

Multiagent Systems

second edition

edited by Gerhard Weiss



The MIT Press
Cambridge, Massachusetts
London, England

Contents

Preface

xxxv

The Subject of This Book • Main Features of This Book •
 Readership and Prerequisites • Changes from the First Edition •
 Structure and Chapters • The Exercises • How to Use This Book •
 Slides and More – The Website of the Book • Acknowledgments

Contributing Authors

xliii

Part I Agent Architectures and Organizations **1**

1 Intelligent Agents **3**

Michael Wooldridge

1	Introduction	3
2	What Are Agents?	4
2.1	Examples of Agents	7
2.2	Intelligent Agents	8
2.3	Agents and Objects	10
2.4	Agents and Expert Systems	12
2.5	Sources and Further Reading	13
3	Architectures for Intelligent Agents	13
3.1	Logic-Based Architectures	14
3.1.1	Sources and Further Reading	19
3.2	Reactive Architectures	20
3.2.1	The Subsumption Architecture	21
3.2.2	Markov Decision Processes	25
3.2.3	Sources and Further Reading	27
3.3	Belief-Desire-Intention Architectures	28
3.3.1	Sources and Further Reading	35
3.4	Layered Architectures	36

3.4.1	TouringMachines	38
3.4.2	InteRRaP	40
3.4.3	Sources and Further Reading	42
4	Conclusions	42
5	Exercises	42
	References	45
2	Multiagent Organizations	51
	<i>Virginia Dignum and Julian Padget</i>	
1	Introduction	51
2	Background	53
2.1	From Intelligent Agents to Multiagent Systems	53
2.2	From Multiagent Systems to Multiagent Organizations	55
2.3	Sources of Inspiration	56
2.3.1	Organization as Structure	56
2.3.2	Organization as Institution	58
2.3.3	Organization as Agent	59
2.4	Autonomy and Regulation	60
2.5	Example Scenario	62
3	Multiagent Organizations	62
3.1	Organization Concepts	64
3.2	Example of Organization Modeling: The OperA Framework	65
3.2.1	The Social Structure	68
3.2.2	The Interaction Structure	70
3.2.3	The Normative Structure	71
3.2.4	The Communication Structure	72
4	Institutions	72
4.1	Organizations, Institutions, and Norms	73
4.2	Events and States	75
4.3	Obligations, Permission, and Power	77
4.4	Example of Institutional Modeling: InstAL	78
4.4.1	The Formal Model	78
4.4.2	The Conference Scenario	78
5	Agents in Organizations	82
6	Evolution of Organizations	85
6.1	Organizational Adaptation	86
6.2	Emergent Organizations	87
7	Conclusions	88
8	Exercises	89
	References	92

Part II Communication **99**

3 Agent Communication **101**

Amit K. Chopra and Munindar P. Singh

1	Introduction	101
1.1	Autonomy and Its Implications	102
1.2	Criteria for Evaluation	105
2	Conceptual Foundations of Communication in MAS	106
2.1	Communicative Acts	106
2.2	Agent Communication Primitives	107
3	Traditional Software Engineering Approaches	108
3.1	Choreographies	110
3.2	Sequence Diagrams	111
3.3	State Machines	112
3.4	Evaluation with Respect to MAS	113
4	Traditional AI Approaches	114
4.1	KQML	115
4.2	FIPA ACL	116
4.3	Evaluation with Respect to MAS	117
5	Commitment-Based Multiagent Approaches	118
5.1	Commitments	118
5.2	Commitment Protocol Specification	119
5.3	Evaluation with Respect to MAS	121
6	Engineering with Agent Communication	122
6.1	Programming with Communications	123
6.2	Modeling Communications	124
6.2.1	Business Patterns	125
6.2.2	Enactment Patterns	125
6.2.3	Semantic Antipatterns	126
6.3	Communication-Based Methodologies	127
7	Advanced Topics and Challenges	128
7.1	Primacy of Meaning	128
7.2	Verifying Compliance	129
7.3	Protocol Refinement and Aggregation	129
7.4	Role Conformance	130
8	Conclusions	130
9	Exercises	133
	References	136

4 Negotiation and Bargaining 143

Shaheen Fatima and Iyad Rahwan

1	Introduction	143
2	Aspects of Negotiation	144
3	Game-Theoretic Approaches for Single-Issue Negotiation	146
3.1	Cooperative Models of Single-Issue Negotiation	147
3.2	Non-Cooperative Models of Single-Issue Negotiation	151
4	Game-Theoretic Approaches for Multi-Issue Negotiation	156
4.1	Cooperative Models of Multi-Issue Negotiation	157
4.2	Non-Cooperative Models of Multi-Issue Negotiation	157
5	Heuristic Approaches for Multi-Issue Negotiation	161
5.1	Heuristics for Generating Counteroffers	161
5.2	Heuristics for Predicting Opponent's Preferences and Generating Counteroffers	163
5.3	Heuristics for Generating Optimal Agendas	164
5.4	Heuristics for Reasoning about Deliberation Cost	164
6	Negotiating with Humans	165
7	Argumentation-Based Negotiation	167
8	Conclusions	169
9	Exercises	170
	References	171

5 Argumentation among Agents 177

Iyad Rahwan

1	Introduction	177
2	What Is an Argument?	178
2.1	Arguments as Chained Inference Rules	178
2.2	Argument as an Instance of a Scheme	180
2.3	Abstract Arguments	181
3	Evaluating an Argument	181
4	Argumentation Protocols	185
4.1	Abstract Argument Games	186
4.2	Dialogue Systems	188
5	Strategic Argumentation and Game Theory	190
5.1	Glazer and Rubinstein's Model	191
5.2	Game Theory Background	192
5.2.1	Mechanism Design	194
5.2.2	The Revelation Principle	195
5.3	Argumentation Mechanism Design	196
5.4	Case Study: Implementing the Grounded Semantics	198

6	The Argument Interchange Format	201
7	Conclusion	204
8	Exercises	205
	References	206

Part III Basic Coordination 211

6 Computational Social Choice 213

Felix Brandt, Vincent Conitzer, and Ulle Endriss

1	Introduction	213
1.1	Introductory Example	214
1.2	History of the Field	216
1.3	Applications	217
1.4	Chapter Outline	219
2	Preference Aggregation	219
2.1	Social Welfare Functions	219
2.2	Social Choice Functions	223
2.2.1	The Weak Axiom of Revealed Preference	223
2.2.2	Contraction and Expansion Consistency	224
3	Voting	226
3.1	Voting Rules	227
3.1.1	Scoring Rules	227
3.1.2	Condorcet Extensions	229
3.1.3	Other Rules	231
3.2	Manipulation	232
3.2.1	The Gibbard-Satterthwaite Impossibility	233
3.2.2	Restricted Domains of Preferences	233
3.2.3	Computational Hardness of Manipulation	235
3.2.4	Probabilistic Voting Rules	238
3.2.5	Irresolute Voting Rules	239
3.3	Possible and Necessary Winners	240
4	Combinatorial Domains	241
4.1	Preference Representation	243
4.2	Sequential Voting	245
4.3	Voting with Compactly Represented Preferences	246
5	Fair Division	247
5.1	Preference Representation	249
5.2	Fairness and Efficiency	251
5.3	Computing Fair and Efficient Allocations	253

5.4	Convergence to Fair and Efficient Allocations	255
6	Conclusion	257
6.1	Additional Topics	258
6.2	Further Reading	259
7	Exercises	260
	References	266
7	Mechanism Design and Auctions	285
	<i>Kevin Leyton-Brown and Yoav Shoham</i>	
1	Introduction	285
2	Mechanism Design with Unrestricted Preferences	286
2.1	Implementation	287
2.2	The Revelation Principle	288
2.3	Impossibility of General, Dominant-Strategy Implementation	291
3	Quasilinear Preferences	291
3.1	Mechanism Design in the Quasilinear Setting	292
4	Efficient Mechanisms	296
4.1	Groves Mechanisms	297
4.2	The VCG Mechanism	299
4.3	Properties of VCG	301
4.3.1	VCG and Individual Rationality	301
4.3.2	VCG and Weak Budget Balance	302
4.3.3	Drawbacks of VCG	302
4.4	Budget Balance and Efficiency	303
5	Single-Good Auctions	304
5.1	Canonical Auction Families	305
5.1.1	English Auctions	305
5.1.2	Japanese Auctions	305
5.1.3	Dutch Auctions	306
5.1.4	Sealed-Bid Auctions	306
5.2	Auctions as Bayesian Mechanisms	306
5.3	Second-Price, Japanese, and English Auctions	308
5.4	First-Price and Dutch Auctions	310
5.5	Revenue Equivalence	311
6	Position Auctions	312
7	Combinatorial Auctions	315
8	Conclusions	318
9	Exercises	320
	References	325

8 Computational Coalition Formation 329

Edith Elkind, Talal Rahwan, and Nicholas R. Jennings

1	Introduction	329
1.1	Coalitional Games: A Bird's Eye View	330
2	Definitions	331
2.1	Outcomes	332
2.2	Subclasses of Characteristic Function Games	333
2.2.1	Monotone Games	333
2.2.2	Superadditive Games	333
2.2.3	Convex Games	334
2.2.4	Simple Games	335
3	Solution Concepts	335
3.1	Shapley Value	335
3.2	Banzhaf Index	337
3.3	Core	338
3.3.1	The Core of Simple Games	339
3.3.2	The Core of Convex Games	340
3.4	The Least Core	341
3.5	Other Solution Concepts	342
4	Representation Formalisms	343
4.1	Weighted Voting Games	344
4.1.1	Computational Issues	345
4.1.2	Expressivity and Vector Weighted Voting Games	346
4.2	Combinatorial Optimization Games	348
4.2.1	Induced Subgraph Games	348
4.2.2	Network Flow Games	348
4.2.3	Matching and Assignment Games	349
4.3	Complete Representation Languages	349
4.3.1	Marginal Contribution Nets	349
4.3.2	Synergy Coalition Groups	350
4.3.3	Skill-Based Representations	351
4.3.4	Agent-Type Representation	351
5	Coalition Structure Generation	352
5.1	Space Representation	353
5.2	Dynamic Programming Algorithms	354
5.3	Anytime Algorithms	356
5.3.1	Identifying Subspaces with Worst-Case Guarantees	356
5.3.2	Integer Partition-Based Search	359
5.3.3	Integer Programming	360

5.4	Metaheuristic Algorithms	361
5.5	Coalition Structure Generation under Compact Representations	362
5.5.1	Distributed Constraint Optimization	362
5.5.2	Marginal Contribution Nets	364
5.5.3	Coalitional Skill Games	366
5.5.4	Agent-Type Representation	368
5.6	Constrained Coalition Formation	369
6	Conclusions	372
7	Exercises	372
	References	374

9 Trust and Reputation in Multiagent Systems 381

Jordi Sabater-Mir and Laurent Vercouter

1	Introduction	381
2	Computational Representation of Trust and Reputation Values . .	382
2.1	Boolean Representation	383
2.2	Numerical Values	383
2.3	Qualitative Labels	384
2.4	Probability Distribution and Fuzzy Sets	384
2.5	Trust and Reputation as Beliefs	385
2.6	The Reliability of a Value	387
3	Trust Processes in Multiagent Systems	388
3.1	General Overview of Trust-Related Processes	388
3.2	Trust Evaluations	390
3.2.1	Filtering the Inputs	391
3.2.2	Statistical Aggregation	392
3.2.3	Logical Beliefs Generation	393
3.3	Trust Decision	394
3.3.1	Single Trust Values and Probability Distributions	395
3.3.2	Trust Beliefs	395
3.4	Coping with the Diversity of Trust Models	396
4	Reputation in Multiagent Societies	396
4.1	Reputation-Building Process	398
4.1.1	Communicated Image as a Source for Reputation	398
4.1.2	Communicated Reputation	400
4.1.3	Inherited Reputation	400
4.1.4	Putting It All Together	401

4.2	Centralized vs. Decentralized Models	402
4.2.1	Centralized Approaches	402
4.2.2	Decentralized Approaches	403
4.3	Using Reputation	404
4.3.1	Reputation as a Source of Trust	404
4.3.2	Reputation for Social Order	405
4.4	Pitfalls When Using Reputation	405
4.4.1	Unfair Ratings	405
4.4.2	Ballot-Stuffing	406
4.4.3	Dynamic Personality	406
4.4.4	Whitewashing	406
4.4.5	Collusion	406
4.4.6	Sybil Attacks	407
4.4.7	Reputation Lag Exploitation	407
5	Trust, Reputation, and Other Agreement Technologies	407
5.1	Argumentation	408
5.2	Negotiation	410
5.3	Norms	410
5.4	Organizations	411
5.5	Ontologies and Semantics	411
6	Conclusions	413
7	Exercises	414
	References	415

Part IV Distributed Cognitive Abilities 421

10 Multiagent Learning 423

Karl Tuyls and Kagan Tumer

1	Introduction	423
2	Challenges in Multiagent Learning	425
2.1	State, Action, and Outcome Space Problems	426
2.2	Multiagent Credit Assignment Problem	426
2.3	Agent Rewards and System Dynamics	427
2.4	Two Simple Multiagent Learning Paradigms	429
2.4.1	Action-Value Learning	430
2.4.2	Direct Policy Adjustment	431
3	Reinforcement Learning for Multiagent Systems	432
3.1	Markov Decision Processes	433
3.2	Action Selection and Exploration-Exploitation Dilemma	434

3.3	Model-Free and Model-Based Approaches	435
3.4	Multiagent MDP Formulations	437
3.5	Markov Games	438
3.6	State-of-the-Art Algorithms	439
3.6.1	Joint Action Learning	439
3.6.2	Nash-Q Learning	440
3.6.3	Gradient Ascent Algorithms	441
3.6.4	Other Approaches	442
4	Evolutionary Game Theory as a Multiagent Learning Paradigm . .	443
4.1	Matrix Games	443
4.2	Solution Concepts	444
4.3	Evolutionary Stable Strategies	445
4.4	Replicator Dynamics	446
4.5	The Role of Game Theory for Multiagent Learning	447
4.6	Evolutionary Game Theory as a Theoretical Framework .	448
5	Swarm Intelligence as a Multiagent Learning Paradigm	451
5.1	Ant Colony Optimization	453
5.2	Bee Colony Optimization	455
6	Neuro-Evolution as a Multiagent Learning Paradigm	457
6.1	Evolutionary Algorithm Basics	458
6.2	Linking Multiagent Reinforcement Learning to the Neuro-Evolutionary Approach	460
7	Case Study: Air Traffic Control	460
7.1	Motivation	460
7.2	Simulation and System Performance	461
7.3	Agent-Based Air Traffic	463
7.4	Multiagent Air Traffic Results	466
7.5	Summary	467
8	Conclusions	468
9	Exercises	468
	References	475
11	Multiagent Planning, Control, and Execution	485
	<i>Ed Durfee and Shlomo Zilberstein</i>	
1	Introduction	485
2	Characterizing Multiagent Planning and Control	487
3	Coordination Prior to Local Planning	488
3.1	Social Laws and Conventions	489
3.2	Organizational Structuring	490
3.2.1	Organizational Design	491

	3.2.2	Organizational Execution and Functionally-Accurate Cooperation	493
	3.3	The Contract-Net Protocol and Role Assignment	495
4		Local Planning Prior to Coordination	497
	4.1	State-Space Techniques	498
	4.2	Plan-Space Techniques	499
	4.2.1	Single-Agent Plans	499
	4.2.2	Multiagent Plans	501
	4.2.3	Multiagent Plan Coordination by Plan Modification	505
	4.3	Hierarchical Multiagent Plan Coordination	510
5		Decision-Theoretic Multiagent Planning	512
	5.1	Models for Decision-Theoretic Multiagent Planning	513
	5.1.1	Solution Representation and Evaluation	515
	5.1.2	The Complexity of DEC-POMDPs	518
	5.2	Solving Finite-Horizon DEC-POMDPs	519
	5.3	Solving Infinite-Horizon DEC-POMDPs	522
	5.3.1	Correlated Joint Controllers	523
	5.3.2	Policy Iteration for Infinite-Horizon DEC-POMDPs	524
	5.3.3	Optimizing Fixed-Size Controllers Using Non-Linear Programming	526
6		Multiagent Execution	527
	6.1	Multiagent Plan Monitoring	527
	6.2	Multiagent Plan Recovery	528
	6.3	Multiagent Continuous Planning	529
7		Conclusions	532
8		Exercises	533
		References	539

12 Distributed Constraint Handling and Optimization 547

Alessandro Farinelli, Meritxell Vinyals, Alex Rogers, and Nicholas R. Jennings

1		Introduction	547
2		Distributed Constraint Handling	549
	2.1	Constraint Networks	549
	2.2	Distributed Constraint Processing	550
3		Applications and Benchmarking Problems	551
	3.1	Real-World Applications	551
	3.1.1	Meeting Scheduling	552

	3.1.2	Target Tracking	553
	3.2	Exemplar and Benchmarking Problems	553
4		Solution Techniques: Complete Algorithms	554
	4.1	Search-Based: ADOPT	555
	4.2	Dynamic Programming: DPOP	561
5		Solution Techniques: Approximate Algorithms	565
	5.1	Local Greedy Approximate Algorithms	565
	5.1.1	The Distributed Stochastic Algorithm	566
	5.1.2	The Maximum Gain Message Algorithm	567
	5.2	GDL-Based Approximate Algorithms	568
	5.2.1	The Max-Sum Algorithm	568
6		Solution Techniques with Quality Guarantees	570
	6.1	Off-line Guarantees	571
	6.2	Online Guarantees	574
7		Conclusions	577
8		Exercises	578
		References	580

Part V Development and Engineering 585

13 Programming Multiagent Systems 587

Rafael H. Bordini and Jürgen Dix

1		Introduction	587
	1.1	Relation to Other Chapters	589
	1.2	Organization of This Chapter	589
2		From AGENT0 to Modern Agent Languages	590
	2.1	A Brief History of Agent-Oriented Programming (AOP)	590
	2.2	Features of Multiagent-Oriented Programming (MAOP)	591
3		Abstractions in the MAOP Paradigm	593
	3.1	Agent Level	593
	3.2	Environment Level	595
	3.3	Social Level	596
4		Examples of Agent Programming Languages	596
	4.1	JASON	596
	4.1.1	Beliefs	597
	4.1.2	Goals	598
	4.1.3	Plans	599
	4.1.4	Semantics	600
	4.2	Other BDI-Based Languages	601

4.3	Approaches Based on Executable Logics	603
4.3.1	METATEM	607
4.3.2	ConGolog and IndiGolog	608
5	Organization and Environment Programming	609
5.1	Organizations	609
5.1.1	MOISE	610
5.1.2	Other Approaches	611
5.2	Environments	611
5.2.1	CARTAGO	613
5.2.2	EIS	617
6	Example of Full MAOP in JaCaMo	620
6.1	The Application Scenario	621
6.2	Organization Program	623
6.3	Agent Programs	625
6.4	Environment Program	628
7	Conclusions	629
8	Exercises	630
	References	633

14 Specification and Verification of Multiagent Systems 641

Jürgen Dix and Michael Fisher

1	Introduction	641
1.1	Why Logic, Specification, and Verification?	642
1.2	Limits and Relation to Other Chapters	644
1.3	Organization of This Chapter	644
2	Agent Specification	644
2.1	Logics of Agency and Specification Languages	644
2.2	Approaches Based on Temporal Logics	646
2.2.1	LTL	648
2.2.2	CTL and CTL*	649
2.2.3	ATL and ATL*	650
2.3	Approaches Based on Dynamic Logic	652
2.4	Combinations	653
2.4.1	BDI	653
2.4.2	KARO	653
2.4.3	Dynamic Epistemic Logic	654
2.5	Sample Specifications	654
3	From Specifications to Implementations	656
3.1	Toward Formal Verification	656
3.2	Refinement	657

3.3	Synthesis	657
3.4	Specifications as Programs	658
4	Formal Verification	659
4.1	What Is Formal Verification?	659
4.2	Deductive Verification	660
4.3	Algorithmic Verification	660
4.4	Program Verification	662
4.5	Runtime Verification	663
5	Deductive Verification of Agents	663
5.1	The Problem	664
5.2	Direct Proof	665
5.3	Use of Logic Programming	666
5.4	Example	667
6	Algorithmic Verification of Agent Models	667
6.1	The Representation and Size of the Model	668
6.2	(Im-)Perfect Information, (Im-)Perfect Recall	669
6.3	Modular Interpreted Systems	671
6.4	MC Complexity for LTL, CTL, ATL, and MIS	672
6.5	Model Checking Agent Language Models	674
7	Algorithmic Verification of Agent Programs	676
7.1	General Problem	676
7.2	AIL Semantic Toolkit	677
7.3	Multiple Semantic Definitions	678
7.4	Model Checking AIL Through MCAPL/AJPF	679
7.5	Example	679
8	Conclusions	680
9	Exercises	681
	References	683
15	Agent-Oriented Software Engineering	695
	<i>Michael Winikoff and Lin Padgham</i>	
1	Introduction	695
1.1	History of AOSE	698
2	Agent Concepts	700
3	Running Example	702
4	Requirements	704
5	Design	710
6	Detailed Design	717
6.1	Example Design: BDI Platform	719
6.1.1	Initial Structure	719

6.1.2	Subgoal Structure for build2	720
6.1.3	Subgoal Structure for addPart and complete . . .	720
6.1.4	Finalizing the Design	722
6.1.5	Multiple Plans	724
6.1.6	Control Information	725
6.2	Example Design: Finite-State Automaton	725
6.3	Final Features	726
7	Implementation	727
8	Assurance	728
8.1	Testing and Debugging	729
8.2	Formal Methods	732
9	Software Maintenance	734
10	Comparing Methodologies	735
11	Conclusions	736
12	Exercises	740
	Appendix: Agent UML Sequence Diagram Notation	742
	References	744

Part VI Technical Background 759

16 Logics for Multiagent Systems 761

Wiebe van der Hoek and Michael Wooldridge

1	Introduction	761
1.1	A Logical Toolkit	764
2	Representing Cognitive States	769
2.1	Intention Logic	770
2.2	BDI Logic	773
2.3	Discussion	778
2.4	Cognitive Agents in Practice	778
2.4.1	Specification Language	778
2.4.2	Implementation	780
2.4.3	Verification	782
3	Representing the Strategic Structure of a System	785
3.1	Coalition Logic	786
3.2	Strategic Temporal Logic: ATL	790
3.3	Knowledge in Strategic Temporal Logics: ATEL	794
3.4	CL-PC	796
4	Conclusion and Further Reading	797

5	Exercises	799
	References	800
17	Game-Theoretic Foundations of Multiagent Systems	811
	<i>Edith Elkind and Evangelos Markakis</i>	
1	Introduction	811
2	Normal-Form Games	812
2.1	Dominant Strategy	814
2.2	Nash Equilibrium	816
2.3	Mixed Strategies and Mixed Nash Equilibrium	818
2.4	Elimination of Dominated Strategies	820
2.5	Games with Infinite Action Spaces	821
2.5.1	Games with Differentiable Payoff Functions	823
2.6	Zero-Sum Games	824
2.7	Computational Aspects	827
3	Extensive-Form Games	828
3.1	Nash Equilibrium and Critiques	831
3.2	Subgame-Perfect Equilibrium	832
3.3	Backward Induction	834
4	Bayesian Games	836
4.1	Two Examples	837
4.2	Formal Definitions	840
5	Conclusions	842
6	Exercises	842
	References	847
	Subject Index	849