it works as a seesaw: efficiency gains in cooling bring PUE down, but efficiency gains in IT equipment (more efficient servers, fewer idle servers, improved memory management in server software and so forth) bring PUE back up again, as shown in Table 7. As a consequence, “*PUE corresponds poorly with energy and carbon efficiency*” (Horner 2016). In turn, this means that – as the Uptime Institute rather diplomatically puts it – to rely “*too heavily on PUE as the industry’s key efficiency metric may reduce operators’ motivation to pursue IT efficiency improvements*” (Davis et al. 2022).

Between 2007 and circa 2014, improvements to the data centre as a facility meant PUE decreased substantially worldwide, from around 2.0 to 1.6, but after this initial decrease, it has proven difficult to lower PUE further (Ascierto and Lawrence 2020). The fact that PUE is flattening out may, to a point, be a sign of a healthy industry sector: efficiency improvements in facility hardware bring PUE down, and – as explained above – server efficiency improvements bring PUE up. This way, energy inefficiency is mitigated from both ways, as shown in Figure 35 and which results in a static PUE. Still, a PUE of 1.6 is quite high (for context, 1.6 is much higher than what the large Big Tech data centres have (e.g. (Google 2020))). To some extent, one reason why world average PUE remains high may be that more data centres are built in the tropics, such as in Singapore, where cooling needs are substantial.

Also Figure 35 illustrates the inadequacy of PUE, when used as a measure of holistic data centre energy efficiency. Despite this misleading notion, PUE is extensively used as an industry benchmark and a leader board-like measure, paving the way for grey-area marketing and greenwashing, and potentially misrepresenting actual data centre environmental impact at a global scale.

Table 7. PUE example. These example figures, from an imagined 1 MW data centre, illustrate a major downside of PUE. When efficiency improvements are made to cooling, PUE decreases (meaning it improves). Any improvements of IT equipment (investing in more energy-efficient servers, undertaking server virtualisation etc) tilt PUE back up again. Since PUE is generally measured over the course of a month or a year, both average power and energy per month data are provided here.

| Power (MW)  |             |                | Energy (MWh/month) |            |                | Metrics     |  |
|-------------|-------------|----------------|--------------------|------------|----------------|-------------|--|
| Cooling etc | Servers     | Facility total | Cooling etc        | Servers    | Facility total | PUE         | Comment                                |
| 0.50        | 0.50        | 1.00           | 360                | 360        | 720            | <b>2.00</b> | Start condition                        |
| <b>0.30</b> | 0.50        | 0.80           | <b>216</b>         | 360        | 576            | <b>1.60</b> | Improved cooling also improves PUE     |
| 0.30        | <b>0.40</b> | 0.70           | 216                | <b>288</b> | 504            | <b>1.75</b> | Improved server efficiency worsens PUE |

In other words, since many sustainability improvements within the data centre have inverse effects on PUE, and since PUE does not account for the source of energy used or possible waste heat use, *PUE cannot be used as an all-encompassing metric for data centre energy efficiency*. Yet, that is exactly how it has been employed for many years, by managers, consultants, climate-conscious clients and politicians alike (Horner 2016).

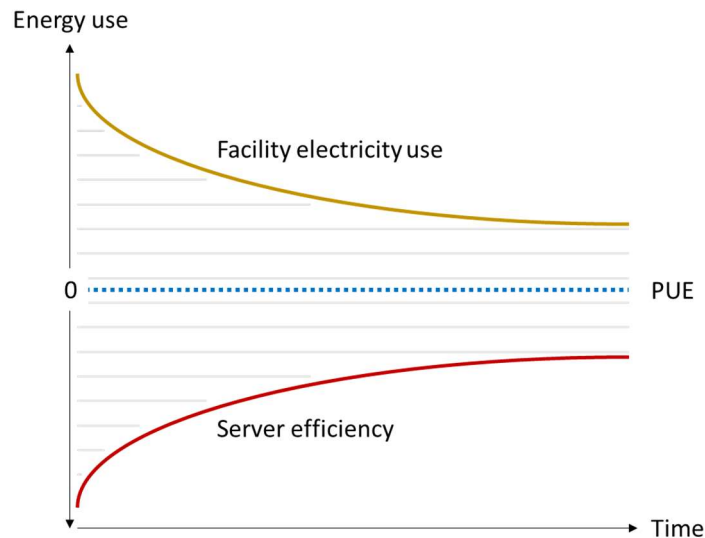


Figure 35. Ideally, the power consumption of both facility and IT equipment become smaller with time. Thus, PUE remains stagnant, although efforts are successfully targeting energy inefficiencies. In case facility improvements are successful, PUE goes down. Conversely, if server efficiency improvements are more successful, PUE goes up. In this illustrative thought experiment, despite significant energy saving measurements taken, PUE remains constant (precisely 2) since cooling and computation use the same amount of energy.

### 8.1.2 The quest for new data centre energy efficiency metrics

Ever since its inception, many researchers and industry people have argued against PUE as an approximate metric for data centre energy efficiency – and even more so, for data centre overall sustainability. The large number of proposed metrics (e.g. (Reddy et al. 2017)) to address the evident limitations of PUE are clear indicators something is not right. This interest is not strange, given today's huge investments in data centres. Further, when possible, morally conscious organisations would look to metrics when deciding on what data centre to choose for their business.

As shown in the above, PUE is inadequately equipped to make holistic measurements. Still, rather than employing other metrics, PUE is used as a yardstick throughout the industry, and unfortunately, also by politicians. As a matter of fact,



Amsterdam placed a cap on data centres already in 2008, based on their PUE ratings (Laan 2008). Singapore, home to many data centre providers in Southeast Asia, followed suit a few years after (Swinhoe 2022). Today, also China bases planning permissions on PUE (Industry and Information Division 2021), and in May 2023 (vom Dahl 2023), Germany's Federal Cabinet approved the draft for a law to do the same (Bundesministerium für Wirtschaft und Klimaschutz 2023). In most if not all of these instances, the maximum allowed PUE value is set to 1.3.

Just how problematic PUE is as a general sustainability metric can be illustrated by Google's sustainability reports (Google 2020), which use yearly PUE rating improvements as a token of successful sustainability efforts. With an increased emphasis on power-hungry AI computation the processors will run hotter, in turn *improving* PUE because of the increased energy use. In other words, the more power Google directs to its servers, the better their sustainability rating.

Several researchers have recognised PUE's path to a misrepresentation of Big Tech's sustainability efforts. Crawford – herself employed at Microsoft Research – partially acknowledges the sustainability work carried out by Big Tech, stating that some corporations “*are responding to growing alarm about the energy consumption of large-scale computation, with Apple and Google claiming to be carbon neutral (which means they offset their GHG emissions by purchasing credits) and Microsoft promising to become carbon negative by 2030*” (Crawford 2021). However, she strongly criticises Big Tech's energy use, and ties this problem to a more deliberate strategy to hide the industry's environmental problems:

*“The tech sector heavily publicizes its environmental policies, sustainability initiatives, and plans to address climate-related problems using AI as a problem-solving tool. It is*

*all part of a highly produced public image of a sustainable tech industry with no GHG emissions. In reality, it takes a gargantuan amount of energy to run the computational infrastructures of Amazon Web Services or Microsoft's Azure, and the carbon footprint of the AI systems that run on those platforms is growing.” (Crawford 2021)*

One point Crawford does not bring to the table is *how* the industry can achieve this misrepresentation. The short answer is – PUE. As shown in the above, PUE rewards heavy compute, which makes it an adequate metric for an AI-centric company's sustainability report, but not for sustainability endeavours. Since data centre configurations and purposes differ, it is further difficult to accurately describe power and energy use. For instance, a data centre mainly used for long time data storage would have other demands and opportunities than one dedicated to cryptocurrency mining. In this case, a stakeholder using PUE as a sustainability indicator would likely incorrectly report the cryptocurrency facility as more sustainable of the two, due to its high IT equipment power usage.

### **8.1.3 Some proposed metrics for data centre work**

Some of the drawbacks of PUE listed above are well-known in the data centre industry and in academia. Therefore, quite shortly after PUE was introduced, alternative metrics began to be suggested to tackle the issues with it. Below are a couple of examples of metrics aiming to calculate useful work within the data centre – metrics that if successful would be able to directly substitute PUE, but which have proven to be surprisingly difficult to gain traction.

### 8.1.4 CPE, Compute Power Efficiency

Also proposed by Belady and Malone (2007), the unit-less CPE is an early attempt to recognise servers' work carried out, rather than their power consumption. The metric was criticised as early as 2008, because of the unclear use of "*utilisation*" (The Green Grid 2008). CPE is formulated this way:

$$\begin{aligned} CPE &= \frac{IT\ Equipment\ utilisation}{PUE} \\ &= \frac{IT\ equipment\ utilisation * IT\ equipment\ power}{Total\ facility\ power} \end{aligned}$$

(eq 5)

### 8.1.5 DPPE, Datacenter Performance Per Energy

DPPE is based on the idea of comparing data centre work with the sustainability of its energy use. The metric was introduced by the Japanese industry in 2011 (Ueoro 2011), and aims to show energy efficiency for the entire data centre, without being overly complicated. However, energy reuse is not accounted for.

DPPE is expressed as

$$DPPE = \frac{Data\ centre\ work}{Data\ centre\ energy - Green\ energy}$$

(eq 6)

DPPE was introduced alongside the Green Energy Coefficient (GEC): Green energy generated on site divided by the data centre's total consumed energy.

### 8.1.6 DCeP, Data Center Energy Productivity

Introduced by The Green Grid in 2008, the focus of DCeP is to define the actual usefulness of a given data centre (The Green Grid 2008). Thus, it needs useful work defined, which is a problematic task. Despite this difficulty, DCeP is close to being an all-encompassing data centre metric. According to The Green Grid’s definition, “*productivity is the quantity of useful information processing done relative to the amount of some resource consumed in producing the work*”. DCeP is formulated this way:

$$DCeP = \frac{\textit{Useful work produced}}{\textit{Total data centre energy consumed producing this work}}$$

(eq 7)

### 8.1.7 FLOPS and other server energy measurement metrics

Attempts to measure server work include Floating Point Operations per Second (FLOPS or flop/s). As the name says, FLOPS calculates how much compute can be done over a second for a certain device or a number of devices. This means that instead of using energy as an indicator of work carried out, one records actual compute directly. However, FLOPS says nothing about how this work has been distributed between server management work and operations for the task at hand, and floating point operations are only useful to calculations dealing with very large or small numbers.

Developed by the Standard Performance Evaluation Corporation (SPEC), the benchmarking Server Efficiency Rating Tool (SERT) measures the energy use of servers. Like FLOPS, SERT has potential as a benchmark, but its use is obviously limited to just one part of the data centre energetics system.

Already in 2012, Segó et al. compared DCeP with FLOPS and Energy-Delay product (EDP), another metric designed for the server level. For data centre energy efficiency measurement, they deemed DCeP superior due to its focus on realistic scenarios, and that transposing FLOPS and EDP

*“to the data center level is problematic at best. If synthetic workloads are used, their performance is difficult to correlate to the performance of actual workloads at the data center level because the latter tend to be a function of multiple interdependent performance aspects, such as processing performance, memory capacity and speed, and network and Input/Output (I/O) bandwidth and latency. Contributors to data center energy use, such as transformers, Uninterruptable Power Supplies (UPS), and data center lighting and cooling might not be properly considered. Consequently, it would be inappropriate to attempt to measure energy productivity at data center level using these types of metrics.”* (Segó et al. 2012)

As the quote shows, data centre energy measurements is anything but a straightforward process, and the gap between measures of server performance and of the data centre as a facility needs more work to be overcome.

### **8.1.8 Other IT equipment energy metrics**

There are also metrics with greater ambition regarding the data centre as a physical entity. One of these is Power Density Efficiency (PDE) (Lajevardi, Haapala, and Junker 2014), which aims to bridge calculations on server performance with server rack airflow, or in other words, bridge computer science and mechanical engineering. Expanding on this idea, a discussion on aligning metrics for thermal management with server performance specifically is found in Capozzoli et al. (2014). There are some more general metrics for measuring ICT energy efficiency and similar matters. For example, how to optimise the Energy Efficiency Ratio (EER) of electrical and possibly other

supporting equipment has been discussed by Kulshrestha and Patel (2019). Other researchers have examined the chain of hardware and software energy use, ranging from processor and server power models to network interfaces, cooling and machine learning (ML) strategies (Dayarathna, Wen, and Fan 2016).

### 8.1.9 Contextualising metrics

There are metrics of higher interest to this work, metrics that put the data centre in a larger context. As this thesis is mostly concerned with waste heat use, a metric of high interest here is Energy Reuse Effectiveness (ERE), a collaborative effort of The Green Grid, the National Renewable Energy Laboratory (NREL) and the Lawrence Berkeley National Laboratory (LBNL) (The Green Grid 2010b). Though not formally spelt out in their white paper (The Green Grid 2010b), ERE can be formulated this way:

$$ERE = \frac{\textit{Total facility energy} - \textit{reused energy}}{\textit{IT equipment energy}}$$

(eq 8)

Though it incorporates the amount of reused energy, not even ERE provides a complete picture of energy reuse matters. For example, it does not indicate transportation losses or conversion losses or the value of the use case. From a financial perspective, it fails to acknowledge that the heat energy retrieved could have been gathered from somewhere else. Introduced together with ERE, the Energy Reuse Factor (ERF; ISO/IEC 30134-6) is the amount of reused energy divided by the data centre's total energy.

In order to look beyond the data centre, and to view the data centre as one part of a sustainable community, one also needs to address issues such as water usage, GHG emissions, the degree of energy reuse, mineral extraction and fair work in the supply

chain and other matters. In terms of environmental and social sustainability, as well as industry financials, these aspects may in fact have more impact than savings made within the data centre's walls. One such metric – Compute Carbon Efficiency (CCE) – concerns GHG emissions from the data centre. Mirroring the PUE formula, CCE is calculated this way:

$$CCE = \frac{GHG[Total\ facility]}{GHG[Compute\ capability]}$$

(eq 9)

Similar metrics exist, such as *Carbon Usage Effectiveness* (CUE; ISO/IEC 30134-8), which is the ratio between the data centre's and its IT equipment's GHG emissions. The *Renewable Energy Factor* (REF; ISO/IEC 30134-3), finally, is the “ratio of the renewable energy ... owned and controlled by a data center to the total data center energy consumption” (International Organization for Standardization 2016b).

From a holistic perspective, it should be noted that none of the metrics presented up to this point, REF included, incorporates the quality of the electrical power going to the data centre. How sustainable is it... *really*? Even when REF is 1, one may wonder whether there could have been better societal use for the energy.

Another problem with REF relating to the argumentation of this thesis, and recognised by the ISO (International Organization for Standardization) (2016b), is the underlying definition of renewable energy. The ISO documentation states – with reason – that criteria to “categorize an energy source as renewable can differ among jurisdictions based on local environmental or other reasons” (International Organization for Standardization 2016b). In other words, REF values can fluctuate

depending on policy decisions. For instance, ever since the EU in 2022 declared fossil gas sustainable power (Hancock 2022), European data centers powered with gas have wrongfully enjoyed substantially improved REFs. REF's off-setting option is also problematic. In the exact words of the ISO REF standard, "*a data centre that receives electricity entirely from a coal-fired plant can purchase RE [renewable energy] certificates to off-set the entire electric use and these certificates are included as RE within the calculation of the REF*" (International Organization for Standardization 2016b).

## **8.2 Introducing the Datacenter Energy Sustainability Score**

A virtue of PUE and ERF is their readability: the formulas can be easily explained and remembered. However, they have a limited view on data centre energy. To address this situation, it is useful to first consider what data centre energy matters actually involve: the promise of work being fed forward. In addition, extreme complexity sometimes makes easily measurable metrics less useful, and that is the reality for the data centre industry and its external stakeholders. In fact, Reddy et al. (2017) conclude that none of the more than hundred metrics they investigated allow for just comparisons between data centres: there are simply too many variables to consider. One solution to both of the above problems is to grade data centre energy concerns (see Figure 36) individually, and add them together.






Figure 36. The three main concerns of energy use relating to a data centre.

Reflecting these three concerns and recognising the issue with overwhelming complexity brought up by Reddy et al., the *Datacenter Energy Sustainability Score* (DESS) has three parts, presented in the following.

#### *EnergyIn*

This ratio is the amount of energy derived from renewable energy sources divided by total energy. It is not to be confused with renewable energy per se: a data centre using up a society’s available renewable energy cannot be regarded sustainable. When sustainable energy is plenty, the ratio approaches 1. That being the case, also sustainable energy sources come with different environmental and societal footprints (see section 4.1). Therefore, *EnergyIn* would always be lower than 1.

#### *EnergyCompute*

This part of the formula is, for now, 1 minus PUE overhead. As stated above, PUE is not an apt metric to show energy use within a data centre, or the usefulness of that work. Future endeavours will hopefully complement PUE with a metric that instead specifies the amount of useful work carried out per energy used. However, a challenge would be to define “useful work”; a computer engineer might think of floating point operations per second or kWh and a sociologist of computation that somehow benefits a society. Moreover, also metrics such as floating point operations per second (FLOPS) and the benchmarking Server Efficiency Rating Tool (SERT) have their limitations: they only see to server efficiency and are thus the opposites of PUE.

*EnergyReuse*

This unit indicates the amount and degree of usefulness of heat reclaimed outside of the data centre for a given purpose, such as commodity dehydration in a certain societal context. ERF is similar, but omits the contextualised usefulness of EnergyReuse (see below). Put differently, the difference between EnergyReuse and ERF is chapter 7 of this thesis.

Each constituent has a few members, as shown in Table 8. DESS is formulated this way:

$$DESS = 3 * EnergyIn + EnergyCompute + 2 * EnergyReuse$$

(eq 10)

As seen from the definition, the constituents are not graded equally. This is because of adherence to the circular economy principles. As a starting point, the incoming resource (electricity in this case) should be sustainably obtained. EnergyCompute relates to the Reduce principle, and waste heat use to Recycle and Reuse. Of these, EnergyCompute is seen as least important, since if all energy is reused, the climate impact is low and the variation of data centres is not that great (meaning that if ungraded, EnergyCompute would have too much weight). That being said, an energy-inefficient data centre puts strain on the climate due to unnecessary hardware usage and that hardware's associated embodied energy. As will be shown in the next chapter, the factors 3 – 1 – 2 provide a crude but efficient way to represent these concerns.

The three constituents of DESS all have values between 0 and 1. For convenience, energy is suggested to be measured in MWh rather than in kWh or the SI unit Joule, as DESS is primarily meant for medium-sized and large data centres, using

megawatt hours or gigawatt hours over a year. The metric's exact constituents and their weights will need further investigation and collaboration between different parties, such as the societal actors and academia. The goal here is not precision: rather than advocating for exactness, DESS was designed to transform a highly engineer-centric worldview into a material social. This transformation, in turn, necessitates interventions with actors from disciplines outside engineering.

Like PUE, DESS relates to energy use only. Therefore, similar metrics may be needed, not least for GHG emissions (Datacenter GHG Emission Sustainability Score, DGESS), land use (Datacenter Land Use Sustainability Score, DLUSS), water use (Datacenter Water Use Sustainability Score, DWUSS), and social issues (Datacenter Social Impact Sustainability Score, DSISS). Weighed together, they would create a scorecard solution (Datacenter Overall Sustainability Score, DOSS). It might be that DOSS should come in two versions, one for the operational sustainability score (cf. OpEx) and one that also includes the embodied energy and other environmental costs in servers and the data centre building (cf. CapEx + OpEx over its lifespan).

As societal complexity is too high to be precisely captured in a formula, accreditation bodies must evaluate the sustainability for the specific case, based on both quantitative and qualitative assessments. This kind of holistic solution has been tried before. In fact, Greenpeace made an ambitious attempt with regards to Big Tech data centre sustainability assessment already in 2012 (Greenpeace International 2012). A holistic rating may also be the only way to avoid what Whitehead et al. (2014) call "*pollution shift*" – moving sustainability-related issues between components to achieve good ratings. However, despite the many existing metrics, PUE has endured throughout the last decade, since scorecard solutions have their own set of problems. First, there is

a risk a scorecard solution leads to fuzzy borders. Second, its many constituents makes a score card difficult to decipher. PUE, on the other hand, was never fuzzy nor unapproachable.

DESS may be regarded a compromise between PUE and a true score card. Thanks of its scope (worldview), audience, explainability and limited parameter set, there is hope that DESS will do as a starting point for a more usable metric – and chapter 9 shows some promise.

Table 8. Constituents of DESS.

| Constituent              | Unit   | Explanation  | Formula   | Typical value range |
|--------------------------|--------|--|---|---------------------|
| <b>EnergyIn</b>          |        |  |   |                     |
| total energy             | MWh    | amount of energy, regardless of source                                   |   | 40–1 000 000        |
| sustainability rating    | 1-10   | how sustainable the energy is, scopes 1-3 included                       |   | 3-9                 |
| total sustainable energy | ratio  |  | $\frac{sustainability\_rating}{total\_energy}$      | 0-0.9               |
| <b>EnergyCompute</b>     |        |  |   |                     |
| used by IT equipment     | MWh    | (for now) calculated as in ISO PUE                                       |   | 60-75% of PowerIn   |
| Overhead                 | factor |  | $\frac{total\_energy}{used\_by\_IT\_equipment} - 1$ | 0.1-0.8             |
| <b>EnergyReuse</b>       |        |  |   |                     |
| Reuse                    | %      | percent of heat energy reused  |   | 0-90                |
| degree of usefulness     | %      | how useful the heat energy is, compared to equivalent from other sources |   | 30-95               |
| proper reuse             | MWh    | properly reused heat energy  | $total\_energy * reuse * degree\_of\_usefulness$    | 0-30                |
| proper reuse             | ratio  |  | $reuse * degree\_of\_usefulness$                    | 0-0.98              |

### 8.3 PUE, ERE, ERF and DESS – a comparison

This chapter has presented a number of metrics, used or proposed. The focus areas of some of the main metrics are shown in Figure 37. This section compares a few of these metrics.

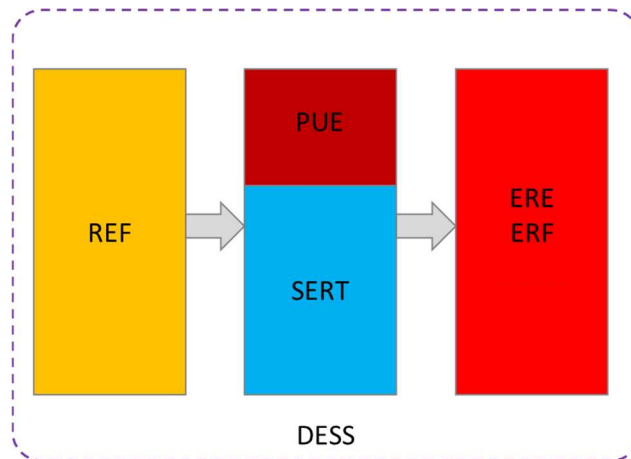


Figure 37. Focus areas of some of the metrics pertaining to data centre energy.

As the overall sustainability key metric PUE certainly has its limitations, in that it cannot cover all aspects of data centre energy use. However, though PUE may be too blunt an instrument for comparison of data centres' power efficiency, it seems to be good at what it does. Quoting The Green Grid, the ratifying body for the PUE metric, *“PUE is an excellent metric for understanding how well a data center is delivering energy to its information technology equipment”* (Avelar, Azevedo, and French 2012). What lies outside the building is unaccounted for, severely impeding the use of PUE as a universal metric.

In other words, if used holistically, it is evident PUE fails as a key metric. ERE and ERF aim to incorporate energy reuse, which is of course benevolent. However, they do not say anything about the *quality* of the energy reuse, and unmonitored and unaccredited energy reuse may have undesired consequences. To quote the ISO ERF standard: *“Processes that take advantage of the reused energy for other uses are outside the data centre boundary and the benefits of that energy and the efficiency of the reuse process are not considered in the ERF”* (International Organization for Standardization 2021). This has implications. For example, to comply with waste heat use requirements (such as those formed in the coming EU Energy Efficiency Directive presented in the

next section), we may end up with swimming pools in isolated places, and data centres may incorrectly report waste heat delivered to district heating systems as full reuse, also when the heat is rejected by the district heating owner. Therefore, ERF's position as an ISO metric is worrisome.

DESS was developed as a reaction to the drawback of a reliance on very specific metrics. Thus, it was designed to particularly appeal to external stakeholders, such as governmental decision-makers at all levels. To politicians, the combination of PUE, ERF and other metrics may seem opaque and impenetrable, thus useless. To the engineer, DESS may appear too blunt and imprecise for serious consideration, at least for day-to-day work. Next chapter will illustrate the difference between PUE and DESS for a multitude of cases.

With time, DESS may hopefully evolve into DOSS (Datacenter Overall Sustainability Score). DOSS would be another key metric, more fluid in its design – and metaphorically speaking, in its *worldview*. Introducing several properties – some of which are qualitative – it may incorporate water use, what source power originates from, what sort of work the servers are doing, where waste heat ends up, and what end-users actually do with the computing power.

The major argument against DESS is its imprecision, in turn, an effect of its qualitative aspects. However, these seem to be needed. Lucivero writes that “*environmental harm is still uncertain and difficult to measure*” (Lucivero 2019), and that for big data and data centres,

*“current measures of environmental impact assessment may fall short when asked to provide a definitive answer or a clear-cut weighing of benefits and drawbacks. This complexity demonstrates the need for an ethical and philosophical reflection on impacts*

*that takes account, not only [of] quantitative aspects, but also of qualitative, behavioural and moral dimensions.” (Lucivero 2019)*

One can use Lucivero’s text to compare DESS with the LCA-focused metric proposed by Tozer et al. (2018), mentioned in section 2.1. But for DESS’s added recognition of waste heat use, they would span the same area. However, their underlying philosophies differ: whereas DESS is holistic both in its coverage and through a recognition of the contextualisation of energy, the other proposed metric is bound to a mechanical engineer’s worldview.

Lucivero’s expanded view (or – put differently – her call for a material social perspective) makes for another reflection. In the introduction to this thesis, it was mentioned that for centuries, the world has been regarded a cornucopia for humanity to use as she pleases. With the data centre industry’s substantial energy demands and GHG footprint, as well as its reliance on rare earth materials, maybe today should be the day when we reflect on what we actually use the world’s computing resources *for*.

## **8.4 The EU Energy Efficiency Directive**

In relation to decision-making of society authorities, the relevance of the metrics presented above may feel as only moderately important. Such a notion would be wrong. In reality, data centre energy efficiency metrics – and other metrics relating to data centre sustainability – are quite high on the political agenda today.

Given their status as ISO metrics, PUE’s and ERF’s inability to truthfully reflect a data centre’s environmental footprint and contribution to societal growth is concerning, since this is what they are essentially used for today. Though not overly detailed in their composition, they only provide partial truths (if that). Still, they are

included in what is now being proposed in the EU – the EU Energy Efficiency Directive (EED) (European Parliament 2023c). The directive is meant to speed up energy savings within the EU generally (European Parliament 2023b), and it particularly pinpoints data centre energy. In fact, the 200-page-document uses the term “data centre[s]” fifty times.

At first glance, what is requested seems reasonable. The addendum stresses a few reasons for the mandatory reports – to report on sustainability-related matters, to raise awareness (and, implicitly, to drive change within management teams and supply chains), to monitor change and to enable (authorities’) decision-making:

*“The data centre sustainability indicators should be used to measure four basic dimensions of a sustainable data centre, namely how efficiently it uses energy, how much of that energy comes from renewable energy sources, the reuse of any waste heat that it produces, the effectiveness of cooling, the effectiveness of carbon usage and the usage of freshwater. The data centre sustainability indicators should raise awareness amongst network operators, data centre owners and operators, manufacturers of equipment, developers of software and services, users of data centre services at all levels as well as entities and organisations that deploy, use or procure cloud and data centre services. It should also give confidence about the actual improvements following efforts and measures to increase the sustainability in new or existing data centres. Finally, it should be used as a basis for transparent and evidence-based planning and decision-making.”* (European Parliament 2022)

Pertaining to the data centre sector, the objective of the directive is to “*result in a considerable reduction of the energy and water consumption, an increase in systems’ efficiency promoting decarbonisation of the grid or in the reuse of waste heat in nearby facilities and heat networks*” (European Parliament 2023c). Interestingly, waste heat use is given unprecedented attention – data centres with rated power of at least 1 MW



will need to use their waste heat “*unless they can show that it is not technically or economically feasible*” (European Parliament 2023c). This ought to provide loopholes in reality, though the underlying idea of the EU EED is fantastic as an ambition.

At the time of this writing, the directive and instructions for measurement have not been finalised, but as it seems, every data centre in the EU with a rated power of some hundred watts or more is expected to deliver detailed specifications concerning many of the metrics dealt with in this chapter: PUE, ERF, REF, CUE and WUE (European Parliament 2022). According to the Uptime Institute’s Director for Sustainability, Jay Dietrich, metrics needed from the data centre industry will be more detailed than this (Dietrich 2023). Some, he concludes, are likely requested to deepen the EU’s own understanding of data centre energetics.

It is impossible to say something definitive about the successes and problems relating to the EED, as the legislative process is ongoing. Will it be a blow in the air or will it become what Dietrich proclaims “*is likely to be the most important legislation yet for data center sustainability and efficiency reporting*” (Dietrich 2023)? With its dependency on underlying metrics of dubious value, and with a disregard of societal complexity, the EU EED is benevolently meant, and it may well result in similar legislation in North America a few years from now. In fact, new legislation for the USA is currently being drafted, called the “Federal Data Center Enhancement Act of 2023” (Rosen (sponsor) 2023). However, the EU EED is nonetheless problematic: its detailed view of complex tasks may not be the right way to address climate change or society concerns.

The currently implemented German Energy Efficiency Law (the EnEfG-E for short) not only sets maximum allowed PUE value but also a minimum value for heat

reclamation (between 10% and 30% for future data centres, depending on when the data centre is being erected) (Bundesministerium für Wirtschaft und Klimaschutz 2023). How this will work in practice is still unclear, and only time will tell how data centre builders can work with this requirement. Whatever the fate of this national law, it is probable that waste heat use will be expanded to the EU EED and later to North America. This assumption is based on something repeatedly mentioned at industry gatherings: that Europe leads the way in data centre waste heat use, and that the American industry monitors European waste heat use cases and EU legislation closely.

As a holistic metric, DESS was designed for an overarching goal: creating just societies that use energy responsibly. With this philosophy, data centres should be built where it makes environmental or societal sense, say, in a region that is a net exporter of low-GHG energy (e.g. Luleå) or where it is plausible that establishment would bring societal advantages in the form of IT-based enterprises, improved data access or the like (e.g. Cyberjaya, Costa Rican rural grounds, or, again, Luleå). In contrast, there is a risk that the EED and the EnEfG-E will result in a focus on measurement for measurement's sake. For example, based on the metrics mentioned in the above, a data centre powered by fossil gas and featuring dubious waste heat use may still be regarded sustainable in the eyes of the EED, and if so, nothing has been gained from its implementation. Thus, it is clear that sustainability efforts are needed not only to define metrics, but also to understand when – and when *not* – to use them.

## 8.5 Chapter conclusion

This chapter is an effort to show the urgent needs for new metrics, and it should be read in light of previous chapters' concerns for climate change and other sustainability challenges. With regards to data centre sustainability metrics especially, something

needs to be done, since the metrics must be better at showing decision-makers and other stakeholders which data centre configuration is preferable given a specific geographical and societal context. To quote Horner, “*the industry’s past focus on PUE is misplaced*” (Horner 2016), and it is clear the EU tries to address this problem.

Recognising environmental and financial concerns, many have approached data centre energy metrics from an engineering perspective. This research, as a contrast, has been conducted with also the non-engineer external stakeholder and the public in mind. Here, accuracy is more important than precision, which is why a singular, holistic metric can be more useful than a whole set of detailed metrics that say rather little. For the politician, a single metric is easier to discuss, and easier to explain to the public. Thus, such a metric promotes not only sustainability but also democracy.

Despite a focus on precision, holistic metrics, being (to cite Carveth Read) “*vaguely right*” rather than “*exactly wrong*” (Read 1914), also increase the quality of the decision-making process. Two natural science researchers concerned with flawed environmental policies, Orrin Pilkey and Linda Pilkey-Jarvis, agree with this view, stating that

*“if we wish to stay within the bounds of reality we must look to a more qualitative future, a future where there will be no certain answers to many of the important questions we have about the future of human interactions with the earth.”* (Pilkey and Pilkey-Jarvis 2007)

The two scientists also note:

*“Even in the unlikely event that sometime in the remote future we will understand each of the numerous parameters and their interactions and feedbacks that control the event we are predicting, we shall never accurately predict the future. No one knows in what*

*order, for what duration, from what direction, or with what intensity the various events that affect a process will unfold. No one can ever know.”* (Pilkey and Pilkey-Jarvis 2007)

Pilkey and Pilkey-Jarvis speak from a geologist’s perspective, where uncertainties are abundant. For example, one can see an increased risk of an earthquake or volcano eruption, but it is impossible to know when it will appear. Similarly, few seem to have anticipated the energy crisis that plagued Europe in 2022, even though signs of the military escalation that eventually led to the war in Ukraine and the fossil gas shortage that followed had been visible for some time.

For conducted experiments in a controlled environment, say, for simulations on server energy use, it seems that the pursuit for exactness may provide an adequate approximation of the truth. Seen on a conceptually higher level, such as in societal contexts where variables are too many to enable usable calculations, an emphasis on exactness instead displaces ambitions and *veils* the truth. In the aim to – as Pilkey and Pilkey-Jarvis put it – “*stay within the bounds of reality*”, detailed specifications and measurements are not what we need. In their stead, three pinches of energy quality, one dash of data centre internals and two sprinkles of heat reclamation may provide a way forward.

# 9 Analysis

This thesis has focused on the prospects of waste heat use and its associated energy savings (optimistically hundreds of terawatt-hours each year), as well as on its societal value. However, a data centre obviously has more value to a society than heat energy only. The societal value – derived from careful integration of data centres in communities – means stable access to information, in turn potentially enabling or facilitating activities such as off-hours school work, e-health, improved farming procedures, online commerce and upholding social connections. Moreover, the waste heat is not only an environmental gain but also a societal: access to dehydration facilities makes possible a stable food supply and economic growth through sales of dehydrated commodities.

A matter that needs acknowledging is that societies are built around people, and that there is a big difference between providing means for change and making changes happen. Indeed, as shown by Montiel and Delgado-Ceballos (2014), we cannot achieve a sustainable future with new technology alone. In fact, with time, we should be able to overcome inequality and solve the problem with, for example, rare earth metals and e-waste. However, even with these problems solved, we will still rely on a *consumption-*

*based* society, which by the very definition puts a strain on the climate. Thus, as pointed out many times in research (e.g. (Shim and Bellomy 2018; Waas et al. 2014; Van Der Leeuw 2020; Murray, Skene, and Haynes 2017; Danielsen Sorup et al. 2019)), the complexities of sustainability matters need holistic solutions, but establishing “*smart sustainable cities*” (Bibri 2018) is far from enough. Surely, we need to take advantage of technology, but also of policy and, very patiently, of change in people’s mindsets and habits.

Figure 38 shows the number of people in the world divided by their daily income, as a logarithmic view. As can be seen, the world is anything but binary, and the world may seem a more hopeful place than one would intuitively have believed. Even so, eleven percent of the world’s population still live in extreme poverty. The simulation shows that by 2040, that is, at the time when data centres built today will reach their expected end-of-life, this will have changed for the better: even with vastly increased populations in low-income countries, the eleven percent will decrease to seven. For some of this decrease, we should probably thank ICT.

Undoubtedly, ICT is a key enabler of future sustainable societies. That said, ICT can also work against them, through skyrocketing energy use, mineral resource depletion and social injustice. It is, therefore, crucial that both data centre industry and policy makers make the best of the unprecedented intelligence-gathering power the data centres supply us with.

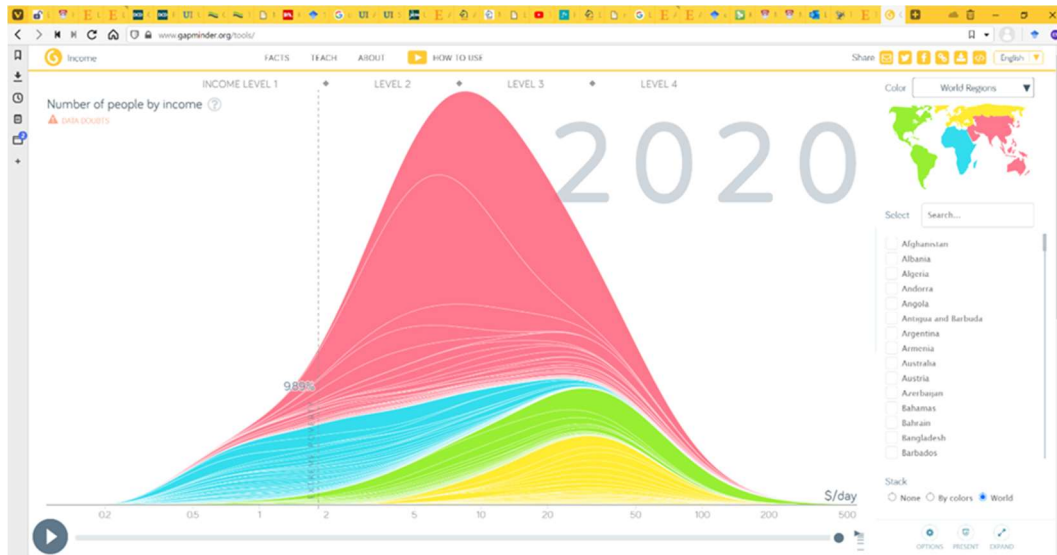


Figure 38. More than 800 million people live under the poverty line (USD 1.90/day). As problematic as that is, this means nine out of ten live above the line. Graph made using Gapminder ([www.gapminder.org](http://www.gapminder.org)).

## 9.1 Testing the analytical framework as a checklist for identifying potential data centre sites

As shown in chapter 5, establishing a data centre industry is not a trivial task. It is also a matter which almost certainly will require governmental involvement, due to intricate dependencies with and between technology- and power-related infrastructure. Chapter 7 introduced three studied cases: large-scale data centres for commodity dehydration in Cyberjaya (Malaysia), small-scale data centres for coffee bean drying in Costa Rica and large-scale data centres for district heating and commodity dehydration in (northern) Sweden. In this section, the analytical framework (Figure 11) is applied to the three cases (note that the headings in this section follow the naming of the analytical framework).

### 9.1.1 Energy quality and use

The source of the energy coming to the data centre varies greatly between the three cases. Malaysia is mainly powered with fossil fuel. Renewable energy initiatives exist,

but are, for now, limited. In Sweden, in stark contrast, data centres are ideally powered fully with renewable energy, and even when they are not, they are not running on fossil fuel (but a mix of hydropower, wind power and nuclear power). Seen from this perspective only, Swedish data centres have the upper hand against Malaysian, as Table 9 illustrates (see section 9.2). The Costa Rican case differs from the others: here, what energy is used by the data centre is not at all discussed. Instead, this topic involves how to replace electricity or fossil fuel used to dry coffee beans, that is, make use of the next step of the energy chain. For all three cases, the question remains whether the energy could have been used better.

### **9.1.2 Preferable data centre configuration**

A large data centre in northern Sweden can to a certain degree make use of free cooling from outside air. In Malaysia, this is definitely not the case, due to the warm ambient temperatures. Hence, data centres may be configured differently between the two countries. Be that as it may, for the time being, it seems likely that the large data centres in these regions will be traditional, air-cooled systems.

For the Costa Rican case container-sized data centres were suggested (see section 7.4). Ideally, because of their small scale and remote locations, they should be as self-contained as possible – not that different from the ocean-based Naticks. This means that they are good candidates for liquid immersion cooling. Immersed in oil, servers are not exposed to residues from plants and insects, which they risked being with decaying air filters. A disadvantage of liquid immersion cooling is the hassle of dealing with oil when replacing, reconfiguring or repairing servers. However, with a self-contained facility, this would not present a problem. Before ultimately requiring replacement, the Costa Rican servers could therefore allow for heavy compute, because



of the elevated temperatures possible with immersed servers. Interestingly, this means that these rurally placed data centres could offset some computation in San José, the capital of Costa Rica.

### **9.1.3 Geographies**

But for the Cameron Highlands, peninsular Malaysia features a quite hot climate most of the year, and humidity is high (Lockard et al. 2024). In other words, purely from a geographical viewpoint, this is not the ideal setting for data centres. Though Costa Rican coffee plantations are found in different climates (Ovalle-Rivera et al. 2020), those are cooler than Kuala Lumpur and its surroundings – and in that sense a better choice for data centre placement from a geographical perspective. The threat to infrastructure from natural disasters is comparatively low in all three settings (Aleksandrova et al. 2021), though volcanoes and earthquakes pose natural hazards to Costa Rica (Elbow et al. 2024).

In terms of infrastructure access, northern Sweden excels (Velkova 2016; GSM Association 2020a). Clean electrical power is abundant and reliable. The places where data centres are positioned have short distances to the Internet backbone, here mainly represented by the Swedish TeliaCarrier network. This resource-rich region, home to iron ore and deep forests, is also known for its lakes, rivers and considerable hydropower resources. In other words, there is an abundance of clean energy and fresh water. The road and railroad infrastructures are excellent, making transportation of goods and people to data centres unproblematic. Something that speaks against Sweden (and its neighbours Norway and Finland) is the distance to end-users. In total, the population of three countries is approximately 20 million, or just a fraction of Europe's total population of about 600 million (Russia excluded).

A bit surprisingly, Cyberjaya shares many geography-related traits with Sweden, despite its vastly different climate. Infrastructure access is first-rate, and so is access to the Internet backbone. Also electricity access is a non-issue. However, the energy mix is much more of an GHG emitter than northern Sweden's and energy reclamation is more difficult due to the high temperatures. Thus, there are important factors that speak against data centre establishment in Cyberjaya. On the other hand, Cyberjaya is close to a large part of the world's population and likely a better site than neighbouring Singapore (a major data centre hub today) due to an overall higher capacity, more space, more available fresh water – in fact, Malaysia exports fresh water to Singapore today (ASEAN Coordinating Centre for Humanitarian Assistance on disaster management and Japan International Cooperation Agency 2015) – and more domestic commodities to dehydrate.

In the Costa Rican case, calculations were made on the container-sized data centre with the highest power handling commercially available at the time (Terenius, Garraghan, and Harper 2021). Through liquid immersion cooling, a power draw of 500 kW would not be impossible. Therewith, there are two options for Costa Rica: small data centres meant to serve the local population with data, and data centres that feature a more active role in Costa Rica's ICT landscape. The former would be less expensive and in less need of power; the latter would – where the electricity grid allows for it – drastically increase compute capability and allow for much greater waste heat use operations.

For the Costa Rican case, a few questions regarding infrastructure access remain. What is the electricity grid capacity in rural areas? And how far away is the closest major Internet exchange point? Here, it should be noted that the very reason why the

Costa Rican project is proposed, is to strengthen data access in rural areas. Thus, it is to be expected that additional network infrastructure is needed for this idea to bear fruition. That being said, there are plenty of fibre optic cables crossing Costa Rica (see Figure 31). Therefore, it should be possible to find a location where coffee growers are in the relative vicinity of both fibre optics and a reliable electricity grid. Or, more succinctly, to find the sweet spot in the crosshair of data needs, dehydration opportunities and infrastructure access.

#### **9.1.4 Deployment, operation**

For Malaysia and Sweden, building new data centres and configuring transmission networks accordingly is not a major obstacle. In Sweden, Big Tech are commonly building and overseeing new data centre establishments. Sweden also has a strong energy sector (ABB and Vattenfall) and one of the world's most experienced ICT network constructors (Ericsson). During the time in Malaysia, it was noted that in that part of the world, the Western enterprises are not as prominent. Indeed, whereas Europeans turn to the USA for collaboration, technology assistance and business or research opportunities, Malaysia cooperates with Japan. In fact, many of the Malaysian researchers visited knew Japanese, and at the UTM, there were numerous strategic partnerships in place with Japan. The Japanese also helped to build the elaborate highway system in Kuala Lumpur and other country infrastructure (depicted in Figure 39).




Figure 39. Left: Thesis writing on the train from Ipoh to Kuala Lumpur. Right: a few of the five or so levels of road and railroad infrastructure in Kuala Lumpur.

It was concluded that in a rural, mountainous area, deployment – installation especially – and operation would pose greater challenges. This is why preconfigured, self-containing containers using liquid immersion cooling are proposed here. To deploy, all that are needed are a) a concrete slab (or solid earth or rock), b) connecting to the existing electricity grid and fibre optic cable and c) if applicable, filling up the immersion tanks. No maintenance should be needed, unless connectivity to the electricity grid or the Internet fails. The setup may seem optimistic, but through the Naticks, Microsoft has already proven that the concept of self-sustaining small-scale data centres is working quite well (see section 5.4.2).

### 9.1.5 Heat use capabilities

Chapter 6 discussed the possibility of waste heat use, and extends far beyond heating buildings, which is the strategy mainly employed today. Waste heat use is quite location-specific. In Sweden, heating buildings, greenhouses, swimming pools and fish tanks is viable. So is dehydration of pellets and possibly timber. Something to keep in mind is that Sweden's large power-hungry industries offer hotter waste heat than what data centres provide, limiting the relative value of data centre waste heat. One must therefore match waste heat uses to the specific area. Ideally, and as observed in the

German Energy Efficiency Law, this should be done when the data centre is planned rather than in hindsight.

For Malaysia, the main challenge for data centre waste heat use is the low difference between waste heat temperature and ambient temperature. However, data centre waste heat is filtered, thus cleaner than outside air, and the temperature and humidity more stable than what sun-drying allows for. Some suitable uses for the waste heat are fruits, cardboards and possibly seaweed. In addition, as shown in section 6.4, data centre waste heat has the potential to increase the work efficiency of OTEC plants.

The Costa Rican case is about coffee bean drying. This is a suitable commodity in that it requires low temperatures for dehydration and that it is harvested in rural areas where data access may be limited. A question is how one can use the waste heat off-season. Hence, further research is needed to identify locally produced commodities in need of drying the rest of the year, to maximise the use for the waste heat. Therewith, finding the “sweet spot” as mentioned above will be a more complex undertaking – which is why the framework was created in the first place.

A matter of concern for waste heat use is the pursuit of a suitable metric. Last chapter showed that today’s endeavours are not enough: they are not precise enough, but more importantly, they do not measure the right thing. Indeed, as evident in section 9.2, a holistic metric is the only appropriate way to deal with this issue. It is also the only reasonable way a politician can evaluate the sustainability of planned data centres that compete for the same resources, as well as move beyond arbitrarily set PUE value limitations and instead establish reality-based sustainability limits for coming data centres.

### 9.1.6 Society profile

The three cases were chosen to relate to three quite different society profiles. As it turns out, differences between the Malaysian and the Swedish society structures are not that big, in part due to the British earlier influence over Malaysian legislation. In fact, there are more similarities than differences in terms of policy, finances, legislation and level of infrastructure. As a nation, Costa Rica is – on paper – rather alike the two other nations. However, the case for Costa Rica excludes large-scale establishments in San José, meaning that statistics on the national level would not say that much.

A trademark of the Nordic culture, and a success factor, is collaboration between local authorities and industry. In contrast, much of large Malaysian infrastructural developments are results of, and decided by, the country leaders. How decision-making is carried out in practice in Costa Rica has not been investigated here.

Of the three nations, ICRI (see section 7.2.2) is lowest for Costa Rica. Again, the proposition is for the rural parts of Costa Rica, where ICRI would be especially low. From a business perspective, this is good news as ICRI identifies potential market spaces. More importantly, the low ICRI indicates that data centre establishment in rural parts has the potential to strengthen local societies and accelerate Costa Rica's transformation to a knowledge-based nation.

### 9.1.7 Commentary

Having applied the analytical framework to the three cases, some patterns occur. For example, it is not strange Asta Vonderau called Luleå the “*perfect place for locating cloud infrastructure*” (Vonderau 2019) – Sweden excels in just about every aspect of the framework. Viewing all three cases, the framework does seem to have some merit, as it has helped identify strengths, weaknesses, opportunities and threats to data centre

establishment despite three very different settings. Doing so, it also proved to be quite flexible: It is applicable to both large-scale and small-scale facilities, and agnostic to technologies used inside the data centre. However, even if the analytical framework seems to function rather successfully, it is for now limited to the view of this thesis. Therefore, it would benefit from further development, based on input from the societal actors as well as from data centre engineers.

A final reflection on this matter: Something discovered during a visit to a Malaysian warehouse, was how much hotter the outside air got, since the warehouse acted as a light-trap. Thus, where greenhouses or warehouses are used, the value of data centre waste heat in Cyberjaya may be less than anticipated from just watching the thermometer. However, this makes for a compelling case for liquid immersion cooling: regardless of suitable uses for the waste heat, the elevated temperature made possible with liquid immersion cooling may be the only way to take full advantage of data centre waste heat in the tropics. In the best of worlds, the tropics may therefore provide the injection the liquid immersion cooling industry needs to grow – and as shown in next section, DESS may assist in this endeavour.

## **9.2 PUE and DESS**

The previous chapter introduced DESS as a complement to existing data centre sustainability metrics, and as a replacement as a decision-making tool for external stakeholders such as politicians. Table 9 and Figure 40 show sixteen scenarios, spanning a variety of climates, economies and usage of renewable energy.

Table 9. The annual PUE and DESS cases mentioned in the thesis, and a few comparison cases. To facilitate comparisons between PUE and DESS, and to better spot the differences in

information conveyed, the scenarios' PUE and DESS rankings are shown in italics, and top ranked scenarios are shaded.

| Scenario                                     |            | EnergyIn          |                       |                           | Compute                   |                  | EnergyReuse |                         |                    |                      | Metrics                  |                          |
|--|------------|-------------------|-----------------------|---------------------------|---------------------------|------------------|-------------|-------------------------|--------------------|----------------------|--------------------------|--------------------------|
| Reuse scenario                               | Location   | Total energy, MWh | Sustainability rating | Sustainable energy, ratio | Used by IT equipment, MWh | Overhead, factor | Reuse, %    | Degree of usefulness, % | Reused energy, MWh | Reused energy, ratio | PUE                      | DESS                     |
| Current energy mix, no reuse                 | Cyberjaya  | 1000              | 2                     | 0.2                       | 625                       | 0.60             | 0%          | 0%                      | 0                  | 0                    | <b>1.60</b><br><i>13</i> | <b>1.0</b><br><i>14</i>  |
| Current energy mix, commodity dehydration    | Cyberjaya  | 1000              | 2                     | 0.2                       | 625                       | 0.60             | 80%         | 70%                     | 560                | 0.56                 | <b>1.60</b><br><i>13</i> | <b>2.12</b><br><i>12</i> |
| 100% renewable energy, commodity dehydration | Cyberjaya  | 1000              | 9                     | 0.9                       | 625                       | 0.60             | 80%         | 70%                     | 560                | 0.56                 | <b>1.60</b><br><i>13</i> | <b>4.22</b><br><i>6</i>  |
| Coffee (and leather etc) dehydration         | Costa Rica | 1000              | 9                     | 0.9                       | 625                       | 0.60             | 65%         | 80%                     | 520                | 0.52                 | <b>1.60</b><br><i>13</i> | <b>4.14</b><br><i>7</i>  |
| District heating                             | Nordics    | 1000              | 9                     | 0.9                       | 870                       | 0.15             | 50%         | 90%                     | 450                | 0.45                 | <b>1.15</b><br><i>6</i>  | <b>4.45</b><br><i>5</i>  |
| Pellet dehydration and district heating      | Nordics    | 1000              | 9                     | 0.9                       | 870                       | 0.15             | 90%         | 80%                     | 720                | 0.72                 | <b>1.15</b><br><i>6</i>  | <b>4.99</b><br><i>2</i>  |
| Californian data centre, no reuse            | California | 1000              | 4                     | 0.4                       | 800                       | 0.25             | 0%          | 0%                      | 0                  | 0                    | <b>1.25</b><br><i>8</i>  | <b>1.95</b><br><i>13</i> |
| Google (fleetwide avg), no reuse             | World      | 1000              | 9                     | 0.9                       | 890                       | 0.12             | 0%          | 0%                      | 0                  | 0                    | <b>1.12</b><br><i>4</i>  | <b>3.58</b><br><i>9</i>  |
| Google (fleetwide avg), reuse                | World      | 1000              | 9                     | 0.9                       | 890                       | 0.12             | 75%         | 80%                     | 600                | 0.6                  | <b>1.12</b><br><i>4</i>  | <b>4.78</b><br><i>3</i>  |
| Blackpool, no reuse                          | UK         | 1000              | 9                     | 0.9                       | 800                       | 0.25             | 0%          | 0%                      | 0                  | 0                    | <b>1.25</b><br><i>8</i>  | <b>3.45</b><br><i>10</i> |
| Blackpool, pool heating                      | UK         | 1000              | 9                     | 0.9                       | 800                       | 0.25             | 95%         | 90%                     | 855                | 0.855                | <b>1.25</b><br><i>8</i>  | <b>5.16</b><br><i>1</i>  |
| Natick, no reuse                             | Ocean      | 1000              | 9                     | 0.9                       | 935                       | 0.07             | 0%          | 0%                      | 0                  | 0                    | <b>1.07</b><br><i>2</i>  | <b>3.63</b><br><i>8</i>  |
| Natick, pond heating                         | Ocean      | 1000              | 9                     | 0.9                       | 935                       | 0.07             | 100%        | 50%                     | 500                | 0.5                  | <b>1.07</b><br><i>2</i>  | <b>4.63</b><br><i>4</i>  |
| Coal-powered, no reuse                       | World      | 1000              | 0                     | 0                         | 800                       | 0.25             | 0%          | 0%                      | 0                  | 0                    | <b>1.25</b><br><i>8</i>  | <b>0.75</b><br><i>16</i> |
| (Coal-powered, full reuse)                   | World      | 1000              | 0                     | 0                         | 800                       | 0.25             | 100%        | 100%                    | 1000               | 1                    | <b>1.25</b><br><i>8</i>  | <b>2.75</b><br><i>11</i> |
| (Coal-powered, PUE=1, no reuse)              | World      | 1000              | 0                     | 0                         | 1000                      | 0                | 0%          | 0%                      | 0                  | 0                    | <b>1.00</b><br><i>1</i>  | <b>1.00</b><br><i>14</i> |



To facilitate comparisons between the data, all share TotalEnergy and most share sustainability coefficients (0.9 is used rather than 1, even in cases where 100% renewable energy is utilised, to compensate for the fact that energy could have served other purposes, and that most energy sources are limited).

To represent future typical installations, PUE values are set to current leading endeavours:

- 1.60 (which happens to be about the world average PUE among monitored data centres (Davis et al. 2022)) is the reported value of NTT's flagship Cyberjaya Data Center 5 (NTT n.d.).

- 1.60 is also used for the container-sized data centres of Costa Rica. This value would differ based on data centre configuration and computational load. The climate is colder than in Cyberjaya, indicating the possibility of lower PUEs, but it is more difficult to attain low PUEs in small data centres, pushing PUE up again.

- 1.25 is used for data centres in the UK and North America. To facilitate comparison, so is the coal-powered data centre.

- The reported PUE for EcoDataCenter, 1.15, is used for the Nordic examples.

- The reported average PUE for Google's data centre fleet is an impressive 1.12 (Google 2020), and for Microsoft's submerged Natick containers it is 1.07.

EnergyReuse estimates are based on the many factors of the analytical framework, the literature study and the research visits carried out during the PhD work. They should be read as follows:

- In the Cyberjayan example, 80% of the waste heat is used, and the usefulness of that heat is 70%.

- In the Costa Rica example, 65% is reused to dry coffee, but also leather, fruit etc. The usefulness of the heat is higher than in Cyberjaya's case, since Costa Rican coffee is grown on high altitudes, featuring less sunshine, a colder climate and a weaker transportation infrastructure.

- In one of the examples from the Nordics, district heating is only used in winter (no district cooling is accounted for), thus 50%. The heat is seen as quite valuable.

- In the other example, since the waste heat is utilised all year around, its energy reuse is set to 90%. However, drying pellets is (here) less efficient than district heating, thereby lowering overall yearly usefulness.

- The Californian data centre lacks waste heat use, and is powered through an energy mix, consisting of the statewide average 40% fossil fuel dependency (Nyberg 2022).

- Google's data centres are sustainably powered (Google 2020). Since Google's data centres are found in many parts of the world, reuse options are here unrelated to specific uses and geographical and societal contexts.

- In 2022, the Fibre Blackpool Cooperative Alliance announced a sustainable data centre nexus on England's western coast. The data centres will be powered by an ocean-based wind farm. Since Blackpool is a beach resort that hosts a number of swimming pools, data centre waste heat can potentially be used as-is, heating these pools. Heat from electricity or fossil fuels would be substituted in the process, which is why both energy reuse and usefulness values are high.

- The Natick tests have not included waste heat use. However, should the outgoing heat be connected to an enclosed environment, such as a pond, it is to be expected that submerged data centres can enhance fish or shellfish production in some locations.

- The coal data centre in this example fares well in terms of PUE but not in terms of DESS.

- To clearly illuminate the differences between PUE and DESS, two unrealistic coal-powered use cases are added, one with 100% reuse and one with no reuse but PUE=1.

Figure 40 highlights the discrepancies even more clearly than Table 9 does, showing PUE and DESS side by side. PUE and DESS are not truly compatible, as PUE has no upper limit. In this figure, however, PUE=2 is set as the maximum tolerated number; anything higher would be unacceptable in a modern data centre.

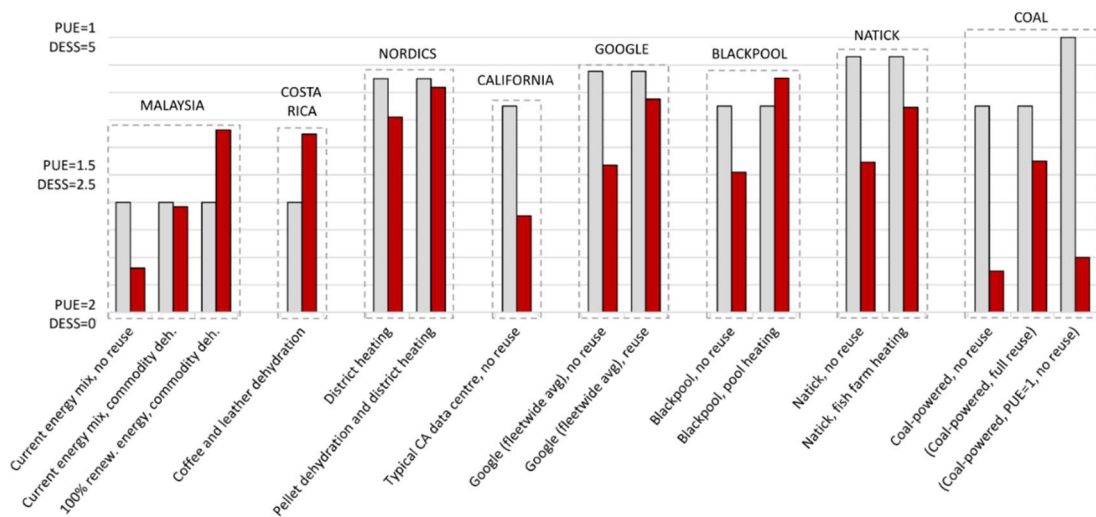


Figure 40. A number of comparison cases for PUE (labelled grey) and DESS (labelled red). To facilitate comparison, PUE is shown from high to low on the y axis. This is because minimum PUE (1) roughly corresponds to maximum DESS (5).

The differences between the PUE and DESS values, and their implications, are discussed in section 9.3.2.

## **9.3 Reflection on the two analysis methods**

In this chapter, the three cases of the thesis have been analysed in two ways: through the analytical framework and through applying PUE and DESS in a metrics comparison.

This section accounts for a reflection on the results.

### **9.3.1 The analytical framework**

There is reason to believe the factors presented in the analytical framework have merit when deciding upon data centre establishment. As Myovella et al. suggest, rather than drawing up template data centres and imposing them on a society, careful consideration of both technical and societal factors may greatly help the society (Myovella, Karacuka, and Hauca 2020). As a consequence, through citizen approval of the data centre, the implementation process would become easier.

It is clear the data centre industry must be able to acknowledge more factors that pertain to data centre establishment, work smarter to better mitigate negative environmental and societal impacts from building new data centres, and better communicate the positive aspects of it. In 2022, Meta's projected 400 MW Dutch data centre was finally rejected, after months of protesting (Moss 2022). With proper recognition of the factors in the analytical framework, the outcome may have been another. For example, since data centres can heat greenhouses, it is possible the giant facility could have prolonged the growing season of Dutch agricultural commodities (a similar strategy can be employed by the UK (Judge 2023)). Hopefully, the analytical framework can spur ideas that truthfully convey also the positive sides of data centre establishment within a society.

### 9.3.2 PUE and DESS

The differences between the resulting PUE and DESS values of Figure 40 are striking. Unsurprisingly, accounting for renewable energy and reused heat energy makes – correctly – a big difference for DESS. For example, the coal-powered data centres perform very well as far as PUE is concerned, but terribly in terms of DESS. Table 9 and Figure 40 also provide several reasons why DESS should interest the industry, clients and legislators. First, the advantage PUE erroneously gives Google over the “Nordics” and “Blackpool” data centres (which use their heat) is corrected with DESS. In other words, the new metric makes it easier for EcoDataCenter and similar enterprises to compete with sustainability arguments and, potentially, new sustainability certificates. Second, DESS provides governments on national, regional and local levels with a more valid key metric for their regulations, reflecting the intentions of the German Energy Efficiency Law and the EU EED. Set correctly, DESS-based regulations should result in newbuilds with higher sustainability foci. Third, with a more accurate metric in use among decision-makers, data centre owners get an inducement to invest more wisely: DESS is designed to encourage savings and even earnings from the three parts of the energy chain, rather than solely point to advanced and expensive cooling systems as the main strategy to improve sustainability ratings. On that note, it should be stressed that while the industry still waits for SERT or a similar metric to replace PUE for data centre internal efficiency, what is inside the data centre is based on PUE also in DESS. Consequently, no additional work needs to be carried out for the time being for data centre managers.

## 9.4 Chapter conclusion

This chapter concludes the six-chapter-long hunt for optimal cases for data centre waste heat in different locations and societal contexts. The point of this chapter has *not* been

to find an ultimate location for data centre establishment. But why not? The obvious “winner” of the three cases would be Sweden; as Vonderau points out, it is “*the perfect place*” (Vonderau 2019). However, this is not where many of the future data centres need to be built, should the desire for low latencies persist. Malaysia is much better suited in this regard, as it has already started to build its “Multimedia Super Corridor” (see section 7.3). In other words, Cyberjaya may eventually become an important ICT nexus in Southeast Asia, potentially serving hundreds of millions of users with ICT access within the decade. However, the high heat presents an obvious challenge for a sustainable data centre sector. In other words, what is “best” is a matter of perspective, of desires of the industry and of societal needs. And in any case, as DESS shows, heating swimming pools in Blackpool may trump district heating and pellet dehydration in the Nordics. So maybe who the “winner” is, is not so clear-cut after all.

Similarly, rural Costa Rica would likely not be the first choice for the data centre industry. However, with the prospects of waste heat use, stable GDP growth and a low ICRI, this market may seem more attractive. DESS also indicates that with a well-designed solution based on a systems perspective on societal opportunities and needs, the idea may be feasible and even financially rewarding. In addition, the environmental benefit would be noticeable, if not significant. Hopefully, if well-implemented, this idea means that societies can prosper as well. If they can, much has been gained, for Costa Ricans, for migrant workers coming from Nicaragua and Honduras, and for many more of the 800 million people worldwide now living under the poverty line.

# 10 Discussion

*“Until recently, the planet was a large world in which human activities and their effects were neatly compartmentalized within nations, within sectors (energy, agriculture, trade), and within broad areas of concern (environment, economics, social). These compartments have begun to dissolve. This applies in particular to the various global ‘crises’ that have seized public concern, particularly over the past decade. These are not separate crises: an environmental crisis, a development crisis, an energy crisis. They are all one.” (Brundtland et al. 1987)*

These words, so appropriate in today’s world, are over three decades old, and taken from the United Nations’ aforementioned *Our Common Future*. The words exhibit many of the problems mentioned in this thesis, and also ties to a systems science-based worldview. The report’s environment–economy–society troika, commonly referred to as the “triple bottom line”, steered sustainability endeavours in everything from governmental and intergovernmental institutions to the education and industry sectors until the advent of the UN SDGs in 2015. In both models, the world and its constituents can indeed be seen as constantly interacting and changing *systems*. Just about every thinkable real-world system has one or more inputs, internal workings, and

one or more outputs. For a data centre, this means data to process, data computation and storage of data, and finally processed or requested data. At the same time, a data centre is an industrial facility dealing with incoming power, computing and cooling, and finally waste heat. Thus, three steps constitute what is of main concern for the data centre industry and for data centre-related research, as shown in Figure 41.

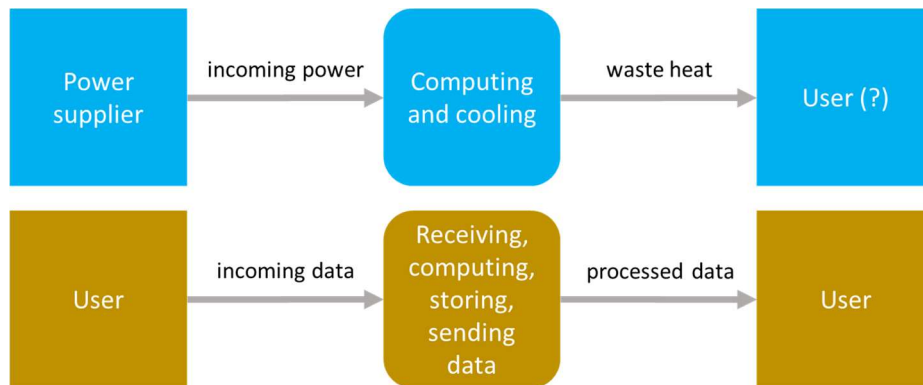


Figure 41. Components of a data centre, focusing on energy (top) and on data (bottom).

Viewed as an industrial facility, a data centre differs from a traditional factory, in that the main output is immaterial rather than physical. In turn, this immaterial output can be delivered worldwide, literally at the speed of light (lag caused by routers excluded). That being the case, the data centre also has a local presence; it has a place in the society in which it is erected.

Being embedded in a local context presents its own opportunities and challenges. In northern Sweden, using waste heat is easy because of many contributing factors. In Malaysia, the very idea of reusing data centre waste heat has never been suggested. Thus, chapter 6 discussed existing applications for waste heat, and proposed several new applications, since much of what works in the Nordics is not feasible in the tropics. Moreover, based on the relationship between available or proposed applications for data centre waste heat and the local societal context, chapter 7 explored the three main geographically bound cases of the thesis: Malaysia, Sweden and Costa Rica.



As can be seen, there is a relationship between the figure above and Figure 19, which showed a data centre as an enabler of sustainable societies, through providing energy and ICT access. The key point here, and something that is mostly omitted (e.g. (Lannelongue, Grealey, and Inouye 2021; Tozer et al. 2018)), is waste heat use. As a concept, data centre waste heat use is nothing new. In fact, the Energy Reuse Effectiveness (ERE) and Energy Reuse Factor (ERF) metrics were formed for this very reason already in 2010 (The Green Grid 2010b). However, ERE failed to gain traction. One possible reason for this may be its lack of perceived usefulness, another the principal aim of the data centre operator: high uptime (see section 5.3.2).

The de facto metric for energy use within data centres, Power Usage Effectiveness (PUE), correctly informs the owner of electricity overhead, that is, electricity that does not reach the servers of a data centre but instead is used for cooling or other things. However – and vital to the argumentation of this thesis – one should keep in mind that also at an (unobtainable) PUE of 1, a data centre still has a substantial GHG footprint. And after several years of making data centre internals more energy-efficient, PUE gains are slowing in. In contrast to what Masenet et al. claim (2020), to *“bring about a near-term plateau in energy use, which provides critical time to prepare for the possibility of future energy demand growth”*, is not enough when the effects of climate change are upon us. As a consequence, other sustainability gains must be carried out, and since we cannot do enough inside the data centre walls, we must move beyond them. But for these sustainability endeavours to be optimally implemented, one needs to be able to measure them. Consequently, chapter 8 dealt with the very particulars of data centre energy efficiency: the metrics that underpin data centre managers’ energy efficiency work.

Luckily, today, environmentally conscious corporations increase their climate change mitigation efforts through reforestation projects and similar campaigns. The leading data centre operators also push for very high energy efficiency. However, though benevolent, these actions do not mean electricity has not been used in the first place. Even in cases where electricity may be hundred percent renewable, there is still an environmental cost in, for example, producing and installing solar panels, wind turbines or similar. Besides, other processes, such as manufacturing or transportation, could also have made use of this energy. Thus, whilst continuous improvement is beneficial, this is not where the main advances can be made today or tomorrow, and the industry now faces a situation where galloping energy use can come to threaten worldwide energy security. So, after a decade of mitigating overhead electricity use considerably, it is appropriate to again observe what energy really *is*. As explained in chapter 4, energy is a promise: a promise of work to sometime be carried out. This promise means electricity arriving to a data centre is not lost energy after compute. On the contrary, the promise should make societal actors focus their attention on energy matters outside of the data centre – the cleanliness of electricity used, and waste heat use. In truth, it is time we see data centres for what they are: not abstract and intangible clouds but pieces of the local and global societal puzzle.

## **10.1 This thesis in retrospect**

This thesis was written based on the assumption that data centre waste heat can be used, and can have value, in many geographical and political contexts. The work commenced by interrogating the use of energy in a data centre. It was concluded that much of the literature concerns energy use within the data centre, not so much data centre energetics in societal and environmental contexts. The intersection of data centre energetics and societal energetics therefore became the domain.

This hitherto unpopulated domain needed a *worldview*. For reasons stated in the beginning of chapter 3, systems science and its derivate systems engineering suited this worldview well. The domain further needed a framework for analysis, based on the intersection of the data centre and the society. However, fields in social sciences (such as science and technology studies, STS) are mainly exploring their topics with methods and worldviews of the social, and engineering fields may not observe the social and societal aspects fully. Hence, recognising the merits of both engineering and social sciences, the thesis belongs to material social studies, where the material and the social/societal are given equal weight.

In this investigation, visual observation has at times been employed. This a hallmark of the “anthropologists” methodology, again and again used by the authors quoted in chapter 2. However, whereas Ortar et al. use such approaches to “*disentangle the relations between materialities, flows, and data and energy capitalism*” (Ortar et al. 2023), visual observation is here used inversely. As seen in Figure 37, the entanglement that a data centre contributes to is, here, welcomed: it is in the combination of heat and processed data that one, with a systems-based worldview, can see opportunities. In short, Dr Azzman’s idea of using data centres to push societies around the Globe into knowledge-based and prosperous systems seems, with this view, reasonable.

With the stage set, exploration began, of applications, societies and metrics. Applications, as data centre waste heat use in tropical and temperate regions had not previously been considered; societies, as a precondition for waste heat use is the integration with a specific society (many of today’s large data centres may be templated, but societies are unique); metrics, since to measure the validity of strategies proposed, adequate energy-related metrics were required. It soon became evident that the data

centre industry struggles with such metrics: PUE only tells a fraction of the energy story, but is nonetheless used for policy decisions on both local and national levels (see section 8.1.2). With the analytical framework and new metrics in place, it was possible to evaluate the feasibility of the ideas proposed.

Investigating the problematics of the data centre industry through a material social lens has required a multi-disciplinary investigation. Systems engineering's strong focus on problem-solving and of making connections between entities (Watson 2019) has facilitated the investigation. To embrace cross-disciplinarity and give equal weights to engineering and societal concerns, the mixed methods paradigm was helpful, since it stresses the use of both qualitative and quantitative approaches and is suiting when addressing "*complex, interdisciplinary research problems*" (Creswell and Garrett 2008).

To conclude: this thesis dives deep into social and material aspects of data centre and societal energetics. The duality of this perspective is complemented – and made more complex – by the juxtaposition of materiality and immateriality of the data centre itself. Hu notes that

*"inside a data center, data are connected only paratactically to each other, porn next to military documents next to banking records next to your e-mail. In that model of physical proximity, we might see an alternate map of community."* (Hu 2015)

What Hu describes is the "*prehistory*" – or perhaps more accurately, the genesis – of the cloud. The term "cloud" was earlier criticised for being suggestive (see section 2.3). However, in Hu's perspective, a "cloud" is a suitable metaphor of duality: a real cloud is neither a tangible nor an intangible object; it is something in between, indeed, something fluffy. As opposed to a cloud, however, a large-scale data centre is *both*, not

*neither* – through its grandiosity (see Figure 42) and its invisible and silent transportation of goods, it is both (very) physical and intangible, both material and social.



Figure 42. A “home for the cloud”. Perhaps, but also the enormous Interxion data centre buildings in the outskirts of Zürich. The tram station on the right looks minuscule in this photo.

A reason for setting this thesis in material social studies is that the highly explorative research undertaken paves new ground. In this case, the research problem has been approached through an extensive and cross-disciplinary literature review, through visual observation, through experimentation, through participation in data centre industry hearings and through other activities. Taken together, the approaches pointed towards the same conclusion: data centre waste heat use is possible in both cold and warm climates and regardless of social structures and local financial situation. Thus, the approaches seem to have converged, pointing to “*the same set of events, facts, or interpretations*” (Yin 2012) that are a hallmark of triangulation.

## 10.2 Reflections on the research questions and purpose

The PhD work has been guided by a set of research questions. This section revisits those questions and comments on them.

### **10.2.1 RQ 1: How can data centre waste heat be used in different geographical and societal contexts?**

There is no short answer to this question. This thesis shows that of the possible applications for data centre waste heat use discussed in chapter 6, only a few are employed today. Further, chapters 7 and 9 show that applications must be mapped to geography and societal context, as there is great variance between suitable uses for the heat.

The main point is that using data centre waste heat is possible, profitable and the environmentally responsible thing to do, for today's data owners and for politicians of all levels. As the question implies: it's not anymore about "if", but about "how".

### **10.2.2 RQ 2: What are the societal implications and environmental benefits of implementing the ideas proposed?**

A key point of the reasoning in this thesis is that the ideas investigated should carry potential *value*. They should either be of value to society or to nature – or preferably to both. The coffee bean drying case (see section 7.4) is an example of this dual value. OTEC integration (see section 6.4) is another.

In addition, the ideas proposed help mitigate climate change. First, using waste heat offsets heat retrieved from other energy sources. With a yearly consumption of 1% of world electricity, this offset is bound to have significant value. Second, good DESS values (see section 8.2) can only be obtained if the energy source is a low-GHG emitter.

### **10.2.3 RQ 3: Which is the optimal design for a sustainable data centre?**

Also this question's answer is dependent on other factors. One of these is the desired temperature of the waste heat. For high temperatures, liquid immersion cooling is the only available option today.

On a more general level, the answer is that the data centre should be designed to optimally fit into the society in which it is placed: a data centre in the Nordics should certainly be designed or retrofitted to take advantage of waste heat. A data centre in rural parts of Costa Rica should be small, to reflect the local needs and the dimensioning of the electricity grid. Rather than discarding of functional infrastructure, an American data centre owner having a reliable power supply should think twice before replacing gasoline-driven backup generators with huge lithium battery packs (see section 5.2.1).

#### **10.2.4 RQ 4: Can there actually be such a thing as a one key metric for data centre energy efficiency, and if so, how should it be designed?**

This question has been returned to over the entire scope of the PhD work. A first metric developed did indeed improve on PUE and ERE/ERF. However, during later refinements, it was concluded that the complexity that placement of sustainable data centres is subjected to is simply too large to be captured by a singular, detail-specific metric. Hence, a higher-level metric, DESS, was construed. Perhaps it was in light of the EUs coming Energy Efficiency Directive (see section 8.4) and the German Energy Efficiency Law that DESS received attention by the data centre industry (see Figure 43) immediately after its journal publication (Terenius, Garraghan, and Harper 2023).



Figure 43. Presenting DESS at the OCP European Summit, Prague 2023. Photo by the OCP Foundation, used with permission.

DESS actually has the potential of becoming a “one key metric”. One reason why it works is because of its dependence on measurement on the lower level (namely: a) origin of and possible other demand for the incoming energy, b) the data centre’s inner workings and c) the way the energy is reused). After almost fifteen years, we should stop the search for a detail-specific metric; we need to accept complexity for what it is, and turn to holistic metrics in order to achieve meaningful goals. The alternative, an ever-increased detailed specification not only causes headaches for the industry, but also gives the opportunity to work around sustainability endeavours. For instance, a data centre owner in Germany could potentially dig a pit outside the data centre, call it a swimming pool and enjoy a maximum ERF value, and at the same time follow the German Energy Efficiency Law.

Accreditation bodies will be needed to evaluate the three constituents of DESS (see section 8.2). Since there are already accreditation bodies such as Uptime Institute for data centre reliability certification, expanding to sustainability concerns is not a major step.



### **10.2.5 RQ 5: To what extent is material social studies suitable for investigating cross-disciplinary problems?**

Earlier attempts to discuss data centre sustainability and data centre energetics have had limited successes. This may partially be a consequence of the scopes of the cultures of the different scientific fields:

- Whilst good at addressing compartmentalised problems, mechanical, electrical and computer engineers have so far been unsuccessful in fully solving the matters associated with data centre energetics.
- Watchdogs more than anything else, anthropologists are excellent in identifying problems, but do not as often seem to attempt to solve them.
- Sustainability science researchers do indeed aim to both identify and solve problems, but data centre energetics is too specific a field – and maybe too far away from their own – to be good a candidate for enquiry.

Instead, a more holistic, explorative and solution-oriented paradigm for investigation has been required to solve this complex and multi-disciplinary enquiry. Systems engineering excels in this regard, founded as it is in systems science principles such as set theory. Therefore, *systems engineering* has been instrumental to the completion of this thesis. However, methods and worldviews of the engineering disciplines have not been enough to explore the topic of the thesis. Hence, also the views and methods of the *social sciences* have been used, and insights from *sustainability science* have been needed. From this follows that the thesis's wide, multi-disciplinary net of themes would have been impossible to throw, had it not been for the use of *mixed methods*: after all, not every methodology is well-suited to all disciplines. The mix of engineering, social sciences, sustainability science and mixed methods pointed to

*material social studies* as the disciplinary home of the thesis. The commentary to research questions 1-4 indicate that material social studies has been able to shed some light on the topics under investigation.

### **10.3 Further research**

Traditionally, academic works end with a call for more research on the subject. In a multi-disciplinary study, this is particularly suitable due to its many connecting parts. Pertaining to this thesis, some themes worth pursuing are presented below.

#### *Data centre waste heat use*

Whilst the thesis has laid a foundation and offered many possibilities to make this happen, actual work needs to be carried out, and, eventually, scaled up. Researchers must monitor – and, to make progress, perhaps partly lead – this progress.

#### *Data centre placement*

The framework offered for data centre placement may be useful already, but academic and industrial collaboration could make it even more valuable. This work would require a number of people with different backgrounds and insights, from engineering, computer science, anthropology, sustainability science, management and other disciplines.

#### *Data centres and related energy technologies*

Both OTEC-integration and salt hydrate “charging” have potential to benefit the world through an energy security perspective, and OTEC has the potential to provide equatorial nations with clean energy and drinking water. How – and where – these technologies can bond needs further investigation. Evaluating data centre-dried seaweed for biofuel production is another (relatively inexpensive) topic to investigate,

and so is the connection between wooden pellet production and direct carbon capture when burning the pellets.

#### *Consumer behaviour*

Another topic, not discussed in this thesis but worth investigation, is how consumer behaviour studies and human-computer interaction (HCI) can benefit data centre sustainability. As discussed, data centres and transmission networks use up approximately one percent of world electricity each. With greater awareness amongst end-users of the energy cost of streaming and online computation, we may be able to limit the accelerating excessiveness of such services.

#### *Sustainability matters of data centre buildings*

Data centre buildings (as well as the servers contained within them) have a certain degree of embodied energy. Embodied energy, in turn, is only one part of the total set of sustainability aspects relating to the building itself. Another is the construction material, typically cement. Building a data centre with a wooden frame can greatly reduce the carbon footprint of the data centre (see section 5.2.1). How to decommission data centre buildings responsibly should also be further explored to obtain a fuller picture of data centre sustainability.

#### *Sustainability matters of data centre internals*

As mentioned in section 1.5, there are other aspects of data centre sustainability, not least water use, land use and the lifecycle of building elements. E-waste from data centres is another source of concern. A large data centre contains thousands of servers, which will eventually need replacement. Producing and recycling these servers come with an environmental cost, and discarding of them unsustainably is equally bad. Every year, the world produces around 50 million tons – the equivalent of 7,000 Eiffel towers

(Eiffel 1889) – of e-waste (Rautela et al. 2021), of which more than four fifths end up illegally in dump sites in, often, low-income countries (Rautela et al. 2021). Further, mining materials and manufacturing components for computers are associated with socio-environmental problems (Boluda, Patitsas, and McMahan 2021). Cooling systems and energy backup solutions are other parts of a data centre that can benefit from additional research endeavours.

#### *Data centre sustainability metrics*

The work on this topic must continue. Hopefully, DESS or a modification of it can serve as a default metric for data centre placement within a societal context. However, new metrics are also needed for the data centre itself. For example, further work is needed to, as Sego et al. write, design new metrics at the data centre level “*that can bridge the gap between unitless data center infrastructure efficiency metrics and component-level performance metrics*” (Sego et al. 2012). Integrating the ideas of Tozer et al. (2018) with contextualised metrics for energy supplied and reused, may be one way forward.

#### *Understanding the fear of data centre managers*

As the thesis has shown, data centre waste heat can be used in a number of ways, and generate profit for the data centre owner. Still, there is a tremendous amount of inertia in the industry. One may wonder why this is – after all, there is money to be made from the selling of heat, and of selling dehydration equipment on a grand scale. It could be that for data centre owners, it is conceptually difficult to look beyond the data centre, and for data centre suppliers, there is more money in advanced cooling equipment than in building piping to a dehydration facility next door. There exists some literature on the subject (e.g. (Klemick, Kopits, and Wolverton 2019)), but not much, and not very

recent. This research, with its connections to knowledge management, needs further study.

#### *Case studies from implementations of concepts*

The thesis contains three case studies. However, the lack of real data and the novel uses for data centre waste heat have made it impossible to demonstrate cases implemented based on the proposed concepts. Especially with the climate emergency in mind, swift implementation of some of them is desired. Should this happen, carrying out multi-disciplinary research on the outcomes of the implementation, related to the three actors of society as well as to nature, would be welcome.

#### *The future of data centres*

The data centre market is moving rapidly, and with new technologies, climate worries and high energy prices, the data centre energy context changes as well. How the data centre industry's energy growth is being shaped today and how this may change in the future are worth further study, especially in the light of future climate impacts and population growth. In essence, both the future of the rapidly changing data centre market and the server and data centre technologies are worth exploration, and so is the relationship between these topics.

### **10.4 Final remarks**

Much scientific thought and effort have gone into trying to prove climate change. Still, the idea that humanity has *not* been able to alter the climate is absurd. In truth, we have managed to take our chosen path of eternally increased consumption using fossil fuel, and – against better knowledge – we still manage to increase the extraction of fossil fuel by about fifteen percent per decade (U.S. Energy Information Administration 2020).

Indeed, burning massive amounts of coal and almost a *hundred million* barrels of oil and fossil gas equivalents yearly, does things to the atmosphere. Oh – was that a hundred million barrels per year? Well, sadly no. According to the EIA (U.S. Energy Information Administration 2022), in addition to large quantities of coal, we use up a hundred million barrels of oil and fossil gas equivalents... *per day*.

In light of this situation, and of the dire need to lessen the strain on our planet, the prospects of commodity dehydration with data centre waste heat are promising. The question is no longer whether the concept works, but what commodity to use in what context: one needs to match suitable commodities with technological advancements, ICT needs, local financial structures and political incentives to enable this process.

The perspective of the thesis means that data centre queries are superimposed on existing societal matters. It should be said though, that the perspective may be flipped. Indeed, instead of looking for uses for waste heat, stakeholders of municipality infrastructure design should ask themselves if they have industries, homes or communal enterprises with a potential use for heat. If they do, there is an indication that they could benefit from the establishment of a data centre (or other industry that generates waste heat). Doing so would show faith in industrial symbiosis and in the circular economy, in which a data centre is but one of many parts.

Further, waste heat use may be the tie that binds the materiality dichotomies of data centres. In this vision, local meets global, and the data centre is brought out of the shadows. One of its outputs, processed data, may be just as intangible as before, but the work performed can now be felt by the worker at the dehydration plant. In addition, the data centre has found its place within the society: rather than an obstacle – or even an

energy-hungry monstrosity – the data centre is now also an enabler, of education, health, interaction and prosperity.

A world relying on sustainable data centres requires large amounts of renewable energy. But as this thesis has shown, when data centres become net-zero users, they provide a fresh perspective on the clean energy infrastructure. Thus, to conclude, it is recommended that data centre owners and legislators look into the possibility of data centre waste heat use. Few, if any, other energy savings can be so easily achieved, and bring the societal gains, that data centre waste heat use can.

\*

The thesis opened with a quote from Rachel Carson’s ominous *Silent Spring*, the catalyst for environmental research that led to so much scientific work – this thesis included. A few hundred pages later, one can conclude that a multi-disciplinary view on information technology can reveal solutions to some of the systematic problems Carson so astutely identified and discussed.

A thesis is shaped like an hourglass: starting big, funnelling down to the very specifics of the core problems, and then expanding to some sort of general finding and reasoning. The running sand within the hourglass may serve as a metaphor too, for the time passing whilst undertaking one’s long – yet never long enough – PhD studies. The sand running through this thesis has come in many colours, and worries concerning the world’s climate underpin its reasons for being. In its very final grains, there is hope. It is a quote from the British renaissance playwright, poet and academic Ben Johnson’s *Discoveries* (1641):

*“I cannot thinke Nature is so spent, and decay’d, that she can bring forth nothing worth her former yeares. She is alwayes the same, like her selfe: And when she collects her strength, is abler still. Men are decay’d, and studies: Shee is not.”*



# Post scriptum: a long and winding road?

In June 2020, a few months after covid hit and all research communication went online, I attended an industry webinar on data centre sustainability. The moderator noted there were a few spaces left at the roundtable discussion that followed. I jumped the opportunity and found myself “sitting” amongst company representatives of the key players in the data centre construction business.

While I sat quiet, the others discussed the galloping energy use and what to do about the problem. I eventually joined the conversation, said I was looking into data centre waste heat use... in particular waste heat use in warm regions.

*“It’s not possible,” said the sustainability manager of a major company.*

*“Oh, I aim to prove it is.”*

*“How?”*

*“Commodity dehydration. Tea leaves in Indonesia, coffee in Brazil, fodder in Australia; the list goes on.”*

*“Commodity dehydration,” the man echoed softly, almost whispering. I’ve never thought of that.”*

That was the exact point in time when I figured this research could be valuable, but also novel: had these ideas been implemented already, the persons in that particular room would have known about it.

Usually, it takes some time for academic research to reach company executives, so the instantaneousness of this discussion made it a rare event. Alas, though I and RISE independently planted the commodity dehydration seed a few years ago, substantial implementation of the idea has yet to be carried out by the industry. Still, the findings of this thesis point to the viability of the concept. If nothing else, the following two-decade-old words of a few Berkley researchers discussing data centre waste heat should still echo among data centre owners: “billions of dollars are left on the table with each year of operation” (Mills et al. 2008). Also, the EU EED and the German Efficiency Law will certainly help, and hopefully soon also the American Federal Data Center Enhancement Act.

So, in short, I hope this thesis, and derivates thereof, can help. The current energy crisis the world faces, and the climate change difficulties hanging over us, need our immediate attention. Starting new coal mines in Germany is a step in the wrong direction. Extending the use for one or a few percent of world electricity would be a major step in the right one.

# References

- Abraham, KM. 2020. "How comparable are sodium-ion batteries to lithium-ion counterparts?" *ACS Energy Letters* 5 (11): 3544-3547.
- Adame, Berhan Oumer. 2021. "The Ethiopian telecom industry: gaps and recommendations towards meaningful connectivity and a thriving digital ecosystem." *Heliyon* 7 (10).
- Adams, M, and AE Ghaly. 2007. "Maximizing sustainability of the Costa Rican coffee industry." *Journal of Cleaner Production* 15 (17): 1716-1729.
- Agarwal, Anup, Jinghan Sun, Shadi Noghabi, Srinivasan Iyengar, Anirudh Badam, Ranveer Chandra, Srinivasan Seshan, and Shivkumar Kalyanaraman. 2021. "Redesigning Data Centers for Renewable Energy." Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks.
- Agliazanov, Ramil, Muhammed Sit, and Ibrahim Demir. 2020. "Hydrology@ Home: a distributed volunteer computing framework for hydrological research and applications." *Journal of Hydroinformatics* 22 (2): 235-248.
- Agrawal, Adarsh, and Ravindra S Kulkarni. 2019. "SETI@ home: A Detailed Analysis and Study." 70th International Astronautical Congress (IAC), Washington, DC.
- Agusdinata, Datu Buyung, Wenjuan Liu, Hallie Eakin, and Hugo Romero. 2018. "Socio-environmental impacts of lithium mineral extraction: towards a research agenda." *Environmental Research Letters* 13 (12): 123001.
- Alderman, Liz, and Monika Pronczuk. 2022. "Europe Plans to Say Nuclear Power and Natural Gas Are Green Investments." *New York Times*, 2022-01-02, 2022. Accessed 2022-01-22. <https://www.nytimes.com/2022/01/02/business/europe-green-investments-nuclear-natural-gas.html>.
- Aleksandrova, Mariya, Sascha Balasko, Markus Kaltenborn, Daniele Malerba, Peter Mucke, Oliver Neuschäfer, Katrin Radtke, Ruben Prütz, Christoph Strupat, Daniel Weller, and Nicola Wiebe. 2021. *World Risk Report 2021*. (Ruhr: Bündnis Entwicklung Hilft, Ruhr University Bochum – Institute for International Law of and Peace and Armed Conflict (IFHV)).

- Arce, Victor Julio Chavez, Raul Raudales, Rich Trubey, David I King, Richard B Chandler, and Carlin C Chandler. 2009. "Measuring and managing the environmental cost of coffee production in Latin America." *Conservation and Society* 7 (2): 141-144.
- Arrhenius, Svante. 1896. "On the Influence of Carbonic Acid in the Air upon the Temperature of the Ground." *Philosophical Magazine and Journal of Science* 41 (5): 237-276.
- Ars Technica. 2018. "Satya Nadella: The cloud is going to move underwater." Ars Technica. Accessed 2022-04-04. <https://arstechnica.com/gadgets/2018/11/satya-nadella-the-cloud-is-going-to-move-underwater/>.
- Ascierto, Rhonda, and Andy Lawrence. 2020. *Uptime Institute global data center survey 2020*. Uptime Institute.
- ASEAN Coordinating Centre for Humanitarian Assistance on disaster management, and Japan International Cooperation Agency. 2015. *Country Report Singapore. Natural Disaster Risk Assessment and Area Business Continuity Plan Formulation for Industrial Agglomerated Areas in the ASEAN Region*. (Jakarta: ASEAN Coordinating Centre for Humanitarian Assistance on disaster management and the Japan International Cooperation Agency).
- ASHRAE Technical Committee. 2016. *Data Center Power Equipment Thermal Guidelines and Best Practices*. (Peachtree Corners, GA: Refrigerating and Air-Conditioning Engineers The American Society of Heating).
- Atteridge, Aaron, and Georgia Savvidou. 2019. "Development aid for energy in Small Island Developing States." *Energy, Sustainability and Society* 9 (10). <https://doi.org/https://doi.org/10.1186/s13705-019-0194-3>.
- AUC/OECD. 2021. *Africa's Development Dynamics 2021: Digital Transformation for Quality Jobs*. African Union Commission and The Organisation for Economic Co-operation and Development (Addis Ababa, Paris).
- Avelar, Victor, Dan Azevedo, and Alan French. 2012. *PUE: a comprehensive examination of the metric*. The Green Grid. <https://www.thegreengrid.org/en/resources/library-and-tools/20-PUE:-A-Comprehensive-Examination-of-the-Metric>.
- Azzman Shariffadeen, Tengku Mohd. 1988. *Microelectronics, information technology and society*. Kuala Lumpur: Malaysian Institute of Microelectronic Systems.
- Bai, Hanyu, and Ziyu Song. 2023. "Lithium-ion battery, sodium-ion battery, or redox-flow battery: A comprehensive comparison in renewable energy systems." *Journal of Power Sources* 580. <https://doi.org/https://doi.org/10.1016/j.jpowsour.2023.233426>.
- Bailey, Kenneth D. 1994. *Sociology and the new systems theory: Toward a theoretical synthesis*. Suny Press.
- Baltzer, Harald. 2010. "Wikileaks flyttar till 'kärnvapensäker' anläggning." *Computer Sweden*, 2010-08-30, 2010.
- Bansode, Pratik V, Jimil M Shah, Gautam Gupta, Dereje Agonafer, Harsh Patel, David Roe, and Rick Tufty. 2018. "Measurement of the Thermal Performance of a Single-Phase Immersion Cooled Server at Elevated Temperatures for Prolonged Time." International Electronic Packaging Technical Conference and Exhibition.
- Barakat, Matthew. 2022. "As data centers proliferate, neighbors knock the noise." AP News. Accessed 2023-06-26. <https://apnews.com/article/technology-virginia-web-services-cloud-computing-fce7d5c5b5eb25a5989171c686578047>.

## References

- Barroso, Luiz André, Urs Hölzle, and Parthasarathy Ranganathan. 2019. *The Datacenter as a Computer: Designing Warehouse-Scale Machines, Third Edition*. Edited by Margaret Martonosi. *Synthesis lectures on computer architecture*: Morgan & Claypool.
- Belady, Christian L, and Christopher G Malone. 2007. "Metrics and an infrastructure model to evaluate data center efficiency." International Electronic Packaging Technical Conference and Exhibition.
- Berners-Lee, Mike. 2021. *There Is No Planet B: A Handbook for the Make or Break Years Updated Edition*. Cambridge University Press.
- Bibri, Simon Elias. 2018. "A foundational framework for smart sustainable city development: Theoretical, disciplinary, and discursive dimensions and their synergies." *Sustainable Cities and Society* 38.
- Birbarah, Patrick, Tarek Gebrael, Thomas Foulkes, Andrew Stillwell, Alexandra Moore, Robert Pilawa-Podgurski, and Nenad Miljkovic. 2020. "Water immersion cooling of high power density electronics." *International Journal of Heat and Mass Transfer* 147: 118918.
- Blackwell, Alan F. 2021. "Ethnographic artificial intelligence." *Interdisciplinary Science Reviews* 46 (1-2). <https://doi.org/https://doi.org/10.1080/03080188.2020.1840226>.
- Boluda, Inès Moreno, Elizabeth Patitsas, and Peter McMahan. 2021. "What do Computer Scientists Know About Conflict Minerals?" Workshop on Computing within Limits.
- Bonds, Eric, and Liam Downey. 2012. "Green technology and ecologically unequal exchange: The environmental and social consequences of ecological modernization in the world-system." *Journal of World-Systems Research*: 167-186.
- Boulding, Kenneth E. 1966. "The Economics of the Coming Spaceship Earth." In *Environmental Quality in a Growing Economy* edited by H. Jarrett, 3-14. Baltimore, MD: Resources for the Future/Johns Hopkins University Press.
- Bradford, Kent J, Peetambar Dahal, Johan Van Asbrouck, Keshavulu Kunusoth, Pedro Bello, James Thompson, and Felicia Wu. 2018. "The dry chain: Reducing postharvest losses and improving food safety in humid climates." *Trends in Food Science & Technology* 71: 84-93.
- Brown, Alan S. 2017. "Heat sink sunk." *Mechanical Engineering*.
- Brundtland, Gro Harlem, M Khalid, S Agnelli, S Al-Athel, and BJNY Chidzero. 1987. "Our common future." *New York* 8.
- Bryce, Robert. 2020. "How Google Powers Its 'Monopoly' With Enough Electricity For Entire Countries." *Forbes*, 2020-10-21, 2020. Accessed 2021-07-02. [www.forbes.com/sites/robertbryce/2020/10/21/googles-dominance-is-fueled-by-zambia-size-amounts-of-electricity](http://www.forbes.com/sites/robertbryce/2020/10/21/googles-dominance-is-fueled-by-zambia-size-amounts-of-electricity).
- Bryman, Alan. 2006. "Integrating quantitative and qualitative research: how is it done?" *Qualitative research* 6 (1): 97-113.
- . 2007. "Barriers to integrating quantitative and qualitative research." *Journal of mixed methods research* 1 (1): 8-22.
- Brückner, Sarah, Selina Liu, Laia Miró, Michael Radspieler, Luisa F Cabeza, and Eberhard Lävemann. 2015. "Industrial waste heat recovery technologies: An economic analysis of heat transformation technologies." *Applied Energy* 151: 157-167. <https://doi.org/http://dx.doi.org/10.1016/j.apenergy.2015.01.147>.

- Bundesministerium für Wirtschaft und Klimaschutz. 2023. Entwurf eines Gesetzes zur Steigerung der Energieeffizienz und zur Änderung des Energiedienstleistungsgesetzes. Berlin: Bundesministerium für Wirtschaft und Klimaschutz.
- Burck, Jan, Thea Uhlich, Christoph Bals, Niklas Höhne, Leonardo Nascimento, Monica Tavares, and Elisabeth Strietzel. 2022. *Climate Change Performance Index 2023*. (Bonn: Germanwatch).
- Busby, Joshua W, and Josh Busby. 2007. *Climate change and national security: an agenda for action*. Vol. 32: Council on Foreign Relations Press.
- Bush, George W. 2001. "President Bush Discusses Global Climate Change." The White House. Accessed 2023-06-08. <https://georgewbush-whitehouse.archives.gov/news/releases/2001/06/20010611-2.html>.
- Cambridge University Press. n.d. Culture. In *Cambridge Dictionary*. Cambridge: Cambridge University Press.
- Capozzoli, Alfonso, Marta Chinnici, Marco Perino, and Gianluca Serale. 2014. "Review on performance metrics for energy efficiency in data center: The role of thermal management." International Workshop on Energy Efficient Data Centers.
- Central Intelligence Agency. 2023. "The World Factbook: Central America:: Costa Rica." Central Intelligence Agency. Accessed 2023-04-01. <https://www.cia.gov/the-world-factbook/countries/costa-rica/>.
- Chalise, Santosh, Amir Golshani, Shekhar Raj Awasthi, Shanshan Ma, Bijen Raj Shrestha, Labi Bajracharya, Wei Sun, and Reinaldo Tonkoski. 2015. "Data center energy systems: Current technology and future direction." 2015 IEEE Power & Energy Society General Meeting.
- Chandel, Sonali, Zang Jingji, Yu Yunnan, Sun Jingyao, and Zhang Zhipeng. 2019. "The Golden Shield Project of China: A Decade Later - An in-Depth Study of the Great Firewall." International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), Guilin, China.
- Cheng, Huiwen, Bo Liu, Weiwei Lin, Zehua Ma, Keqin Li, and Ching-Hsien Hsu. 2021. "A survey of energy-saving technologies in cloud data centers." *The Journal of Supercomputing* 77. <https://doi.org/10.1007/s11227-021-03805-5>.
- Chilukuri, MV, Masliza Mohd Dahlan, and Chan Chuey Hwye. 2018. "Benchmarking Energy Efficiency in Tropical Data Centres—Metrics and Measurements." 2018 International Conference and Utility Exhibition on Green Energy for Sustainable Development (ICUE).
- Churchman, Charles West. 1967. "Wicked problems." *Management Science* 14 (4).
- Clidas, Jimmy, David W. Stiver, and William Hamburg. 2009. Water-based Data Center [Google]. USA US 7,525,207 B2.
- Conapto, 2022-10-04, 2022, "Conapto adds 20MW data center capacity in Stockholm," <https://www.conapto.com/conapto-adds-20mw-data-center-capacity-in-stockholm/>.
- Cotterill, Joseph. 2021. "Cabling Africa: the great data race to serve the 'last billion'." *Financial Times*, 2021-01-31, 2021. <https://www.ft.com/content/adb1130e-2844-4051-b1df-a691fc8a19b8>.
- Crawford, Kate. 2021. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven and London: Yale University Press.
- Creswell, John W, and Amanda L Garrett. 2008. "The "movement" of mixed methods research and the role of educators." *South African journal of education* 28 (3): 321-333.

## References

- Cutler, Benjamin F, Norman Ashton Whitaker, Spencer G Fowers, and Jeffrey Alex Kramer. 2016. Artificial reef datacenter.
- Daily, Ian. 2010. "Nuclear bunker houses world's toughest server farm." *Wired*.
- Danielsen Sorup, Hjalte Jomo, Ole Fryd, Li Liu, Karsten Arnbjerg-Nielsen, and Marina Jensen. 2019. "An SDG-based framework for assessing urban stormwater management systems."
- Data Centre Magazine. 2021. "Data centres are drinking the desert dry." *Data Centre Magazine*.
- David Harman V, Thomas. 2008. "Waste heat recovery in datacenters: ejector heat pump analysis." School of Mechanical Engineering, Georgia Institute of Technology.
- Davies, G F, G G Maidment, and R M Tozer. 2016. "Using data centres for combined heating and cooling: An investigation for London." *Applied Thermal Engineering* 94: 296-304.
- Davies, Paul. 2019. *The Demon in the Machine*. London: Penguin Random House.
- Davis, Jacqueline, Daniel Bizo, Andy Lawrence, Owen Rogers, and Max Smolaks. 2022. *Global Data Center Survey 2022*. (New York: Uptime Institute).
- Dayarathna, Miyuru, Yonggang Wen, and Rui Fan. 2016. "Data Center Energy Consumption Modeling: A Survey." *IEEE Communications Surveys & Tutorials* 18 (1): 732-794.
- De Napoli, Carmine, Agostino Forestiero, Demetrio Laganà, Giovanni Lupi, Carlo Mastroianni, and Leonardo Spataro. 2016. "Efficiency and green metrics for distributed data centers." *Report P-26, ICAR*.
- Deemer, Bridget R., John A. Harrison, Siyue Li, Jake J. Beaulieu, Tonya DelSontro, Nathan Barros, José F. Bezerra-Neto, Stephen M. Powers, Marco A. dos Santos, and J. Arie Vonk. 2016. "Greenhouse Gas Emissions from Reservoir Water Surfaces: A New Global Synthesis." *BioScience* 66 (11): 949-964. <https://doi.org/https://doi.org/10.1093/biosci/biw117>.
- Dencer-Brown, AM, RM Jarvis, AC Alfaro, and S Milne. 2021. "The Mixed Methods Practical Sustainability Research Framework: An Illustration From Research on the Creeping Problem of Coastal Complexity and Mangrove Management." *Journal of Mixed Methods Research*.
- Department of the Environment Climate and Communications. 2021. *Climate Action Plan 2021: Securing Our Future*. (Dublin: Government of Ireland).
- Deschenes, P. J., and Marian Chertow. 2004. "An island approach to industrial ecology: towards sustainability in the island context" *Journal of Environmental Planning and Management* 47 (2).
- Dietrich, Jay. 2023. "A wake up call: the EU Efficiency Directive reporting will affect you." Datacenter Dynamics. Accessed 2023-05-31. <https://www.datacenterdynamics.com/en/opinions/a-wake-up-call-the-eu-efficiency-directive-reporting-will-affect-you/>.
- Diguet, Cécile, and Fanny Lopez. 2019. *The spatial and energy impact of data centers on the territories*. (Angers: Synthesis).
- Dourish, Paul. 2022. *The Stuff of Bits: An Essay on the Materialities of Information*. MIT Press.
- Drucker, Peter F. 1969. *The age of discontinuity: Guidelines to our changing society*. New York, NY: Harper & Row.
- Dybdal Christensen, Jakob, Jens Therkelsen, Ivo Georgiev, and Henrik Sand. 2018. *Data centre opportunities in the Nordics: An analysis of the competitive advantages*. (Copenhagen: Nordiska ministerrådet).

- Ebrahimi, Khosrow, Gerard F Jones, and Amy S Fleischer. 2014. "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities." *Renewable and Sustainable Energy Reviews* 31: 622-638.
- EcoDataCenter. 2021. "Why EcoDataCenter?". EcoDataCenter. Accessed 2022-04-04. <https://ecodatacenter.se/sites/main-site/>.
- . 2022. "EcoDataCenter in Falun secures a future power supply of 80 MW from renewable sources." EcoDataCenter. Accessed 2022-04-04.
- . n.d. "Sustainability that really makes a difference." Accessed 2022-07-29. <https://ecodatacenter.se/sustainability/>.
- Economic Planning Unit. 2021. *Twelfth Malaysia Plan*. (Putrajaya: Economic Planning Unit).
- Eiffel, Gustave. 1889. "The Eiffel Tower." *The New Review* 1 (2).
- Eiland, Richard, John Fernandes, Marianna Vallejo, Dereje Agonafer, and Veerendra Mulay. 2014. "Flow Rate and inlet temperature considerations for direct immersion of a single server in mineral oil." Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm).
- Elbow, Gary S., Franklin D. Parker, Charles L. Stansifer, and Thomas L. Karnes. 2024. "Costa Rica." Encyclopedia Britannica. Accessed 2024-04-14. <https://www.britannica.com/place/Costa-Rica>.
- Energimarknadsinspektionen. 2016. *Åtgärder för ökad efterfrågeflexibilitet i det svenska elsystemet*. Energimarknadsinspektionen (Eskilstuna).
- Enterprise Ireland. 2022. *Towards net-zero construction in energy intensive sectors. Experiences from the decarbonisation of mission-critical facilities*. (Dublin: Enterprise Ireland).
- European Environment Agency. 2021. "Municipal waste landfill rates in Europe by country." European Environment Agency. Accessed 2021-11-01. <https://www.eea.europa.eu/data-and-maps/figures/municipal-waste-landfill-rates-in>.
- European Parliament. 2022. Amendments by the European Parliament to the Commission proposal Proposal for a directive of the European Parliament and of the council on energy efficiency (recast). Brussels: European Parliament.
- . 2023a. Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 (recast). Strasbourg: European Parliament.
- . 2023b. "Energy efficiency directive." Accessed 2023-05-30. [https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive\\_en](https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en).
- . 2023c. Proposal for a Directive of the European Parliament and of the Council on Energy Efficiency (recast) - Analysis of the final compromise text with a view to agreement. Brussels: European Parliament.
- Evans, David M, Rorie Parsons, Peter Jackson, Sarah Greenwood, and Anthony Ryan. 2020. "Understanding plastic packaging: The co-evolution of materials and society." *Global Environmental Change* 65. <https://doi.org/https://doi.org/10.1016/j.gloenvcha.2020.102166>.
- Fan, Yuqi, Hongli Ding, Lusheng Wang, and Xiaojing Yuan. 2016. "Green latency-aware data placement in data centers." *Computer Networks* 110: 46-57.



## References

- Foley, Richard. 2018. *The Geography of Insight: The Sciences, the Humanities, How They Differ, Why They Matter*. Oxford University Press.
- Freitag, Charlotte, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S Blair, and Adrian Friday. 2021. "The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations." *Patterns* 2 (9): 100340.
- Fuenmayor, Ramsés. 2001. "The oblivion of Churchman's plea for a systems approach to world problems. I. The inseparability of systems thinking and world issues in the modern epoch." *Systemic Practice and Action Research* 14 (1): 11-28.
- Gonzalez Monserrate, Steven. 2022. "The Cloud Is Material: On the Environmental Impacts of Computation and Data Storage." *MIT Case Studies in Social and Ethical Responsibilities of Computing* Winter 2022. <https://doi.org/https://doi.org/10.21428/2c646de5.031d4553>.
- Google. 2020. *Google Environmental Report 2020*. Google.
- Greenpeace International. 2012. *How Clean is Your Cloud?* (Amsterdam: Greenpeace International).
- Gregson, Nicky, Mike Crang, Sara Fuller, and Helen Holmes. 2015. "Interrogating the circular economy: the moral economy of resource recovery in the EU." *Economy and society* 44 (2): 218-243.
- GSM Association. 2020a. "Data - GSMA Intelligence." GSM Association. Accessed 2023-01-19. <https://www.gsmainelligence.com>.
- . 2020b. "Mobile Connectivity Index." GSM Association. Accessed 2023-01-19. <https://www.mobileconnectivityindex.com/#year=2018&dataSet=indexScore>.
- Guitart, Jordi. 2017. "Toward sustainable data centers: a comprehensive energy management strategy." *Computing* 99. <https://doi.org/https://doi.org/10.1007/s00607-016-0501-1>.
- Hancock, Alice. 2022. "EU parliament votes to designate gas and nuclear as sustainable." *Financial Times*, 2022-07-06, 2022. <https://www.ft.com/content/0df04289-1014-406e-81c7-1e4a6b1ea5bc>.
- Hansson, Staffan. 2006. "Technology and social change: a technological megasystem in the north of Sweden." In *Migration, industrialisation and regionalisation: papers II from the conference the Use and Abuse of History in the Barents Region at Luleå University of Technology, Luleå, Sweden 2004*, edited by Lars Elenius, 20-31. Luleå: Luleå tekniska universitet.
- Harding, Luke. 2015. "The node pole: inside Facebook's Swedish hub near the Arctic Circle." *The Guardian*, 2015. <https://www.theguardian.com/technology/2015/sep/25/facebook-datacentre-lulea-sweden-node-pole>.
- Hoes, Olivier A. C., Lourens J. J. Meijer, Ruud J. van der Ent, and Nick C. van de Giesen. 2017. "Systematic high-resolution assessment of global hydropower potential" *PLoS ONE* 12 (2).
- Holt, Jennifer, and Patrick Vonderau. 2015. "Where the internet lives: Data centers as cloud infrastructure." In *Signal traffic: Critical studies of media infrastructures*, edited by Lisa Parks and Nicole Starosielski, 71-93. Champaign: University of Illinois Press.
- Horner, Nathaniel Charles. 2016. "Powering the information age: metrics, social cost optimization strategies, and indirect effects related to data center energy use." Carnegie Mellon University.
- Hosseini, Seyed Ehsan, and Mazlan Abdul Wahid. 2020. "Hydrogen from solar energy, a clean energy carrier from a sustainable source of energy." *International Journal of Energy Research* 44 (6): 4110-4131.

- Houngbonon, Georges V, Erwan Le Quentrec, and Stefania Rubrichi. 2021. "Access to electricity and digital inclusion: evidence from mobile call detail records." *Humanities and Social Sciences Communications* 8. <https://doi.org/https://doi.org/10.1057/s41599-021-00848-0>.
- Howard, Philip N, Aiden Duffy, Deen Freelon, Muzammil Hussain, Will Mari, and Marwa Mazaid. 2011. *Opening closed regimes: What was the role of social media during the Arab Spring?* . (Seattle, WA: University of Washington).
- Hu, Tung-Hui. 2015. *A prehistory of the cloud*. Cambridge, Massachusetts: MIT Press.
- Huang, Pei, Benedetta Copertaro, Xingxing Zhang, Jingchun Shen, Isabelle Löfgren, Mats Rönnelid, Jan Fahlen, Dan Andersson, and Mikael Svanfeldt. 2020. "A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating." *Applied Energy* 258: 114109.
- Huber, Matthew. 2015. "Theorizing Energy Geographies." *Geography Compass*.
- Huyer, Sophia, Mariola Acosta, Tatiana Gumucio, and Jasmin Irisha Jim Ilham. 2020. "Can we turn the tide? Confronting gender inequality in climate policy." *Gender & Development* 28 (3): 571-591.
- Hwit, Emil. 2019. "Rekommenderad framledningstemperatur i fjärrvärmenät baserat på rökgaskondensering. En beräkningsundersökning av rökgaskondensering och fjärrvärme i en medelstor svensk stad." Akademin för teknik och miljö, Högskolan i Gävle.
- IEA. 2019. *Africa Energy Outlook 2019*. (Paris: IEA).
- . 2022. *Data Centres and Data Transmission Networks*. (Paris: International Energy Agency). <https://www.iea.org/reports/data-centres-and-data-transmission-networks>.
- Industry and Information Division. 2021. *New Data Center Development Three-Year Action Plan (2021-2023)*. (Beijing: Ministry of Industry and Information Technology of the People's Republic of China). [www.mit.gov.cn/jgsj/txs/wjfb/art/2021/art\\_12cc04dc9daf4d57a7038811a57383b6.html](http://www.mit.gov.cn/jgsj/txs/wjfb/art/2021/art_12cc04dc9daf4d57a7038811a57383b6.html).
- Institute for Economics & Peace. 2021. *Global Peace Index 2021: Measuring Peace in a Complex World*. (Sydney: Institute for Economics & Peace).
- International Coffee Organization. 2003. *Sustainable Coffee Executive Summary*.
- . 2019. *Coffee Development Report 2019: Growing for prosperity*. International Coffee Organization (London).
- . 2020. *Coffee Production Report April 2020*.
- International Monetary Fund. 2021. *Regional Economic Outlook. Sub-Saharan Africa: one planet, two worlds, three stories*. (Washington, DC: International Monetary Fund).
- . 2022. *World Economic Outlook database: April 2022*. Washington, DC.
- International Organization for Standardization. 2016a. *ISO/IEC 30134-2:2016(en) Information technology - Data centres - Key performance indicators - Part 2: Power usage effectiveness (PUE)*. Geneva: International Organization for Standardization.
- . 2016b. *ISO/IEC 30134-3:2016(en) Information technology - Data centres - Key performance indicators - Part 3: Renewable energy factor (REF)* Geneva: International Organization for Standardization.

## References

- . 2021. *ISO/IEC 30134-6:2021(en) Information technology - Data centres key performance indicators - Part 6: Energy Reuse Factor (ERF)*. Geneva: International Organization for Standardization.
- IPCC. 2013. "Chapter 8: Anthropogenic and Natural Radiative Forcing." In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by T F Stocker, D Qin, Plattner G-K, M Tignor, S K Allen, J Boschung, A Nauels, Y Xia, Bex V and P M Midgley. Cambridge and New York: Cambridge University Press.
- Ipsen, Heather A. 2018. "Catching the Cloud and Pinning It Down: the Social and Environmental Impacts of Data Centers." Syracuse University.
- Jamalzadeh, Morteza, and Navid Behravan. 2012. "An exhaustive framework for better data centers' energy efficiency and greenness by using metrics." *Indian Journal of Computer Science and Engineering* 2 (6).
- Johansson, Per-Olov, and Bengt Kriström. 2021. "The Costs and Benefits of Supporting Data Centers: A General Equilibrium Analysis." *Center for Environmental and Resource Economics, CERÉ, The Costs and Benefits of Supporting Data Centers: A General Equilibrium Analysis (November 15, 2021)*.
- Johnson, R Burke. 2017. "Dialectical pluralism: A metaparadigm whose time has come." *Journal of Mixed Methods Research* 11 (2): 156-173.
- Johnson, R Burke, A J Onwuegbuzie, and Lisa A Turner. 2007. "Toward a definition of mixed methods research." *Journal of Mixed Methods Research* 1 (2).  
<https://doi.org/https://doi.org/10.1177/1558689806298224>.
- Judge, Peter. 2016. "DCD Enterprise: Microsoft expands underwater data center plans." Data Centre Dynamics. Accessed 2024-01-06. <https://www.datacenterdynamics.com/en/news/dcd-enterprise-microsoft-expands-underwater-data-center-plans/>.
- . 2022a. "Drought-stricken Holland discovers Microsoft data center slurped 84m liters of drinking water last year." Data Centre Dynamics. Accessed 2022-08-23.  
<https://www.datacenterdynamics.com/en/news/drought-stricken-holland-discovers-microsoft-data-center-slurped-84m-liters-of-drinking-water-last-year/>.
- . 2022b. "Work begins on Chinese underwater data center." Data Centre Dynamics. Accessed 2022-03-10. <https://www.datacenterdynamics.com/en/news/work-begins-on-chinese-underwater-data-center/>.
- . 2023. "Could data centers have helped the UK tomato crisis?". Datacenter Dynamics. Accessed 2023-05-02. <https://www.datacenterdynamics.com/en/opinions/could-data-centers-have-helped-the-uk-tomato-crisis/>.
- Kaijser, Anna, and Annica Kronsell. 2014. "Climate change through the lens of intersectionality." *Environmental politics* 23 (3): 417-433.
- Kennedy, Joel, Assane Lo, Haile-Selassie Rajamani, and Saad Lutfi. 2021. "Solar and sand: Dust deposit mitigation in the desert for PV arrays." *Sustainable Energy, Grids and Networks* 28: 100531.
- Kim, Nam Jin, Kim Choon Ng, and Wongee Chun. 2009. "Using the condenser effluent from a nuclear power plant for Ocean Thermal Energy Conversion (OTEC)." *International Communications in Heat and Mass Transfer* 36 (10): 1008-1013.
- King, Anthony. 2023. "Northvolt to bring sodium-ion batteries to European market." Royal Society of Chemistry. Accessed 2024-01-05. <https://www.chemistryworld.com/news/northvolt-to-bring-sodium-ion-batteries-to-european-market/4018576.article>.

- Kleinman, Zoe. 2023. "Tiny data centre used to heat public swimming pool." BBC. Accessed 2023-03-22. <https://www.bbc.com/news/technology-64939558>.
- Klemick, Heather, Elizabeth Kopits, and Ann Wolverton. 2019. "How do data centers make energy efficiency investment decisions? Qualitative evidence from focus groups and interviews." *Energy efficiency* 12 (5): 1359-1377.
- Klingert, Sonja, and Sebastian Szilvas. 2020. "Spinning gold from straw - evaluating the flexibility of data centres on power markets." *Energy Informatics* 3 (7). <https://doi.org/https://doi.org/10.1186/s42162-020-00110-y>.
- Klir, George J. 1991. *Facets of systems science*. Vol. 7. Springer Science & Business Media.
- Koomey, Jonathan, Kenneth Brill, Pitt Turner, John Stanley, and Bruce Taylor. 2008. *A simple model for determining true total cost of ownership for data centers*. (Uptime Institute).
- Korhonen, Santeri. 2018. "Energy Efficiency of Modern Datacenter." School of Science, Aalto University.
- Koronen, Carolina, Max Åhman, and Lars J Nilsson. 2020. "Data centres in future European energy systems—energy efficiency, integration and policy." *Energy Efficiency* 13 (1): 129-144.
- Korpela, Eric, Dan Werthimer, David Anderson, Jeff Cobb, and Matt Leboisky. 2001. "SETI@ home-massively distributed computing for SETI." *Computing in science & engineering* 3 (1): 78-83.
- Kortetmäki, Teea, and Marja Järvelä. 2021. "Social vulnerability to climate policies: Building a matrix to assess policy impacts on well-being." *Environmental Science & Policy* 123: 220-228.
- Kulshrestha, Sudhanshu, and Sanjeev Patel. 2019. "A Study on Energy Efficient Resource Allocation for Cloud Data Center." Twelfth International Conference on Contemporary Computing (IC3), Noida, Noida.
- Kummu, Matti, and Olli Varis. 2011. "The world by latitudes: A global analysis of human population, development level and environment across the north-south axis over the past half century." *Applied geography* 31 (2): 495-507.
- Laan, Marc. 2008. "Stroomslurpers aan banden." *Het Parool*, 2008-03-26, 2008. Accessed 2022-03-28. <https://www.parool.nl/nieuws/stroomslurpers-aan-banden~b49d6bc8/>.
- Lajevardi, Babak, Karl R Haapala, and Joseph F Junker. 2014. "An energy efficiency metric for data center assessment." IIE Annual Conference. Proceedings.
- Lannelongue, Loïc, Jason Grealey, and Michael Inouye. 2021. "Green algorithms: Quantifying the carbon footprint of computation." *Advanced Science* 8 (12): 2100707.
- Larman, Craig, and Victor R Basili. 2003. "Iterative and incremental developments. a brief history." *Computer* 36 (6): 47-56.
- Larsson, Mats R. 2012. "Development Opportunities for Well-Established Technologies." In *The Business of Global Energy Transformation*, 148-158. Springer.
- Lehdonvirta, Vili. 2022. *Cloud Empires: How Digital Platforms Are Overtaking the State and How We Can Regain Control* Cambridge, MA: MIT Press.
- Lei, Nuo, and Eric Masanet. 2020. "Statistical analysis for predicting location-specific data center PUE and its improvement potential." *Energy* 201: 117556.
- Levin, II, AI Dordopulo, YI Doronchenko, MK Raskladkin, AM Fedorov, and ZV Kalyaev. 2016. "Immersion liquid cooling FPGA-based reconfigurable computer system." *IFAC-PapersOnLine* 49 (25): 366-371.

## References

- Levy, Moises. 2019. "New family of data center metrics using a multidimensional approach for a holistic understanding." Faculty of College of Engineering and Computer Science, Florida Atlantic University.
- Li, Xiang, Peter Garraghan, Xiaohong Jiang, Zhaohui Wu, and Jie Xu. 2018. "Holistic Virtual Machine Scheduling in Cloud Datacenters towards Minimizing Total Energy." *IEEE Transactions on Parallel and Distributed Systems* 29 (6).
- Lin, Yi, Xiaojun Duan, Chengli Zhao, and Li Da Xu. 2012. *Systems science: methodological approaches*. CRC press.
- Live Blackpool. 2021. "Blackpool connects to the North Atlantic Loop." Live Blackpool. Accessed 2023-01-16. <https://www.liveblackpool.info/about/history/blackpool-connects-to-the-north-atlantic-loop/>.
- Lockard, Craig A., Ooi Jin Bee, Thomas R. Leinbach, and Zakaria Bin Ahmad. 2024. "Malaysia." Encyclopedia Britannica. Accessed 2024-04-14. <https://www.britannica.com/place/Malaysia>.
- Lombardi, D Rachel, and Peter Laybourn. 2012. "Redefining industrial symbiosis: Crossing academic-practitioner boundaries." *Journal of Industrial Ecology* 16 (1): 28-37.
- Lu, Yi, and Jiuping Xu. 2017. "Cybernetic paradigm based innovative approaches towards coping with climate change." *Journal of Systems Science and Systems Engineering* 26 (3): 359-382.
- Lucivero, Federica. 2019. "Big Data, Big Waste? A Reflection on the Environmental Sustainability of Big Data Initiatives." *Science and Engineering Ethics*. <https://doi.org/https://doi.org/10.1007/s11948-019-00171-7>. <https://doi.org/10.1007/s11948-019-00171-7>.
- López-Hortas, Lucía, Maxine Gely, Elena Falqué, Herminia Domínguez, and María Dolores Torres. 2019. "Alternative environmental friendly process for dehydration of edible *Undaria pinnatifida* brown seaweed by microwave hydrodiffusion and gravity." *Journal of Food Engineering* 261.
- Macdonald, Lee T. 2020. "Proposals to Move the Royal Observatory, Greenwich, 1836–1944." *Journal for the History of Astronomy* 51 (3): 272-304.
- Mahbod, Muhammad Haiqal Bin, Chin Boon Chng, Poh Seng Lee, and Chee Kong Chui. 2022. "Energy saving evaluation of an energy efficient data center using a model-free reinforcement learning approach." *Applied Energy* 322. <https://doi.org/https://doi.org/10.1016/j.apenergy.2022.119392>.
- Malone, Christopher, and Christian Belady. 2006. "Metrics to characterize data center & IT equipment energy use." Proceedings of the Digital Power Forum. Dallas, TX: September.
- Manganelli, Matteo, Alessandro Soldati, Luigi Martirano, and Seeram Ramakrishna. 2021. "Strategies for Improving the Sustainability of Data Centers via Energy Mix, Energy Conservation, and Circular Energy." *Sustainability* 13.
- Mansoor, Sanya. 2023. "Greta Thunberg Detained in Germany as Police Crack Down on Coal Mine Protest." *Time Magazine*, 2023-01-17, 2023. <https://time.com/6247769/greta-thunberg-detained-germany-coal/>.
- Masanet, Eric, Arman Shehabi, Nuo Lei, Sarah Smith, and Jonathan Koomey. 2020. "Recalibrating global data center energy-use estimates." *Science* 367 (6481): 984-986.
- McCammon, Sarah. 2017. "Google Moves In And Wants To Pump 1.5 Million Gallons Of Water Per Day." *National Public Radio*, 2017.

- Mendonca, Karl. 2022. "Beyond Representation: Using Infrastructure Studies to Reframe Ethnographic Agendas and Outcomes." *Ethnographic Praxis in Industry Conference*, Amsterdam.
- Meyer, J, and R Jepperson. 2000. "The "Actors" of Modern Society: The Cultural Construction of Social Agency." *Sociological Theory* 18 (1).
- Microsoft. 2019. "Project Natick Phase 2." Microsoft. Accessed 2021-01-30.  
<https://natick.research.microsoft.com>.
- Miller, Rich. 2008. "Data Center Used to Heat Swimming Pool." Accessed 2023-01-22.  
<https://www.datacenterknowledge.com/archives/2008/04/02/data-center-used-to-heat-swimming-pool>.
- Mills, Evan, Gary Shamshoian, Michele Blazek, Phil Naughton, Robert S Seese, William Tschudi, and Dale Sartor. 2008. "The business case for energy management in high-tech industries." *Energy Efficiency* 1 (1).
- Ministry of Energy Green Technology and Water Malaysia (KeTTHA). 2017. *Green Technology Master Plan Malaysia 2017-2030*. (Putrajaya: Green Technology and Water (KeTTHA) Ministry of Energy).
- Miramonti, Lino. 2009. "Solar neutrino detection." AIP Conference Proceedings.
- Miyamoto, Motoe, Mamat Mohd Parid, Zakaria Noor Aini, and Tetsuya Michinaka. 2014. "Proximate and underlying causes of forest cover change in Peninsular Malaysia." *Forest Policy and Economics* 44.
- Mohamed, Abdul Razak Naina, Mohammad Hadi Abd Halim, and Mohd Helmy Zakaria. 2021. "Well to Wheel Comparison Study between the Electric Vehicle (Tesla Model S 100D) and Internal Combustion Engine Vehicle (Hyundai Sonata 2.4 L)." *Politeknik & Kolej Komuniti Journal of Engineering and Technology* 6 (1): 80-89.
- Montiel, Ivan, and Javier Delgado-Ceballos. 2014. "Defining and measuring corporate sustainability: Are we there yet?" *Organization & Environment* 27 (2): 113-139.
- Moore, Brandon. 2021. "Simpler (and better) than you think. Demystifying immersion-cooled data centers". Green Revolution Cooling. Accessed 2022-02-17.  
<https://www.grcooling.com/blog/simpler-and-better-than-you-think-demystifying-immersion-cooled-data-centers/>.
- Morgan, Kelly. 2022. "Improving datacenter efficiency in Europe – the role of PUE." S&P Global. Accessed 2023-03-18. <https://www.spglobal.com/marketintelligence/en/news-insights/research/improving-datacenter-efficiency-in-europe-the-role-of-pue>.
- Moss, Sebastian. 2022. "Facebook parent Meta officially kills Zeewolde data center after pushback in the Netherlands." Datacenter Dynamics. Accessed 2023-01-16.  
<https://www.datacenterdynamics.com/en/news/facebook-parent-meta-officially-kills-zeewolde-data-center-after-pushback-in-the-netherlands/>.
- Mujawar, Anjum, Shanti Sankara Krishnan, Sandhya Kumar, and Apurva Sawant. 2018. "IoT: Green Data Center Strategies." *International Journal on Future Revolution in Computer Science & Communication Engineering* 4 (5): 170-174.
- Mumford, Lewis. 1966. "The First Megamachine." *Diogenes* 14 (55): 1-15.  
<https://doi.org/10.1177/039219216601405501>.
- Murdock, Beth E, Kathryn E Toghil, and Nuria Tapia-Ruiz. 2021. "A Perspective on the Sustainability of Cathode Materials used in Lithium-Ion Batteries." *Advanced Energy Materials* 11 (39).  
<https://doi.org/https://doi.org/10.1002/aenm.202102028>.

## References

- Murray, Alan, Keith Skene, and Kathryn Haynes. 2017. "The circular economy: an interdisciplinary exploration of the concept and application in a global context." *Journal of business ethics* 140 (3): 369-380.
- Myovella, Godwin, Mehmet Karacuka, and Justus Hauca. 2020. "Digitalization and economic growth: A comparative analysis of Sub-Saharan Africa and OECD economies." *Telecommunications Policy* 44. <https://doi.org/https://doi.org/10.1016/j.telpol.2019.101856>.
- Myrén, Karin. 2021. "Här är Sveriges nya Klondike - men därför kan jobben stoppas." *Tidningen Näringslivet*, 2021-04-21, 2021. <https://www.tn.se/fordjupning/har-ar-sveriges-nya-klondike-men-darfor-kan-jobben-stoppas/>.
- Mytton, David. 2021. "Data centre water consumption." *Nature Parter Journals Clean Water* 4 (1): 1-6.
- Nautilus Data Technologies. 2022. "Our Patented Technology." Accessed 2022-06-16. <https://nautilusdt.com/the-technology-and-patents/>.
- Nkulu, Célestin Banza Lubaba, Lidia Casas, Vincent Haufroid, Thierry De Putter, Nelly D Saenen, Tony Kayembe-Kitenge, Paul Musa Obadia, Daniel Kyanika Wa Mukoma, Jean-Marie Lunda Ilunga, and Tim S Nawrot. 2018. "Sustainability of artisanal mining of cobalt in DR Congo." *Nature sustainability* 1 (9): 495-504.
- Noël, John A, Samer Kahwaji, Louis Desgrosseilliers, Dominic Groulx, and Mary Anne White. 2022. "Phase change materials." In *Storing Energy*, 503-535. Elsevier.
- NTT. n.d. *Cyberjaya Data Centers*. (Tokyo: NTT).
- Nyberg, Michael. 2022. *2021 Total System Electric Generation*. (Sacramento, CA: California Energy Commission). <https://www.energy.ca.gov/data-reports/energy-almanac/california-electricity-data/2021-total-system-electric-generation>.
- O’Kane, Sean. 2021. "Apple will use Tesla’s ‘megapack’ batteries at its California solar farm." *The Verge*. Accessed 2022-01-24. <https://www.theverge.com/2021/3/31/22360839/apple-tesla-megapack-energy-storage-grid-solar-batteries>.
- Obringer, Renee, Benjamin Rachunok, Debora Maia-Silva, Maryam Arbabzadeh, Roshanak Nateghi, and Kaveh Madani. 2021. "The overlooked environmental footprint of increasing Internet use." *Resources, Conservation and Recycling* 167: 105389.
- Oltmanns, Johannes, David Sauerwein, Frank Dammel, Peter Stephan, and Christoph Kuhn. 2020. "Potential for waste heat utilization of hot-water-cooled data centers: A case study." *Energy Science & Engineering* 8 (5): 1793-1810.
- Onder, E, and Nihal Sarier. 2015. "Thermal regulation finishes for textiles." *Functional Finishes for Textiles: Improving Comfort, Performance and Protection*: 17-98.
- Ortar, Nathalie, A. R. E. Taylor, Julia Velkova, Patrick Brodie, Alix Johnson, Clément Marquet, Andrea Pollio, and Liza Cirolia. 2023. "Powering ‘smart’ futures: data centres and the energy politics of digitalisation." In *Energy Futures: Anthropocene Challenges, Emerging Technologies and Everyday Life*, edited by Simone Abram, Karen Waltorp, Nathalie Ortar and Sarah Pink, 125-168. Berlin, Boston: De Gruyter.
- Ovalle-Rivera, Oriana, Marcel Van Oijen, Peter Läderach, Olivier Roupsard, Elias de Melo Virginio Filho, Mirna Barrios, and Bruno Rapidel. 2020. "Assessing the accuracy and robustness of a process-based model for coffee agroforestry systems in Central America." *Agroforestry Systems* 94. <https://doi.org/https://doi.org/10.1007/s10457-020-00521-6>.
- Oxford University Press. n.d. Culture. In *Oxford Reference*. Oxford: Oxford University Press.

- Pandey, Sudhanshu, Ritesh Gautam, Sander Houweling, Hugo Denier Van Der Gon, Pankaj Sadavarte, Tobias Borsdorff, Otto Hasekamp, Jochen Landgraf, Paul Tol, and Tim Van Kempen. 2019. "Satellite observations reveal extreme methane leakage from a natural gas well blowout." *Proceedings of the National Academy of Sciences* 116 (52): 26376-26381.
- Parmar, Aditya, Oliver Hensel, and Barbara Sturm. 2016. "Post-harvest handling practices and associated food losses and limitations in the sweetpotato value chain of southern Ethiopia." *NJAS - Wageningen Journal of Life Sciences* 80: 65-74. <https://doi.org/http://dx.doi.org/10.1016/j.njas.2016.12.002>.
- Peeters, Marjan, and Thomas Schomerus. 2014. *Renewable energy law in the EU. New horizons in environmental and energy law*. Cheltenham: Elgar.
- Periola, AA, OA Osanaiye, and AT Olusesi. 2021. "Future cloud: spherical processors for realizing low-cost upgrade in underwater data centers." *The Journal of Supercomputing*: 1-27.
- Pfeiffer, Jessica. 2020. "Datascapes: Envisioning a new kind of data center." School of Architecture and Interior Design, University of Cincinnati.
- Phillips, Allan L. 1963. "A solar-energy method for reducing coffee-drying costs." *The Journal of Agriculture of the University of Puerto Rico* 47 (4): 226-235.
- Pilkey, Orrin H, and Linda Pilkey-Jarvis. 2007. *Useless Arithmetic: Why Environmental Scientists Can't Predict the Future*. New York, NY: Columbia University Press.
- Pinto, Helen, and Ian D Gates. 2022. "Why is it so difficult to replace diesel in Nunavut, Canada?" *Renewable and Sustainable Energy Reviews* 157. <https://doi.org/https://doi.org/10.1016/j.rser.2021.112030>.
- Porter, Michael E. 1990. *The competitive advantage of nations*. London: Macmillan Press.
- Ramamoorthy, Sripriya, and Shankar Krishnan. 2018. "Towards thermal-acoustic co-design of noise-reducing heat sinks." *IEEE transactions on components, packaging and manufacturing technology* 8 (8): 1411-1419.
- Randall, Dave, Richard Harper, and Mark Rouncefield. 2007. *Fieldwork for design: Theory and practice*. London: Springer.
- Rao, Abhishek, Rohit Pillai, Monto Mani, and Praveen Ramamurthy. 2014. "Influence of dust deposition on photovoltaic panel performance." *Energy Procedia* 54: 690-700.
- Rautela, Rahul, Shashi Arya, Shilpa Vishwakarma, Jechan Lee, Ki-Hyun Kim, and Sunil Kumar. 2021. "E-waste management and its effects on the environment and human health." *Science of the Total Environment* 773: 145623.
- Raworth, Kate. 2012. "A Safe and Just Space for Humanity." *Oxfam Discussion Papers*.
- Read, Carveth. 1914. *Logic. Deductive and inductive*. 4th ed. ed. London: Simpkin, Marshall, Hamilton, Kent & Co. Ltd.
- Reddy, V Dinesh, Brian Setz, G Subrahmanya VRK Rao, GR Gangadharan, and Marco Aiello. 2017. "Metrics for sustainable data centers." *IEEE Transactions on Sustainable Computing* 2 (3): 290-303.
- Reynolds, Matt. 2022. "Gravity Could Solve Clean Energy's One Major Drawback." *Wired*.
- Ritchie, Hannah. 2020. "What are the safest and cleanest sources of energy?". Our World in Data. Accessed 2022-01-22. <https://ourworldindata.org/safest-sources-of-energy>.



## References

- Roach, John. 2020. "Microsoft finds underwater datacenters are reliable, practical and use energy sustainably." Microsoft. Accessed 2022-04-04. <https://news.microsoft.com/innovation-stories/project-natick-underwater-datacenter/>.
- Rockström, Johan, Mattias Klum, and Hubert Mania. 2016. *Big World Small Planet*. Ullstein.
- Rockström, Johan, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin III, Eric Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, and Hans Joachim Schellnhuber. 2009. "Planetary boundaries: exploring the safe operating space for humanity." *Ecology and society* 14 (2).
- Rosen (sponsor), Jacky. 2023. Federal Data Center Enhancement Act of 2023. edited by US Senate. Washington, DC.
- Rosling, Hans, Ola Rosling, and Anna Rosling-Rönnlund. 2018. *Factfulness: Ten Reasons We're Wrong About The World. And Why Things Are Better Than You Think*. New York: Sceptre.
- Ruch, Patrick, Jens Ammann, Stephan Paredes, Nicolay Wiik, Emanuel Lörtscher, Gerhard I Meijer, and Bruno Michel. 2017. "Sustainable data centers and energy conversion technologies." *12th IEA Heat Pump Conference*.
- Sabatini, Francesca. 2019. "Culture as fourth pillar of sustainable development: Perspectives for integration, paradigms of action." *European Journal of Sustainable Development* 8 (3). <https://doi.org/https://doi.org/10.14207/ejsd.2019.v8n3p31>.
- Sachs, Jeffrey D, Kaitlin Y Cordes, James Rising, Perrine Toledano, and Nicolas Wolfram Maennling. 2019. "Ensuring Economic Viability & Sustainability of Coffee Production."
- Salviati, Sergio, Carosio Federico, Guido Saracco, and Alberto Fina. 2019. "Hydrated Salt/Graphite/Polyelectrolyte Organic-Inorganic Hybrids for Efficient Thermochemical Storage." *Nanomaterials* 9 (3). <https://doi.org/https://doi.org/10.3390/nano9030420>.
- Sattiraju, Nikitha. 2020. "The Secret Cost of Google's Data Centers: Billions of Gallons of Water to Cool Servers." *Time Magazine*.
- Seele, Peter, and Irina Lock. 2017. "The game-changing potential of digitalization for sustainability: possibilities, perils, and pathways." *Sustainability Science* 12 (2): 183-185.
- Sego, Landon H, Andrés Márquez, Andrew Rawson, Tahir Cader, Kevin Fox, William I Gustafson, and Christopher J Mundy. 2012. "Implementing the Data Center Energy Productivity Metric." *ACM Journal on Emerging Technologies in Computing Systems* 8 (4). <https://doi.org/https://doi.org/10.1145/2367736.2367741>.
- Shah, Jimil M, Richard Eiland, Ashwin Siddarth, and Dereje Agonafer. 2016. "Effects of mineral oil immersion cooling on IT equipment reliability and reliability enhancements to data center operations." 2016 15th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm).
- Shim, Yeon-soo, and Donald C Bellomy. 2018. "Thinking and acting systematically about the Anthropocene." *Systemic Practice and Action Research* 31 (6): 599-615.
- Shorrocks, Anthony, James Davies, and Rodrigo Lluberias. 2021. *Global Wealth Databook 2021*. (Zürich: Credit Suisse Research Institute).
- Shuja, Junaid, Abdullah Gani, Shahaboddin Shamshirband, Raja Wasim Ahmad, and Kashif Bilal. 2016. "Sustainable Cloud Data Centers: A survey of enabling techniques and technologies." *Renewable and Sustainable Energy Reviews* 62. <https://doi.org/http://dx.doi.org/10.1016/j.rser.2016.04.034>.

- Simons, RE. 1994. "Microelectronics cooling and SEMI-THERM: A look back." Proceedings of 1994 IEEE/CHMT 10th Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM).
- Sioen, Giles Bruno, Toru Terada, and Makoto Yokohari. 2016. "Sustainability science as the next step in urban planning and design." In *Sustainability Science: Field Methods and Exercises*, 117-135. Springer.
- Skaburskis, Andrejs. 2008. "The origin of "wicked problems"." *Planning Theory & Practice* 9 (2): 277-280.
- Skatteverket. 2015. "Betänkandet Energiskatt på el – En översyn av det nuvarande systemet (SOU 2015:87)." Skatteverket. Accessed 2021-11-19. <https://www.skatteverket.se/omoss/varverksamhet/rapporterremissvarochskrivelser/remissvar/2015/remissvar2015/13157527715112.5.3810a01c150939e893f14c5c.html?q=datacenter>.
- Skoglund, Nils, Kajsa Werner, Nylund Göran M, Henrik Pavia, Albers Eva, and Markus Broström. 2017. "Combustion of seaweed - A fuel design strategy." *Fuel Processing Technology* 165.
- Smil, Vaclav. 2017. *Energy and civilization: a history*. MIT Press.
- Smoucha, E A, K Fitzpatrick, S Buckingham, and O G G Knox. 2016. "Life cycle analysis of the embodied carbon emissions from 14 wind turbines with rated powers between 50kW and 3.4MW." *Journal of Fundamentals of Renewable Energy and Applications*.
- Soto, Rodrigo, and Julio Vergara. 2014. "Thermal power plant efficiency enhancement with Ocean Thermal Energy Conversion." *Applied Thermal Engineering* 62 (1): 105-112.
- Sovacool, Benjamin K, Paul Upham, and Chukwuka G Monyei. 2022. "The "whole systems" energy sustainability of digitalization: Humanizing the community risks and benefits of Nordic datacenter development." *Energy Research & Social Science*. <https://doi.org/https://doi.org/10.1016/j.erss.2022.102493>.
- Sperling, Ed. 2009. "Data Centers In The Desert." *Forbes*.
- Star, Susan Leigh. 1999. "The Ethnography of Infrastructure." *American Behavioral Scientist* 43, no. 3: 377-391.
- Star, Susan Leigh, and Karen Ruhleder. 1994. "Steps Towards an Ecology of Infrastructure: Complex Problems in Design and Access for Large-Scale Collaborative Systems." CSCW '94, Chapel Hill, NC.
- Statens energimyndighet. 2004. *Styrning av el till prioriterade användare vid bristsituationer*. Statens energimyndighet (Eskilstuna).
- Stockholm Data Parks. 2019. *350 million people within 30 ms - is Sweden the gateway to a sustainable future?* (Stockholm: Stockholm Data Parks).
- Suomen standardisoimisliitto SFS RY. 2021. "Harnessing waste heat - a new standard improves energy efficiency of data centres." Suomen standardisoimisliitto SFS RY. Accessed 2022-08-29. <https://sfs.fi/en/harnessing-waste-heat-a-new-standard-improves-energy-efficiency-of-data-centres/>.
- Swedish Data Center Industry Association. 2023. "Swedish Data Center News." Swedish Data Center Industry Association. Accessed 30 May 2023. <https://www.sdia.se/events/>.
- Swinhoe, Dan. 2022. "Singapore enters pilot phase to restart data center development; will accept some new applications." Data Centre Dynamics. Accessed 2022-03-07. <https://www.datacenterdynamics.com/en/news/singapore-enters-pilot-phase-to-restart-data-center-development-will-accept-some-new-applications/>.

## References

- Taherzadeh, Mohammad J, and Tobias Richards. 2019. *Resource recovery to approach zero municipal waste*. CRC Press.
- Tapper, Bonno. 1925. "Dilthey's Methodology of the Geisteswissenschaften." *The Philosophical Review* 34 (4): 333-349.
- Taylor, A R E. 2017. "The Technoaesthetics of Data Centre 'White Space'." *Imaginations Journal of Cross-Cultural Image Studies* 8 (2). <https://doi.org/https://doi.org/10.17742/IMAGE.LD.8.2.5>.
- Terenius Dessne, Petter. 2015. "An introduction to OTEC technology." In *OTEC Matters 2015*, edited by Petter Terenius Dessne and Lars G Golmen, In Vetenskap för profession, 12-32. Borås: Högskolan i Borås.
- Terenius, Petter, Peter Garraghan, and Richard Harper. 2020. "Using data centre waste heat to dry coffee whilst supplying small-scale farmers with ICT." International Conference for Sustainable Development, New York.
- . 2021. "Using data centre waste heat to dry coffee whilst supplying small-scale farmers with ICT: A novel idea and a case study based on a systems approach." *Journal of Strategic Innovation and Sustainability* 16 (2).
- . 2022. "Novel Strategies for Data Centre Waste Heat Use." Towards a Cleaner Earth: 18th International Conference on Clean Energy, Kuching.
- . 2023. "A material social view on data center waste heat: Novel uses and metrics." *Frontiers in Sustainability* 3. <https://doi.org/https://doi.org/10.3389/frsus.2022.1008583>.
- Terenius, Petter, Lars G Golmen, Peter Garraghan, and Richard Harper. 2021. "Heat energy from datacenters: an opportunity for marine energy." International Conference on Ocean Energy 2021, Washington, DC.
- Tesla. 2021. "Megapack." Accessed 2022-01-24. [https://www.tesla.com/en\\_gb/megapack](https://www.tesla.com/en_gb/megapack).
- The Green Grid. 2007. "The Green Grid data center power efficiency metrics: PUE and DCiE." *Green Grid report*.
- . 2008. *A framework for data center energy productivity*. (The Green Grid).
- . 2010a. *Carbon Usage Effectiveness (CUE): A Green Grid data center sustainability metric*. (The Green Grid).
- . 2010b. *ERE: A metric for measuring the benefit of reuse energy from a data center*. (The Green Grid).
- The World Bank. 2016. *Digital Dividends: World Development Report 2016*. (Washington, DC).
- . 2021. *Data for Better Lives: World Development Report 2021*. (Washington, DC: The World Bank).
- Thorén, Henrik, and Johannes Persson. 2013. "The philosophy of interdisciplinarity: sustainability science and problem-feeding." *Journal for General Philosophy of Science* 44 (2): 337-355.
- Tozer, Robert, Sophia Flucker, Beth Whitehead, Deborah Andrews, and Jon Summers. 2018. "Data Center Sustainability Index (CH-18-C055)." ASHRAE 2018 Winter Conference, Chicago, IL.
- Transition ApS. 2021. *Guide for industrial symbiosis facilitators*. (Kalundborg Symbiosis).
- Tudoroiu, Theodor. 2014. "Social Media and Revolutionary Waves: The Case of the Arab Spring." *New Political Science* (3): 346-365. <https://doi.org/https://doi.org/10.1080/07393148.2014.913841>.

- U.S. Energy Information Administration. 2020. *Electric Power Monthly with Data for November 2019*. (Washington, DC: U.S. Energy Information Administration).
- . 2021. "What is energy." U.S. Energy Information Administration. Accessed 2022-01-19. <https://www.eia.gov/energyexplained/what-is-energy/>.
- . 2022. "Petroleum and other liquids." U.S. Energy Information Administration. Accessed 2023-01-18. <https://www.eia.gov/international/data/world/petroleum-and-other-liquids/annual-petroleum-and-other-liquids-production>.
- Ueoro, Takeshi. 2011. "Introduction of Datacenter Performance Per Energy (DPPE)."
- United Nations - Department of Economic and Social Affairs - Population Division. 2022. *World Population Prospects 2022*. New York, NY: United Nations.
- United Nations Development Programme. 2022. *The 2021/2022 Human Development Report*. (New York: United Nations).
- Van Der Leeuw, Sander. 2020. *Social sustainability, past and future: undoing unintended consequences for the Earth's survival*. Cambridge University Press.
- Vega, Luis A. 2013. "Ocean thermal energy conversion." *Renewable Energy Systems*: 1273-1305.
- Velkova, Julia. 2016. "Data that warms: Waste heat, infrastructural convergence and the computation traffic commodity." *Big Data & Society* 3 (2): 2053951716684144.
- . 2021. "Thermopolitics of data: cloud infrastructures and energy futures." *Cultural Studies*. <https://doi.org/https://doi.org/10.1080/09502386.2021.1895243>.
- Vesterlund, Mattias, Stanislava Borisová, Erik Lundmark, Virpi Leinonen, Henna Tiensuu, Hampus Markeby-Ljungqvist, Gábor Takács, and Jaakko Suutala. 2021. *Data center for biomass drying*. (Luleå: Oulu University RISE Research Institutes of Sweden AB, SFTec, Luleå University of Technology and H1 Systems).
- Vickers, Geoffrey. 1983. *Human Systems are Different*. London: Harper & Row.
- vom Dahl, Antonia. 2023. "Das neue Energieeffizienzgesetz – "Energy Efficiency First" nun auch in Deutschland?". CMS Deutschland. Accessed 2023-06-05. <https://www.cmshs-bloggt.de/rechtsthemen/sustainability/sustainability-environment-and-climate-change/das-neue-energieeffizienzgesetz-energy-efficiency-first-nun-auch-in-deutschland/>.
- Vonderau, Asta. 2019. "Scaling the cloud: Making state and infrastructure in Sweden." *Ethnos* 84 (4): 698-718.
- . 2021. "Att få molnet att växa: En studie av den digitala framtidens megaprojekt i Sverige." In *Megaprojekt - Kritiska perspektiv på storskalig infrastruktur*, edited by Gabriella Körling and Susann Baez Ullberg, 161. Svenska sällskapet för antropologi och geografi.
- Värri, Konsta, and Petra Seppälä. 2019. *Small modular reactors. Report 2019:625*. (Stockholm: Energiforsk).
- Waas, Tom, Jean Hugé, Thomas Block, Tarah Wright, Francisco Benitez-Capistros, and Aviel Verbruggen. 2014. "Sustainability assessment and indicators: Tools in a decision-making strategy for sustainable development." *Sustainability* 6 (9): 5512-5534.
- Wahlroos, Mikko, Matti Pärssinen, Jukka Manner, and Sanna Syri. 2017. "Utilizing data center waste heat in district heating—Impacts on energy efficiency and prospects for low-temperature district heating networks." *Energy* 140: 1228-1238.

## References

- Wahlroos, Mikko, Matti Pärssinen, Samuli Rinne, Sanna Syri, and Jukka Manner. 2018. "Future views on waste heat utilization: Case of data centers in Northern Europe." *Renewable and Sustainable Energy Reviews* 82: 1749-1764.
- Wan, Jianxiong, Jie Zhou, and Xiang Gui. 2021. "Sustainability Analysis of Green Data Centers with CCHP and Waste Heat Reuse Systems." *IEEE Transactions on Sustainable Computing* 6, no. 1.
- Warf, Barney. 2021. "Internet censorship: Shaping the world's access to cyberspace." In *Routledge Handbook of Media Geographies*, edited by Paul C Adams and Barney Warf. Milton Park, UK: Routledge.
- Watson, Michael D. 2019. "Systems Engineering Principles and Hypotheses." *Insight* 22 (1). <https://doi.org/https://doi.org/10.1002/inst.12233>.
- Weller, Zachary D, Steven P Hamburg, and Joseph C von Fischer. 2020. "A national estimate of methane leakage from pipeline mains in natural gas local distribution systems." *Environmental science & technology* 54 (14): 8958-8967.
- Weston, Sabina. 2021. "Microsoft is submerging servers in boiling liquid to prevent Teams outages." *ITPro*.
- Whitehead, Beth, Deborah Andrews, Amip Shah, and Graeme Maidment. 2014. "Assessing the environmental impact of data centres part 1: Background, energy use and metrics." *Building and Environment* 82: 151-159.
- Wiener, Norbert. 1948. *Cybernetics or Control and Communication in the Animal and the Machine*. New York: John Wiley & Sons.
- Wit, Cary W de. 2003. "Field methods for investigating sense of place." *The North American Geographer* 5 (1-2). <https://doi.org/https://www.researchgate.net/publication/306151929>.
- World Community Grid. 2021. "About us." World Community Grid. Accessed 2022-04-04. <https://www.worldcommunitygrid.org/about/about>.
- World Economic Forum. 2020. *The Global Competitiveness Report Special Edition 2020: How Countries are Performing on the Road to Recovery*. (Geneva: World Economic Forum).
- WWF Malaysia and Boston Consulting Group. 2021. *Securing our Future: Net Zero Pathways for Malaysia*. (Klang: WWF Malaysia).
- Yeh-Yun Lin, Carol, and Leif Edvinsson. 2011. *National intellectual capital. A comparison of 40 countries*. New York: Springer.
- Yin, Robert K. 2012. *Applications of case study research*. 3rd ed. Thousand Oaks, CA: SAGE.
- Zabeu, Sheila. 2023. "Microsoft wants to replicate project to share datacenter batteries." Network King. Accessed 2024-01-05. <https://network-king.net/microsoft-wants-to-replicate-project-to-share-datacenter-batteries/>.
- Zakarya, Muhammad. 2018. "Energy, performance and cost efficient datacenters: A survey." *Renewable and Sustainable Energy Reviews* 94: 363-385.
- Zeuner, Brett. 2018. "An obsolescing bargain in a rentier state: multinationals, artisanal miners, and cobalt in the Democratic Republic of Congo." *Frontiers in Energy Research* 6: 123.

# Appendices

|  |     |
|--|-----|
| Appendix A: Engagement.....                            | 273 |
| Appendix B: Research visits and fieldwork .....        | 275 |
| Appendix C: Interview questionnaire for Dr Azzman..... | 276 |

## Appendix A: Engagement

- Part of the Leverhulme Centre for Material Social Futures (MSF), Lancaster University, 2019-2022
- Post-graduate representative in the Faculty of Science and Technology's Sustainability Advisory Committee, Lancaster University, 2020-2022
- Presenter at the International Conference of Sustainable Development annual conference, Columbia University (online), September 2020
- Poster presenter at the Faculty of Science and Technology Annual Conference, Lancaster University (online), March 2021
- Presenter at the International Conference of Ocean Energy (online), April 2021
- Presenter at the Material Science Institute PhD Conference, Lancaster University, October 2021
- Presenter at the Connecting on Climate Change summit, Lancaster University, October 2021
- Visiting Researcher, Universiti Teknologi Malaysia, Kuala Lumpur, May 2022
- Presentation for school children, Batu Caves Elementary School, Kuala Lumpur, May 2022
- Presenter at the International Conference on Clean Energy, Sawang (online), July 2022
- Presenter at the Postgraduate Research Conference, Lancaster University, November 2022

- Presenter at the Open Compute Project European Summit, Prague, April 2023
- Participant in the Swiss Next Generations working group, May 2023 onwards
- Participant in the OCP data centre energy efficiency metrics and heat reuse polices working groups, June 2023 onwards



## **Appendix B: Research visits and fieldwork**

This thesis has broad connections to infrastructure, societies and academic disciplines. Therefore, it has been necessary to engage in fieldwork and other sorts of travel. When reasonable, travel has been conducted by train. When possible, flights have been climate compensated, and a few have partly been made with biofuel.

- 2-7 April 2020 – research assistant in a forestry research project; forests of the counties of Småland and Västergötland, Sweden
- 15 September 2020 – visit to Lindwalls coffee roastery; Uppsala, Sweden
- 26-28 October 2020 – visit to several data centres and to the RISE SICS North AB research institute; Luleå and Boden, Sweden
- 6-8 July 2021 – seaweed dehydration experiments; lake Sommen, Sweden
- 23 August 2021 – visit to EcoDataCenter; Falun, Sweden
- 11 November 2021 – visit to COP26; Glasgow, UK
- 30 April-2 June 2022 – visiting researcher at Universiti Teknologi Malaysia
- 17 August 2022 – visit to Falu Energi & Vatten AB; Falun, Sweden
- 3 March 2023 – visit to a liquid immersion facility courtesy of DeepSquare; Basel, Switzerland

## Appendix C: Interview questionnaire for Dr Azzman

### Interview questions: Data centres and society

*If appropriate, the following questions may be used during the interview. You are encouraged to ponder these questions prior to the interview.*

*The interview may be conducted in a “semi-structured” manner, that is, the questions may be touched upon, but not necessarily in the written order or in this exact form, and additional questions may be brought up for discussion. In this way, the interview differs from – and may provide more value than – a static questionnaire.*

#### Question 1.

Can you please tell me about your background experience of relevance to this interview (a little of where you’ve been and what you’ve done, and an approximate timeframe)?

#### Question 2.

What, if any, is your experience from working with matters relating to data centres?

#### Question 3.

What, if any, is your experience from working with matters relating to energy?

#### Question 4.

What, if any, is your experience from reliance on performance metrics (such as energy efficiency)? If you do have such experience, what are your thoughts on this matter?

#### Question 5.

Do you have any experiences, thoughts, insights or visions regarding holistic approaches to energy within a societal structure (renewables, energy reuse, etc)?

#### Question 6.

Do you have other thoughts or ideas relating to opportunities for data centres, such as data centres’ geographical placement, water usage, energy source, energy use in the data centre, waste heat reuse or local benefits from ICT (information and communication technology) opportunities that a local data centre may bring with it?