Chapter 1

The Iterated Prisoner's Dilemma: 20 Years On

Siang Yew Chong[1], Jan Humble[2], Graham Kendall[2], Jiawei Li[2,3], Xin Yao[1]

University of Birmingham[1], University of Nottingham[2], Harbin Institute of Technology[3]

1.1. Introduction

In 1984, Robert Axelrod reported the results of two iterated prisoner's dilemma (IPD) competitions [Axelrod (1984)]. The booked was to be a catalyst for much of the research in this area since that time. It is unlikely that you would write a scientific paper about IPD, without citing Axelrod's 1984 book. The book is even more remarkable in that it is just as accessible to a general audience, as well as being an important source of inspiration for the scientific community.

In 2001, whilst attending the Congress on Evolutionary Computation (CEC) conference, we were discussing some of the presentations we had seen which reported recent some of the latest work on the iterated prisoner's dilemma. We were paying tribute to the fact that Axelrod's book had stood the test of time when somebody made a casual comment suggesting that we should re-run the competition in 2004, to celebrate the 20th anniversary. And, so, this book was born.

Of course, since the conversation in Hawaii and the publication of this book, there have been a lot of people doing a lot of work. Not least of all Robert Axelrod who was good enough to give up his time to present a plenary talk at the CEC conference in 2004. At that talk he presented his latest work which is investigating evolution on a grid based world.

We owe a debt of thanks to the UK's EPSRC (Engineering and Physical Sciences Research Council). This is the largest of the UK research councils which funds research in the UK. When we returned from Hawaii, we submitted a proposal,[a] which requested a small amount of funds (£23,718) in order to re-run, and extend, the competitions that Axelrod had run 20 years earlier. The funds we received from EPSRC allowed us to run two competitions, one in 2004 and one in 2005. The entrants to the competitions were invited to submit a chapter for consideration in this book. These chapters underwent a peer review process (see later in this chapter for an acknowledgement of the reviewers) and those chapters that were successful form the latter part of this book.

As editors, we feel fortunate to have several winners, second and third place entries reported in this book. This affords the reader the opportunity to learn, first hand from the authors, what made these strategies so successful and, perhaps, use some of the ideas and innovations in their own strategies for future competitions.

1.2. Iterated Prisoner's Dilemma

Almost every chapter in this book has its own description of the iterated prisoner's dilemma. As each chapter can be read in isolation and, for completeness, we present our own interpretation of the IPD here, along with a short review of some of the important work in the area.

The prisoner's dilemma (PD) and iterated prisoners dilemma (IPD) have been a rich source of research material since the 1950's. However, the publication of Axelrod's book [Axelrod (1984)] in the 1980's was largely responsible for bringing this research to the attention to other areas, outside of game theory, including evolutionary computing, evolutionary biology, networked computer systems and promoting cooperation between opposing countries [Goldstein (1991); Fogel (1993); Axelrod and D'Ambrosio (1995)]. Despite the large literature base that now exists (see, for example, [Poundstone (1992); Boyd and Lorberbaum (1987); Maynard Smith (1982); Davis (1997), Axelrod (1997)], this is an on-going area of research, with Darwen and Yao [Darwen and Yao (1995, 2001); Yao and Darwen (1999)] carrying out some recent work. Their 2001 work [Darwen and Yao (2001)] extends the prisoner's dilemma by offering more choices, other than simply "cooperate" or "defect," and by providing indirect interactions (reputation).

When you play the prisoner's dilemma you have to decide whether to cooperate with an opponent, or defect. Both you and your opponent make a

[a]The EPSRC grant reference numbers are GR/S63465/01 and GR/S63472/01.

choice and then your decisions are revealed. You receive a payoff according to the following matrix (where the top line is the payoff to the column).

	Cooperate	Defect
Cooperate	$R = 3$ $R = 3$	$T = 5$ $S = 0$
Defect	$S = 0$ $T = 5$	$P = 1$ $P = 1$

- R is a ***R**eward* for mutual cooperation. Therefore, if both players cooperate then both receive a reward of 3 points.
- If one player defects and the other cooperates then one player receives the ***T**emptation* to defect payoff (5 in this case) and the other player (the cooperator) receives the ***S**ucker* payoff (zero in this case).
- If both players defect then they both receive the ***P**unishment* for mutual defection payoff (1 in this case).

The question arises: what should you do in such a game?

- Suppose you think the other player will cooperate. If you cooperate then you will receive a payoff of 3 for mutual cooperation. If you defect then you will receive a payoff of 5 for the Temptation to Defect payoff. Therefore, if you think the other player will cooperate then you should defect, to give you a payoff of 5.
- But what if you think the other player will defect? If you cooperate, then you get the Sucker payoff of zero. If you defect then you would both receive the Punishment for Mutual Defection of 1 point. Therefore, if you think the other player will defect, you should defect as well.

So, you should defect, no matter what option your opponent chooses.

Of course, the same logic holds for your opponent. And, if you both defect you receive a payoff of 1 each, whereas, the better outcome would have been mutual cooperation with a payoff of 3. The payoff for an individual is less than that could have been achieved by two cooperating players, thus the dilemma and the research challenge of finding strategies that promote mutual cooperation.

In defining a prisoner's dilemma, certain conditions have to hold. The values we used above, to demonstrate the game, are not the only values that could have been used, but they do have to adhere to the conditions listed below.

Firstly, the order of the payoffs is important. The best a player can do is T (temptation to defect). The worst a player can do is to get the sucker payoff, S. If the two players cooperate then the reward for that mutual cooperation, R, should be better than the punishment for mutual defection, P. Therefore, the following must hold.

$$T > R > P > S\,. \tag{1.1}$$

Secondly, players should not be allowed to get out of the dilemma by taking it in turns to exploit each other. Or, to be a little more precise, the players should not play the game so that they end up with half the time being exploited and the other half of the time exploiting their opponent. In other words, an even chance of being exploited or doing the exploiting is not as good an outcome as both players mutually cooperating. Therefore, the reward for mutual cooperation should be greater than the average of the payoff for the temptation and the sucker. That is, the following must hold.

$$R > (S + T)/2\,. \tag{1.2}$$

Playing a "one-shot" prisoners dilemma, it is not difficult to decide which strategy to adopt, but the question arises: can cooperation evolve from playing the game over and over again, against the same opponent?

If you know how many times you are to play, then there is an argument that the game is exactly the same as playing the "one-shot" prisoners dilemma. This is based on the observation that you will defect on the last iteration as that is the sensible thing to do as, you are in effect playing a single iteration. Knowing this, it is sensible to defect on the second to last one as well; and this logic can be applied all the way to the first iteration.

However, this reasoning cannot be used when the number of iterations is infinite as you know there is always another iteration. In practise, this translates to not knowing when the game will end.

Experiments, using human players [Scodel (1962, 1963); Minas et al. (1960); Scodel and Philburn (1959), Scodel et al. (1959); Scodel et al. (1960)] showed that they, generally, did not cooperate even when it should have been obvious that the other person was going to cooperate, just as long as you do. It has been a long term aim to find strategies which causes players to cooperate. If players would only cooperate then their payoff, over an indefinite number of games could be maximised, rather than tending towards defection and hoping the other player would cooperate. In 1979 Axelrod organised a prisoner's dilemma competition and invited game theorists to submit their strategies [Axelrod (1980a)]. Fourteen entries were received with an extra one being added (defect or cooperate with equal probability). The strategies were competed against each other, including itself. The winner was Anatol Rapoport who submitted the simple strategy (Tit-for-Tat) which cooperates on the first move, then does whatever your opponent did on the previous move. In a second tournament [Axelrod (1980b)], 62 entries were received but, again, the winner was Tit-for-Tat. These two competitions formed the basis of his important book [Axelrod (1984)].

The prisoners dilemma has a modern day version in the form of the TV show "Shafted" - a game show recently screened on terrestrial TV in the UK (note that this show is not a true prisoners dilemma as defined by Rapoport [Rapoport (1996)], but does demonstrate that the ideas have wider applicability). At the end of the show two contestants have accumulated a sum of money and they have to decide if to share the money or to try and get all the money for themselves. Their decision is made without the knowledge of what the other person has decided to do. If both contestants cooperate then they share the money. If they both defect then they both receive nothing. If one cooperates and the other defects, the one that defected gets all the money and the contestant that cooperated gets nothing.

Although the prisoners dilemma, in the context of game theory, has been an active research area for at least 50 [Scodel (1962); Scodel (1963); Minas et al. (1960); Scodel and Philburn (1959); Scodel et al. (1959); Scodel et al. (1960)] years (it can be traced back to von Neumann and Morgenstern [von Neumann and Morgenstern (1944)] and, of course, John Nash [Nash (1950, 1953)]), it is still an active research area with, among other research aims, researchers trying to evolve strategies [O'Riordan (2000)] that promote cooperation.

Recent research has also considered the prisoner's dilemma where there are more than two choices and more than two players. Darwen and Yao have shown that offering more choices leads to less cooperation [Darwen and Yao (2001)], although reputation may help [Darwen and Yao (2002); Yao and Darwen (1999)]. Birk [Birk (1999)] used a multi-payer IPD. His model had continuous degrees of cooperation (as opposed to the binary; cooperate or defect). He used a robotic environment and showed that a *justified-snobism* strategy, that tries to cooperate slightly more than the average, is a successful strategy and is evolutionarily stable (that is, it cannot be invaded by another strategy). O'Riordan and Bradish (2000) also simulated a multi-player game where the players are involved in many types of games.

They show that cooperation can emerge in a high percentage of 2-player games.

As well as the academic papers on the subject, there are many books devoted to game theory and/or the prisoners dilemma. The 1997 book by Axelrod (1997) re-produces a range of his papers (with commentary) ranging from 1986 through to 1997. The papers consider areas such as promoting cooperation using a genetic algorithm, coping with noise and promoting norms.

1.3. Contents of the Book

This book does not have to be read from cover to cover. Each chapter can be read independently, with most of the chapters describing the IPD. This was a conscious decision by the editors as we realised that the book would be dipped into and we did not want to make any chapter dependent on any other. Also, each chapter has its own set of references, rather than having one complete list of references at the end of the book. The book is structured as follows

Chapter 1

This chapter provides a general introduction to the book. In keeping with the rest of the book, we also briefly describe the IPD. As well as briefly describing each chapter. This chapter also presents the results of the two competitions that we ran in 2004 and 2005.

Chapter 2

Chapter 2 ("*Iterated Prisoner's Dilemma and Evolutionary Game Theory*") reviews some of the important work in IPD, with particular emphasis (in the latter part of the chapter) on evolutionary game theory. The chapter contains over 250 references, which we hope will be a good starting point for other researchers who are looking to start work in this area.

We have concentrated on the evolutionary aspects of IPD for two reasons. Firstly, this seemed to be an area that was exploited in the entries we received. Secondly, the literature on IPD is truly vast (perhaps only exceeded by literature on the traveling salesman problem), and we had to draw some boundaries and, given the close links that this competition had with the Congress on Evolutionary Computation, it seemed appropriate to report on the evolutionary aspects of IPD.

We apologise to any authors who feel their work should have been included in this chapter. We hope you understand that we simply could not list every paper. However, if you would like to drop us an EMAIL, we would be happy to consider the inclusion of the reference in any later editions.

Chapter 3

Chapter 3 ("*Learning IPD Strategies Through Co-evolution*") reviews another area of IPD that has received scientific interest in recent years; that of co-evolution. This chapter also discusses an extension to the classic IPD formulation. That is when there are more than two players and when they have more than two choices. Similar to chapter two, there is an extensive list of references for the interested reader.

Chapter 4

This chapter reports the winning strategy from competition 4, from the event held in 2005. This competition mimics the original ones held by Axelrod. Only one entry was allowed per person, to stop the cooperating strategies that had dominated the first competition. Although we believe that having cooperating strategies is a valid tactic, some competitors felt that this did not truly mimic the original competitions. For this reason we introduced an additional competition for the 2005 event. The result was a win for Jiawei Li, who details his winning strategy in chapter 4, which is entitled *How to Design a Strategy to Win an IPD Tournament.*

Chapter 5

The strategy in this chapter attempts to model its opponent using an artificial immune system. It is interesting to see how relatively new methodologies are being used for problems such as IPD, demonstrating that there is a continuous flow of new ideas which might just be shown to be superior to all other methods so far. Whilst not appearing in the top ten of any of the competitions that it entered, it does present an exciting new research direction for IPD tournaments.

Chapter 6

Michael Filzmoser, reports on a variation of tit-for-tat, which he calls *Exponential Smoothed Tit-for-Tat*. Whereas tit-for-tat only considers the last move of the opponent, exponential smoothed tit-for-tat considers the complete history of the opponent. This discussion is extended to IPD with noise, as well as the more common IPD, where the actions by the player are reliably reported.

Chapter 7

In chapter 7 ("*Opponent modelling, Evolution, and the Iterated Prisoner's Dilemma*"), the authors explore the idea of modeling an opponent. It does this by playing tit-for-tat for the first 50 moves, whilst trying to model the moves played by the opponent. After 50 moves, subsequent moves are then based on the model that has been built.

It is interesting to compare this strategy (which came 3rd in competition 4 in 2005), with the strategy described in chapter 4, which also uses a type of modeling but over a shorter time period. Perhaps this explains why it was able to achieve better payoffs, as it was able to exploit opponents much earlier in the game?

Chapter 8

The strategies reported in this chapter were entered in both the 2004 and 2005 events, and performed well in many of the competitions, winning competition 1 in the 2005 event.

This chapter, more than any other, touches on the debate about cooperating strategies, which is why we introduced competition 4 in the 2005 event. If you followed the discussion at the time, many entrants (with some justification) questioned if allowing multiple strategies from one person was in the spirit of the original Axelrod competitions. Whilst we agreed with this, so introduced a single entry rule in 2005, we also argue that these competitions were about the research that was being carried out and some of the chapters in this book report on those results. Of course, as the authors of chapter 8 admit, there are still ways of flouting the rules by submitting cooperating entries under different names. We hope that the other entrants will accept this in the spirit of research under which this was done. As the authors point out, the organisers failed to recognise that cooperating strategies had been submitted, but, as they also say, this is a theoretically difficult problem.

We would also like to take this opportunity to the authors of chapter 8 for missing their OTFT strategy from some of the competitions. It is still unclear to us why this happened.

Chapter 9

A team from Southampton, who took the first three places in competition 1, in the 2004 competition present chapter 9. Their chapter is an excellent example of how strategies can cooperate. As strategies have no mechanism to interact directly, the only way to recognise one of your collaborators is to somehow communicate through the defect/cooperate choices that you make.

Chapter 10

One of the competitions that we run included noise, with some low probability. By noise, we mean that a defect or cooperate signal might be misinterpreted. This final chapter by Tsz-Chiu Au and Dana Nau explores this issue using a strategy they call *Derived Belief Strategy.* It attempts to model their opponent and then judge if their choice has been affected by noise. They performed very well in the competition, even when up against strategies which were cooperating.

1.4. Celebrating the 20th Anniversary: The Competitions

We ran two events. The first was held during the Congress of Evolutionary Computation Conference in 2004 (June 19-23, Portland, Oregon, USA) and the next at the Computational Intelligence and Games Conference in 2005 (April 4-6, 2005, Essex UK). At the 2004 event we ran three competitions, with an additional competition being held in 2005.

(1) The first competition aimed to emulate the original Axelrod competition. We received some enquiries about whether multiple entries were allowed. As we had not stated this as a restriction, we allowed it (but did state we had the right to limit the number, else running the competition may become intractable). At the time, we did not realise the

controversy that this decision would cause, which is why we modified the competitions in the 2005 event.

(2) The second competition had noise in it. Each decision had a 0.1 probability of being mis-interpreted.
(3) The third competition allowed competitors to submit a strategy to an IPD that has more than one player and more than one payoff, that is, multi player and multi-choice.
(4) The fourth competition (which was only run in 2005) emulated the original Axelrod competition. The definition was exactly the same as competition 1, but we only allowed one entry per person.

The payoff table we used for competitions 1, 2 and 4 is shown in table 1.1. The payoff table for competition 3 is shown in table 1.2.

Table 1.1. Payoff table for all IPD competitions except for the IPD with multiple players and multiple choices.

	Cooperate	Defect
Cooperate	$R = 3$ $R = 3$	$T = 5$ $S = 0$
Defect	$S = 0$ $T = 5$	$P = 1$ $P = 1$

Table 1.2. Payoff table for IPD competition with multiple players and multiple payoffs Player BLevels of Cooperation.

		Player B				
	Levels of Cooperation	1	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{4}$	0
Player A	1	4	3	2	1	0
	$\frac{3}{4}$	$4\frac{1}{4}$	$3\frac{1}{4}$	$2\frac{1}{4}$	$1\frac{1}{4}$	$\frac{1}{4}$
	$\frac{1}{2}$	$4\frac{1}{2}$	$3\frac{1}{2}$	$2\frac{1}{2}$	$1\frac{1}{2}$	$\frac{1}{2}$
	$\frac{1}{4}$	$4\frac{3}{4}$	$3\frac{3}{4}$	$2\frac{3}{4}$	$1\frac{3}{4}$	$\frac{3}{4}$
	0	5	4	3	2	1

To support the competitions, we developed a software framework. This is discussed in the Appendix, and a URL is supplied so that the software can be downloaded.

1.5. Competition Results

In the following tables we present the top ten entries from each of the competitions. The full listings of the results can be seen at http://www.prisoners-dilemma.com. Also available on the web site is a log containing all the interactions that took place.

Table 1.3. Results from 2004 event, competition 1. There were 223 entries (19 web based entries, 195 java based entries and 9 standard entries (RAND, NEG, ALLC, ALLD, TFT, STFT, TFTT, GRIM, Pavlov)).

Rank	Player	Strategy	Won	Drawn	Lost	Total Points
1	Gopal Ramchurn	StarSN (StarSN)	105	21	98	117,057
2	Gopal Ramchurn	StarS (StarS)	113	48	63	110,611
3	Gopal Ramchurn	StarSL (StarSL)	115	46	63	110,511
4	GRIM (GRIM Trigger)_1 Wolfgang Kienreich	GRIM (GRIM Trigger)	120	76	28	100,611
5	Wolfgang Kienreich	OTFT (Omega tit for tat)	90	70	64	100,604
6	Wolfgang Kienreich	ADEPT (ADEPT Strategy)	95	72	57	96,291
7	Emp_1	EMP (Emperor)	90	73	61	95,927
8	Bingzhong Wang	()	31	94	99	94,161
9	Hannes Payer	PRobbary (PRobbary Historylength 2)	95	75	54	94,123
10	Nanlin Jin	HCO (HCO)	27	95	102	93,953

Table 1.4. Results from 2004 event, competition 2. There were 223 entries (19 web based entries, 195 java based entries and 9 standard entries (RAND, NEG, ALLC, ALLD, TFT, STFT, TFTT, GRIM, Pavlov)).

Rank	Player	Strategy	Won	Drawn	Lost	Total Points
1	Gopal Ramchurn	StarSN (StarSN)	42	2	180	93,962
2	Colm O'Riordan	Mem1 (Mem1)	5	1	218	83,049
3	Gopal Ramchurn	CoordinateCDCSIAN (CoordinateCDCSIAN)	158	6	60	83,015
4	Gopal Ramchurn	PoorD (PoorD)	190	7	27	82,890
5	Wolfgang Kienreich	OTFT (Omega tit for tat)	158	8	58	82,838
6	Wayne Davis	ltft (ltft)	66	8	150	82,765
7	GRIM (GRIM Trigger)_1	GRIM (GRIM Trigger)	184	7	33	82,591
8	Gopal Ramchurn	MooD (MooD)	193	3	28	82,578
9	Gopal Ramchurn	AITFT (AITFT)	60	9	155	82,504
10	Gopal Ramchurn	GSTFT (GSTFT)	64	9	151	82,502

Table 1.5. Results from 2004 event, competition 3. There were 15 entries. Note that there is only one round in this competition.

Rank	Player	Strategy	Total Points
1	Gopal Ramchurn	AgentSoton (SOTON AGENT)	3,756
2	Gopal Ramchurn	HarshTFT (HarshTFT)	3,756
3	Deirdre Murrihy	PCurvepower1Memory2 (Penalty Curve of 1 using opponent's previous 2 moves)	3,738
4	Deirdre Murrihy	PCurvepower2Memory2 (Penalty Curve of 2 using opponent's previous 2 moves)	3,738
5	Deirdre Murrihy	PCurvepower0.5Memory2 (Penalty Curve of 0.5 using opponent's previous 2 moves)	3,738
6	Enda Howley	PCurvepower2 (Penalty Curve of 2 using opponent's previous move)	3,738
7	Enda Howley	PCurvepower1 (Penalty Curve of 1 using opponent's previous move)	3,738
8	Enda Howley	PCurvepower0.5 (Penalty Curve of 0.5 using opponent's previous move)	3,738
9	Wolfgang Kienreich	CNHM (CosaNostra Hitman)	3,738
10	Wolfgang Kienreich	CNHM (CosaNostra Hitman)	3,738

Table 1.6. Results from 2005 event, competition 1. There were 192 entries (41 web based entries, 142 java based entries and 9 standard entries (RAND, NEG, ALLC, ALLD, TFT, STFT, TFTT, GRIM, Pavlov)).

Rank	Player	Strategy	Won	Drawn	Lost	Total Points
1	Wolfgang Kienreich	CNGF (CosaNostra Godfather)	48	96	49	100,905
2	Jia-wei Li	IMM01 (Intelligent Machine Master 01)	46	112	35	98,922
3	Carlos G. Tardon	CLAS- (CLAS-)	23	95	75	92,174
4	Perukrishnen Vytelingum	SWIN (Soton Agent RA - Competition 1)	61	44	88	90,918
5	Constantin Ionescu	LORD (the lord strategy)	20	102	71	87,617
6	GRIM (GRIM Trigger)_1	GRIM (GRIM Trigger)	73	114	6	84,805
7	Tsz-Chiu Au	LSF (Learning of opponent strategy with forgiveness)	28	94	71	84,698
8	Tsz-Chiu Au	DBStft (DBS with TFT)	23	97	73	83,867
9	Richard Brunauer	PRobberyL2 (PRobberyL2)	14	98	81	83,837
10	Carlos G. Tardon	CLAS2 (CLAS2)	72	96	25	83,746

Table 1.7. Results from 2005 event, competition 2. There were 165 entries (26 web based entries, 130 java based entries and 9 standard entries (RAND, NEG, ALLC, ALLD, TFT, STFT, TFTT, GRIM, Pavlov)).

Rank	Player	Strategy	Won	Drawn	Lost	Total Points
1	Perukrishnen Vytelingum	BWIN (S2Agent1_ZEUS - Competition 2)	85	1	80	73,330
2	Jia-wei Li	IMM01 (Intelligent Machine Master 01)	108	7	51	70,506
3	Tsz-Chiu Au	DBSy (DBS (version y))	35	3	128	68,370
4	Tsz-Chiu Au	DBSz (DBS (version z))	27	3	136	68,339
5	Tsz-Chiu Au	DBSpl (DBS with learning prevention)	37	2	127	67,979
6	Tsz-Chiu Au	DBSd (Derivative Belief Strategy (version d))	42	6	118	67,392
7	Tsz-Chiu Au	DBSx (DBS (version x))	19	9	138	66,719
8	Tsz-Chiu Au	TFTIc (TFT improved (ver. c))	41	4	121	66,409
9	Tsz-Chiu Au	DBSf (Derivative Belief Strategy (version f))	48	2	116	66,269
10	Tsz-Chiu Au	TFTIm (TFT improved (ver. m))	38	3	125	66,239

Table 1.8. Results from 2005 event, competition 3. There were 34 entries. Note that there is only one round in this competition.

Rank	Player	Strategy	Total Points
1	Perukrishnen Vytelingum	$AgentSoton ($SOTON AGENT)	7,558
2	Deirdre Murrihy	PCurvepower1Memory2 (Penalty Curve of 1 using opponent's previous 2 moves)	7,521
3	Deirdre Murrihy	PCurvepower2Memory2 (Penalty Curve of 2 using opponent's previous 2 moves)	7,521
4	Deirdre Murrihy	PCurvepower0.5Memory2 (Penalty Curve of 0.5 using opponent's previous 2 moves)	7,521
5	Enda Howley	PCurvepower2 (Penalty Curve of 2 using opponent's previous move)	7,521
6	Enda Howley	PCurvepower1 (Penalty Curve of 1 using opponent's previous move)	7,521
7	Enda Howley	PCurvepower0.5 (Penalty Curve of 0.5 using opponent's previous move)	7,521
8	Wolfgang Kienreich	CNHM (CosaNostra Hitman)	7,521
9	Wolfgang Kienreich	CNHM (CosaNostra Hitman)	7,521
10	Wolfgang Kienreich	CNHM (CosaNostra Hitman)	7,521

Table 1.9. Results from 2005 event, competition 4. There were 50 entries (26 web based entries, 15 java based entries and 9 standard entries (RAND, NEG, ALLC, ALLD, TFT, STFT, TFTT, GRIM, Pavlov)).

Rank	Player	Strategy	Won	Drawn	Lost	Total Points
1	Jia-wei Li	APavlov (Adaptive Pavlov)	11	34	6	30,096
2	Wolfgang Kienreich	OTFT (Omega tit for tat)	9	36	6	29,554
3	Philip HingstonMod	(Modeller)	7	36	8	29,003
4	Bruno Beaufils	GRAD (Gradual)	8	32	11	28,707
5	Tim Romberg	tro1 (tro1)	13	32	6	28,692
6	Richard Brunauer	DETerminatorL6C4 (DETerminatorL6C4)	12	32	7	28,523
7	Hannes Payer	DETerminatorL4C4 (DETerminatorL4C4)	11	33	7	28,292
8	Bennett McElwee	LOOKDB (LookaheadDB)	22	11	18	28,110
9	Gerhard Mitterlechner	PRobberyM5C4 (PRobberyM5C4)	11	32	8	27,893
10	Wayne Davis	ltft (ltft)	1	44	6	27,834

1.6. Acknowledgements

We would like to thank the following people who acted as reviewers for the chapters in this book.

- Muhammad A. Ahmad
- Oscar Alonso
- Dan Ashlock
- Tsz-Chiu Au
- Carlos Eduardo Rodriguez Calderon
- Michel Charpentier
- Wayne Davis
- Jörg Denzinger
- Eugene Eberbach
- Michael Filzmoser
- Nelis Franken
- Nicholas Gessler
- Michal Glomba
- Philip Hingston
- Enda Howley
- Nick Jennings
- Nanlin Jin
- Jacint Jordana
- Wolfgang Kienreich
- Eun-Youn Kim
- Jia-wei Li
- Helmut A. Mayer
- Bennett McElwee
- Gerhard Mitterlechner
- Colm O'Riordan
- Sarvapali Ramchurn,
- Alex Rogers
- Tim Romberg
- Darryl A. Seale
- Wolfgang Slany
- Elpida Tzafestas
- Perukrishnen Vytelingum
- Georgios N. Yannakakis
- Lukas Zebedin

Appendix: Software Framework

A software library and corresponding application was developed to easily implement prisoner's dilemma strategies and tournament competitions between populations of these. Although a vast array of software is available for the same purpose they did not contain all our feature requirements. For several of our experiments we required a game engine that would, among other things, handle a continuous [normalised] range of moves, arbitrarily sized payoff matrices, different types of signal noise, multiple (> 2) strategies per game, and logging of partial and completed game results.

The software suite was developed in Java, allowing ease in development and web deployment. New strategies are easily implemented by implementing a subclass of the Strategy class. The principal requirements are the implementations of the getMove() and reset() methods which returns the current strategy move and clears the strategy state between games respectively.

Currently we define two types of games: standard and multi-player. A standard game involves two competing strategies playing for a number of rounds, and should mimic the basic game mechanics in the competitions run by Axelrod. A multi-player game involves several competing strategies obtaining payoffs for every other opponent it plays against on each round. A tournament involves every participating strategy and differs for standard and multi-player type games. A standard tournament pits every strategy against every other (including self) in a standard game [a la Round Robin]. A multi-player tournament plays a single multi-player game.

An option is available to introduce a Gaussian distributed random number of rounds to be played, so as to discourage strategies from using the knowledge of a predefined or static parameter for an unfair advantage. There is also an option to introduce noise into the output moves, in principle to test the robustness of the algorithms. Besides the programming API, a graphical user interface is available to set up and run PD tournament competitions (see Figure 1.1).

The software monitors and allows users to log the output of a tournament with different degrees of detail. However, detailed logs will degrade performance.

Besides the standard 2×2 payoff matrix for classic games, there is the ability to define an arbitrarily sized payoff matrix allowing for a wider range of allowable moves. Moves are normalised and payoffs are calculated from the closest allowable move in the payoff matrix.

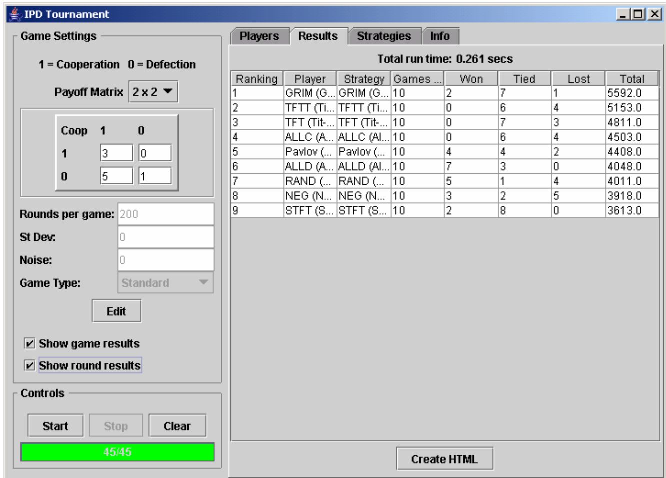

Fig. 1.1. IPD tournament application.

A number of standard classic strategies are included in the library. The software can be downloaded for http://prisoners-dilemma.com.

References

Axelrod, R. (1980a). Effective Choices in the Prisoner's Dilemma, *J. Conflict Resolution*, **24**, pp. 3-25.

Axelrod, R. (1980b). More Effective Choices in the Prisoner's Dilemma, *J. Conflict Resolution*, **24**, pp. 379-403.

Axelrod R. M. (1984). *The Evolution of Cooperation* (BASIC Books, New York).

Axelrod R. and D'Ambrosio L. (1995). Announcement for Bibliography on the Evolution of Cooperation, *Journal of Conflict Resolution* **39**, pp. 190.

Axelrod R. (1997). *The Compleity of Cooperation* (Princeton University Press).

Birk A. (1999). Evolution of Continuous Degrees of Cooperations in an N-Player Iterated Prisoner's Dilemma, Technical Report under review, Vrije Universiteit Brussel, AI-Laboratory.

Boyd R. and Lorberbaum J. P. (1987). No Pure Strategy is Evolutionary Stable in the Repeated Prisoner's Dilemma, **Nature**, 327, pp. 58-59.

Darwen P and Yao X. (2002). Co-Evolution in Iterated Prisoners Dilemma with Intermediate Levels of Cooperation: Application to Missile Defense, *International Journal of Computational Intelligence and Applications*, **2**, 1, pp. 83-107.

Darwen P. and Yao X. (1995). On Evolving Robust Strategies for Iterated Prisoners Dilemma, *In Progress in Evolutionary Computation, LNAI*, **956**, pp. 276-292.

Darwen P. and Yao X. (2001). Why More Choices Cause Less Cooperation in Iterated Prisoner's Dilemma, Proc. *Congress of Evolutionary Computation*, pp. 987-994.

Davis M. Game Theory. (1997). *A Nontechnical Introduction* (Dover Publications).

Fogel D. (1993). Evolving Behaviours in the Iterated Prisoners Dilemma. *Evolutionary Computation*, **1**, 1, pp. 77-97.

Goldstein J. (1991). Reciprocity in Superpower Relations: An Empirical Analysis, *International Studies Quarterly*, **35**, pp. 195-209.

Maynard Smith J. (1982). *Evolution and the Theory of Games* (Cambridge University Press).

Minas J. S., Scodel A., Marlowe D. and Rawson H. (1960). Some Descriptive Aspects of Two-Person, Non-Zero-Sum Games, II, *Journal of Conflict Resolution*, **4**, pp. 193-197.

Nash J. (1950). The Bargaining Problem, *Econometrica*, **18**, pp. 150-155.

Nash J. (1953). Two-Person Cooperative Games, *Econometrica*, **21**, pp. 128-140.

O'Riordan and Bradish S. (2000). Experiments in the Iterated Prisoner's Dilemma and the Voter's Paradox. *11th Irish Conference on Artificial Intelligence and Cognitive Science.*

O'Riordan C. (2000). A Forgiving Strategy for the Iterated Prisoner's Dilemma, *Journal of Artificial Societies and Social Simulation*, **3**, 1.

Poundstone W. (1992). Prisoner's Dilemma, Doubleday

Rapoport A. (1996). Optimal policies for the prisoners dilemma, Tech report No. 50, Psychometric Laboratory, Univ. North Carolina, NIH Grant, MH-10006.

Scodel A. and Philburn R. (1959). Some Personality Correlates of Decision Making under Conditions of Risk, *Behavioral Science*, **4**, pp. 19-28.

Scodel A., Minas J. S., Ratoosh P.and Lipetz M. (1959). Some Descriptive Aspects of Two-Person, Non-Zero-Sum Games, *Journal of Conflict Resolution*, **3**, pp. 114-119.

Scodel A. and Minas J. S. (1960). The Behavior of Prisoners in a "Prisoner's Dilemma" Game, *Journal of Psychology*, **50**, pp. 133-138.

Scodel A. (1962). Induced Collaboration in Some Non-Zero-Sum Games, *Journal of Conflict Resolution*, **6**, pp. 335-340.

Scodel A. (1963). Probability Preferences and Expected Values. *Journal of Psychology*, **56**, pp. 429-434.

von Neumann J. and Morgenstern O. (1944). *Theory of Games and Economic Behavior* (Princeton University Press).

Yao, X and Darwen P. (1999). How Important is Your Reputation in a Multi-Agent Environment. *Proc. Of the 1999 IEEE Conference on Systems, Man and Cybernetics*, IEEE Press, Piscataway, NJ, USA, pp. II-575 – II-580, Oct.