

16th International Conference on Autonomous Agents and Mutliagent Systems (AAMAS 2017)

Sao Paulo, Brazil
8 - 12 May 2017

Volume 1 of 3

ISBN: 978-1-5108-5507-6

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2017) by International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) All rights reserved.

Printed by Curran Associates, Inc. (2018)

For permission requests, please contact International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) at the address below.

info@ifaamas.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

AAMAS 2017

SIXTEENTH International Conference on
Autonomous Agents & Multiagent Systems

www.aamas2017.org



São Paulo • Brazil

May 8 - 12, 2017

Table of Contents

Chairs' Welcome

Kate Larson (*University of Waterloo*)

Michael Winikoff (*University of Otago*)

Sanmay Das (*Washington University in St. Louis*)

Ed Durfee (*University of Michigan*)

Keynote Talks	Session 3E: Virtual Agents
Session 1A: Computational Social Choice 1	Session 3F: Transportation and Routing
Session 1B: Computational Game Theory & Mechanism Design 1	Session 4A: Computational Social Choice 3
Session 1C: Logics 1	Session 4B: Computational Game Theory & Mechanism Design 3
Session 1D: Novel Applications in Safety and Security	Session 4C: Judgment Aggregation and Argumentation
Session 1E: Constraint Reasoning	Session 4D: Novel Applications in Smart Grids and Mobility
Session 1F: Engineering Agent Systems	Session 4E: Networking and Communication
Session 2A: Fair Division and Cooperative Game Theory	Session 4F: Teamwork, Coordination, and Cooperation
Session 2B: Assignment and Matching Problems	Session 5A: Computational Social Choice 4
Session 2C: Negotiation and Virtual Agents	Session 5B: Logic & Game Theory
Session 2D: Agent-Based and Social Simulation	Session 5C: Logics 3
Session 2E: Multi-Robot Coordination and Swarm Robotics	Session 5D: Social Networks
Session 2F: Learning	Session 5E: Human-Robot Interaction and Learning in Robotics

Session 3A: Computational Social Choice 2	Session 5F: Emergence
Session 3B: Computational Game Theory & Mechanism Design 2	Extended Abstracts
Session 3C: Logics 2	Demonstration Abstracts
Session 3D: Planning	Doctoral Consortium Abstracts

Keynote Presentations

Trying to Fix Traffic: Past, Present, and Some Future Trends (Page 1)

Ana L. C. Bazzan (*Univ. Federal do Rio Grande do Sul (UFRGS)*)

On AI, Markets and Machine Learning (Page 2)

David C. Parkes (*Harvard University*)

Active Optimization and Self Driving Cars (Page 3)

Jeff Schneider (*Carnegie Mellon University / Uber ATG*)

Enhancing Human Capability with Intelligent Machine Teammates (Page 4)

Julie A. Shah (*Massachusetts Institute of Technology*)

Optimal Social Decision Making (Page 5)

Nisarg Shah (*Harvard University*)

Session 1A: Computational Social Choice 1

Bribery as a Measure of Candidate Success: Complexity Results for Approval-Based Multiwinner Rules (Page 6)

Piotr Faliszewski (*AGH University*)

Piotr Skowron (*TU Berlin*)

Nimrod Talmon (*Weizmann Institute of Science*)

Proportional Representation in Vote Streams (Page 15)

Palash Dey (*Tata Institute of Fundamental Research*)

Nimrod Talmon (*Weizmann Institute of Science*)

Otniel van Handel (*Weizmann Institute of Science*)

Beyond Electing and Ranking: Collective Dominating Chains, Dominating Subsets and Dichotomies (Page 24)

Jérôme Lang (*Université Paris-Dauphine*)

Jérôme Monnot (*Université Paris-Dauphine*)

Arkadii Slinko (*University of Auckland*)

William S. Zwicker (*Union College*)

Condorcet Consistent Bundling with Social Choice (Page 33)

Shreyas Sekar (*Rensselaer Polytechnic Institute*)

Sujoy Sikdar (*Rensselaer Polytechnic Institute*)

Lirong Xia (*Rensselaer Polytechnic Institute*)

Parameterized Dichotomy of Choosing Committees Based on Approval Votes in the Presence of Outliers (Page 42)

Palash Dey (*Tata Institute of Fundamental Research*)

Neeldhara Misra (*Indian Institute of Technology*)

Y. Narahari (*Indian Institute of Science*)

Session 1B: Computational Game Theory & Mechanism Design 1

SATS: A Universal Spectrum Auction Test Suite (Page 51)

Michael Weiss (*University of Zurich*)

Benjamin Lubin (*Boston University*)

Sven Seuken (*University of Zurich*)

Generalizing Demand Response Through Reward Bidding (Page 60)

Hongyao Ma (*Harvard University*)

David C. Parkes (*Harvard University*)

Valentin Robu (*Heriot Watt University*)

Mechanism Design with Unknown Correlated Distributions: Can We Learn Optimal Mechanisms? (Page 69)

Michael Albert (*Duke University*)

Vincent Conitzer (*Duke University*)

Peter Stone (*University of Texas at Austin*)

Practical versus Optimal Mechanisms (Page 78)

Weiran Shen (*Tsinghua University*)

Pingzhong Tang (*Tsinghua University*)

Thompson Sampling Based Mechanisms for Stochastic Multi-Armed Bandit Problems (Page 87)

Ganesh Ghalme (*Indian Institute of Science*)
Shweta Jain (*Indian Institute of Science*)
Sujit Gujar (*International Institute of Information Technology*)
Y. Narahari (*Indian Institute of Science*)

Session 1C: Logics 1

A Declarative Modular Framework for Representing and Applying Ethical Principles (Page 96)

Fiona Berreby (*Sorbonne Universités & UPMC*)
Gauvain Bourgne (*Sorbonne Universités & UPMC*)
Jean-Gabriel Ganascia (*Sorbonne Universités & UPMC*)

The Event Calculus in Probabilistic Logic Programming with Annotated Disjunctions (Page 105)

Kevin McAreavey (*Queen's University Belfast*)
Kim Bauters (*University of Bristol*)
Weiru Liu (*University of Bristol*)
Jun Hong (*University of the West of England*)

Symbolic Model Checking Multi-Agent Systems against CTL*K Specifications (Page 114)

Jeremy Kong (*Imperial College London*)
Alessio Lomuscio (*Imperial College London*)

A Succinct Language for Dynamic Epistemic Logic (Page 123)

Tristan Charrier (*IRISA*)
François Schwarzentruber (*ENS Rennes*)

Exploring the Bidimensional Space: A Dynamic Logic Point of View (Page 132)

Philippe Balbiani (*Toulouse University*)
David Fernández-Duque (*Toulouse University*)
Emiliano Lorini (*Toulouse University*)

Session 1D: Novel Applications in Safety and Security

Coordinating Vessel Traffic to Improve Safety and Efficiency (Page 141)

Teck-Hou Teng (*Singapore Management University*)
Hoong Chuin Lau (*Singapore Management University*)
Akshat Kumar (*Singapore Management University*)

Influence Maximization in the Field: The Arduous Journey from Emerging to Deployed Application (Page 150)

Amulya Yadav (*University of Southern California*)
Bryan Wilder (*University of Southern California*)
Eric Rice (*University of Southern California*)
Robin Petering (*University of Southern California*)
Jaih Craddock (*University of Southern California*)
Amanda Yoshioka-Maxwell (*University of Southern California*)
Mary Hemler (*University of Southern California*)
Laura Onasch-Vera (*University of Southern California*)
Milind Tambe (*University of Southern California*)
Darlene Woo (*University of Southern California*)

Cloudy with a Chance of Poaching: Adversary Behavior Modeling and Forecasting with Real-World Poaching Data (Page 159)

Debarun Kar (*University of Southern California & USC Center for AI in Society*)
Benjamin Ford (*University of Southern California & USC Center for AI in Society*)
Shahzad Gholami (*University of Southern California & USC Center for AI in Society*)
Fei Fang (*Harvard University*)
Andrew Plumptre (*Wildlife Conservation Society*)
Milind Tambe (*University of Southern California & USC Center for AI in Society*)
Margaret Driciru (*Uganda Wildlife Authority*)
Fred Wanyama (*Uganda Wildlife Authority*)
Aggrey Rwetsiba (*Uganda Wildlife Authority*)
Mustapha Nsubaga (*Wildlife Conservation Society*)
Joshua Mabonga (*Wildlife Conservation Society*)

Prioritized Allocation of Emergency Responders based on a Continuous-Time Incident Prediction Model (Page 168)

Ayan Mukhopadhyay (*Vanderbilt University*)
Yevgeniy Vorobeychik (*Vanderbilt University*)
Abhishek Dubey (*Vanderbilt University*)
Gautam Biswas (*Vanderbilt University*)

A Game Theoretic Approach to Strategy Generation for Moving Target Defense in Web Applications (Page 178)

Sailik Sengupta (*Arizona State University*)
Satya Gautam Vadlamudi (*Arizona State University & Capillary Technologies*)
Subbarao Kambhampati (*Arizona State University*)

Adam Doupé (*Arizona State University*)
Ziming Zhao (*Arizona State University*)
Marthony Taguinod (*Arizona State University*)
Gail-Joon Ahn (*Arizona State University*)

Session 1E: Constraint Reasoning

A Partial Decision Scheme for Local Search Algorithms for Distributed Constraint Optimization

Problems (Page 187)

Zhepeng Yu (*Chongqing University*)
Ziyu Chen (*Chongqing University*)
Jingyuan He (*Chongqing University*)
Yancheng Deng (*Chongqing University*)

An Iterative Refined Max-sum_AD Algorithm via Single-side Value Propagation and Local Search

 (Page 195)

Ziyu Chen (*Chongqing University*)
Yanchen Deng (*Chongqing University*)
Tengfei Wu (*Chongqing University*)

A Deterministic Distributed Algorithm for Reasoning with Connected Row-Convex Constraints

 (Page 203)

Shufeng Kong (*University of Technology, Sydney*)
Jae Hee Lee (*University of Technology, Sydney*)
Sanjiang Li (*University of Technology, Sydney*)

Infinite-Horizon Proactive Dynamic DCOPs

 (Page 212)

Khoi D. Hoang (*New Mexico State University*)
Ping Hou (*New Mexico State University*)
Ferdinando Fioretto (*University of Michigan*)
William Yeoh (*New Mexico State University*)
Roie Zivan (*Ben Gurion University of the Negev*)
Makoto Yokoo (*Kyushu University*)

Proactive Project Scheduling with Time-dependent Workability Uncertainty

 (Page 221)

Wen Song (*Nanyang Technological University*)
Donghun Kang (*Nanyang Technological University*)
Jie Zhang (*Nanyang Technological University*)
Hui Xi (*Rolls-Royce Singapore Pte Ltd*)

Session 1F: Engineering Agent Systems

Arnor: Modeling Social Intelligence via Norms to Engineer Privacy-Aware Personal Agents

 (Page 230)

Nirav Ajmeri (*North Carolina State University*)
Pradeep K. Murukannaiah (*Rochester Institute of Technology*)
Hui Guo (*North Carolina State University*)
Munindar P. Singh (*North Carolina State University*)

Decentralizing MAS Monitoring with DecAMon

 (Page 239)

Angelo Ferrando (*University of Genova*)
Davide Ancona (*University of Genova*)
Viviana Mascardi (*University of Genova*)

Integrating BDI Agents with Agent-Based Simulation Platforms: (JAAMAS Extended Abstract)

 (Page 249)

Dhirendra Singh (*MIT*)
Lin Padgham (*MIT*)
Brian Logan (*University of Nottingham*)

Debugging Agent Programs with Why? Questions

 (Page 251)

Michael Winikoff (*University of Otago*)

BDI Agent Testability Revisited (JAAMAS Extended Abstract)

 (Page 260)

Michael Winikoff (*University of Otago*)

Session 2A: Fair Division and Cooperative Game Theory

Approximate Solutions To Max-Min Fair and Proportionally Fair Allocations of Indivisible Goods

 (Page 262)

Nhan-Tam Nguyen (*Heinrich-Heine-Universität*)
Trung Thanh Nguyen (*New York University*)
Jörg Rothe (*Heinrich-Heine-Universität*)

The Atkinson Inequality Index in Multiagent Resource Allocation

 (Page 272)

Sebastian Schneckeburger (*University of Tübingen*)
Britta Dorn (*University of Tübingen*)
Ulle Endriss (*University of Amsterdam*)

Fair Division via Social Comparison

 (Page 281)

Rediet Abebe (*Cornell University*)
Jon Kleinberg (*Cornell University*)

David C. Parkes (*Harvard University*)

Stability of Generalized Two-sided Markets with Transaction Thresholds (Page 290)

Zhiyuan Li (*Tsinghua University*)

Yicheng Liu (*Tsinghua University*)

Pingzhong Tang (*Tsinghua University*)

Tingting Xu (*Tsinghua University*)

Wei Zhan (*Tsinghua University*)

Supermodular Games on Social Networks (Page 299)

Ayumi Igarashi (*University of Oxford*)

Coalition Structure Generation and CS-core: Results on the Tractability Frontier for games represented by MC-nets (Page 308)

Julien Lesca (*Université Paris-Dauphine*)

Patrice Perny (*Sorbonne Universités & UPMC*)

Makoto Yokoo (*Kyushu University*)

Session 2B: Assignment and Matching Problems

Weighted Matching Markets with Budget Constraints (Page 317)

Naoto Hamada (*Kyushu University*)

Anisse Ismaili (*Kyushu University*)

Takamasa Suzuki (*Kyushu University*)

Makoto Yokoo (*Kyushu University*)

Random Assignment with Optional Participation (Page 326)

Florian Brandl (*TU München*)

Felix Brandt (*TU München*)

Johannes Hofbauer (*TU München*)

Majority Graphs of Assignment Problems and Properties of Popular Random Assignments (Page 335)

Felix Brandt (*TU München*)

Johannes Hofbauer (*TU München*)

Martin Suderland (*USI*)

Stable Matching with Uncertain Pairwise Preferences (Page 344)

Haris Aziz (*Data61, CSIRO & UNSW*)

Péter Biró (*Hungarian Academy of Sciences*)

Tamás Fleiner (*Loránd Eötvös University*)

Serge Gaspers (*Data61, CSIRO & UNSW*)

Ronald de Haan (*Technische Universität Wien*)

Nicholas Mattei (*IBM Research*)

Baharak Rastegari (*University of Glasgow*)

Parameterized Complexity of Group Activity Selection (Page 353)

Hooyeon Lee (*Moloco, Inc.*)

Virginia Vassilevska Williams (*Massachusetts Institute of Technology*)

Efficient Near-optimal Algorithms for Barter Exchange (Page 362)

Zhipeng Jia (*Tsinghua University*)

Pingzhong Tang (*Tsinghua University*)

Ruosong Wang (*Tsinghua University*)

Hanrui Zhang (*Tsinghua University*)

Session 2C: Negotiation and Virtual Agents

Automated Negotiations for General Game Playing (Page 371)

Dave de Jonge (*Western Sydney University*)

Dongmo Zhang (*Western Sydney University*)

An Automated Negotiation Agent for Permission Management (Page 380)

Tim Baarslag (*Centrum Wiskunde & Informatica*)

Alper T. Alan (*University of Southampton*)

Richard Gomer (*University of Southampton*)

Muddasser Alam (*University of Oxford*)

Charith Perera (*The Open University*)

Enrico H. Gerding (*University of Southampton*)

m.c. schraefel (*University of Southampton*)

The Value of Information in Automated Negotiation: A Decision Model for Eliciting User Preferences (Page 391)

Tim Baarslag (*Centrum Wiskunde & Informatica*)

Michael Kaisers (*Centrum Wiskunde & Informatica*)

Grumpy & Pinocchio: Answering Human-Agent Negotiation Questions through Realistic Agent Design (Page 401)

Johnathan Mell (*University of Southern California*)

Jonathan Gratch (*University of Southern California*)

Towards An Autonomous Agent that Provides Automated Feedback on Students' Negotiation Skills (Page 410)

Emmanuel Johnson (*University of Southern California*)

Jonathan Gratch (*University of Southern California*)

David DeVault (*University of Southern California*)

Increasing Fairness by Delegating Decisions to Autonomous Agents (Page 419)

Celso M. de Melo (*University of Southern California*)

Stacy Marsella (*Northeastern University*)

Jonathan Gratch (*University of Southern California*)

Session 2D: Agent-Based and Social Simulation

Agent-Based Modeling and Simulation of Mosquito-Borne Disease Transmission (Page 426)

Akshay Jindal (*International Institute of Information Technology, Bangalore*)

Shrisha Rao (*International Institute of Information Technology, Bangalore*)

What Becomes of the Broken Hearted?: An Agent-Based Approach to Self-Evaluation, Interpersonal Loss, and Suicide Ideation (Page 436)

Jiamou Liu (*The University of Auckland*)

Luqi Li (*The University of Auckland*)

Kyle Russell (*AUT University*)

The Importance of Modelling Realistic Human Behaviour When Planning Evacuation Schedules (Page 446)

Chaminda Bulumulla (*MIT University*)

Lin Padgham (*MIT University*)

Dhirendra Singh (*MIT University*)

Jeffrey Chan (*MIT University*)

Learning Conventions via Social Reinforcement Learning in Complex and Open Settings (Page 455)

George A. Vouros (*University of Piraeus*)

Multi-agent Reinforcement Learning in Sequential Social Dilemmas (Page 464)

Joel Z. Leibo (*DeepMind*)

Vinicius Zambaldi (*DeepMind*)

Marc Lanctot (*DeepMind*)

Janusz Marecki (*DeepMind*)

Thore Graepel (*DeepMind*)

Phase Transitions in Possible Dynamics of Cellular and Graph Automata Models of Sparsely Interconnected Multi-Agent Systems (Page 474)

Predrag T. Tošić (*Washington State University*)

Session 2E: Multi-Robot Coordination and Swarm Robotics

A GRASP Metaheuristic for the Coverage of Grid Environments with Limited-Footprint Tools (Page 484)

Alessandro Riva (*Politecnico di Milano*)

Francesco Amigoni (*Politecnico di Milano*)

Decentralised Online Planning for Multi-Robot Warehouse Commissioning (Page 492)

Daniel Claes (*University of Liverpool*)

Frans Oliehoek (*University of Liverpool & University of Amsterdam*)

Hendrik Baier (*University of York*)

Karl Tuyls (*University of Liverpool*)

Multirobot Symbolic Planning under Temporal Uncertainty (Page 501)

Shiqi Zhang (*Cleveland State University & University of Texas at Austin*)

Yuqian Jiang (*University of Texas at Austin*)

Guni Sharon (*University of Texas at Austin*)

Peter Stone (*University of Texas at Austin*)

Exploiting Robotic Swarm Characteristics for Adversarial Subversion in Coverage Tasks (Page 511)

Navyata Sanghvi (*Carnegie Mellon University*)

Sasanka Nagavalli (*Carnegie Mellon University*)

Katia Sycara (*Carnegie Mellon University*)

Three Years of the RoboCup Standard Platform League Drop-In Player Competition: Creating and Maintaining a Large Scale Ad Hoc Teamwork Robotics Competition (JAAMAS Extended Abstract) (Page 520)

Katie Genter (*University of Texas at Austin*)

Tim Laue (*University of Bremen*)

Peter Stone (*University of Texas at Austin*)

Scaling Expectation-Maximization for Inverse Reinforcement Learning to Multiple Robots under Occlusion (Page 522)

Kenneth Bogert (*University of North Carolina at Asheville*)

Prashant Doshi (*University of Georgia*)

Session 2F: Learning

Asynchronous Data Aggregation for Training End to End Visual Control Networks (Page 530)

Mathew Monfort (*Massachusetts Institute of Technology*)
Matthew Johnson (*Microsoft Research, Cambridge*)
Aude Oliva (*Massachusetts Institute of Technology*)
Katja Hofmann (*Microsoft Research, Cambridge*)

Bootstrapping with Models: Confidence Intervals for Off-Policy Evaluation (Page 538)

Josiah P. Hanna (*The University of Texas at Austin*)
Peter Stone (*The University of Texas at Austin*)
Scott Niekum (*The University of Texas at Austin*)

Reasoning about Hypothetical Agent Behaviours and their Parameters (Page 547)

Stefano V. Albrecht (*The University of Texas at Austin*)
Peter Stone (*The University of Texas at Austin*)

Forward Actor-Critic for Nonlinear Function Approximation in Reinforcement Learning (Page 556)

Vivek Veeriah (*University of Alberta*)
Harm van Seijen (*University of Alberta*)
Richard S. Sutton (*University of Alberta*)

Reward Shaping in Episodic Reinforcement Learning (Page 565)

Marek Grzes (*University of Kent*)

Learning to Partition using Score Based Compatibilities (Page 574)

Arun Rajkumar (*Conduent Labs - India*)
Koyel Mukherjee (*IBM Research - India*)
Theja Tulabandhula (*University of Illinois, Chicago*)

Session 3A: Computational Social Choice 2

Manipulation of Hamming-based Approval Voting for Multiple Referenda and Committee Elections (Page 597)

Nathanaël Barrot (*Kyushu University*)
Jérôme Lang (*Dauphine University*)
Makoto Yokoo (*Kyushu University*)

New Approximation for Borda Coalitional Manipulation (Page 606)

Orgad Keller (*Bar-Ilan University*)
Avinatan Hassidim (*Bar-Ilan University*)
Noam Hazon (*Ariel University*)

Complexity of Control by Partition of Voters and of Voter Groups in Veto and Other Scoring Protocols (Page 615)

Cynthia Maushagen (*Heinrich-Heine-Universität Düsseldorf*)
Jörg Rothe (*Heinrich-Heine-Universität Düsseldorf*)

Divide and Conquer: Using Geographic Manipulation to Win District-Based Elections (Page 624)

Yoad Lewenberg (*Hebrew University of Jerusalem*)
Omer Lev (*University of Toronto*)
Jeffrey S. Rosenschein (*Hebrew University of Jerusalem*)

Structured Proportional Representation (Page 633)

Nimrod Talmon (*Weizmann Institute of Science*)

Session 3B: Computational Game Theory & Mechanism Design 2

Referral-Embedded Provision Point Mechanisms for Crowdfunding of Public Projects (Page 642)

Prabhul Chandra (*Indian Institute of Science*)
Sujit Gujar (*International Institute of Information Technology*)
Yadati Narahari (*Indian Institute of Science*)

Spoofing the Limit Order Book: An Agent-Based Model (Page 651)

Xintong Wang (*University of Michigan*)
Michael P. Wellman (*University of Michigan*)

Hotelling-Downs Model with Limited Attraction (Page 660)

Weiran Shen (*Tsinghua University*)
Zihe Wang (*Shanghai University of Finance and Economics*)

Stop Nuclear Smuggling Through Efficient Container Inspection (Page 669)

Xinrun Wang (*Nanyang Technological University*)
Qingyu Guo (*Nanyang Technological University*)
Bo An (*Nanyang Technological University*)

Coordinating Multiple Defensive Resources in Patrolling Games with Alarm Systems (Page 678)

Nicola Basilico (*University of Milan*)

Andrea Celli (*Politecnico di Milano*)
Giuseppe De Nittis (*Politecnico di Milano*)
Nicola Gatti (*Politecnico di Milano*)

Session 3C: Logics 2

A Path in the Jungle of Logics for Multi-agent System: On the Relation between General Game-playing

Logics and Seeing-to-it-that Logics (Page 687)

Emiliano Lorini (*Toulouse University*)
François Schwarzentruber (*ENS Rennes*)

Iterated Games with LDL Goals over Finite Traces (Page 696)

Julian Gutierrez (*University of Oxford*)
Giuseppe Perelli (*University of Oxford*)
Michael Wooldridge (*University of Oxford*)

Iterated Boolean Games for Rational Verification (Page 705)

Tong Gao (*University of Oxford*)
Julian Gutierrez (*University of Oxford*)
Michael Wooldridge (*University of Oxford*)

Reasoning about Natural Strategic Ability (Page 714)

Wojciech Jamroga (*Polish Academy of Sciences*)
Vadim Malvone (*Università degli studi di Napoli Federico II*)
Aniello Murano (*Università degli studi di Napoli Federico II*)

Coalition Power in Epistemic Transition Systems (Page 723)

Pavel Naumov (*Vassar College*)
Jia Tao (*College of New Jersey*)

Session 3D: Planning

Exploiting Anonymity and Homogeneity in Factored Dec-MDPs through Precomputed Binomial

Distributions (Page 732)

Rajiv Ranjan Kumar (*Singapore Management University*)
Pradeep Varakantham (*Singapore Management University*)

GUBS: a Utility-Based Semantic for Goal-Directed Markov Decision Processes (Page 741)

Valdinei Freire (*Universidade de São Paulo*)
Karina Valdivia Delgado (*Universidade de São Paulo*)

Cost-Based Goal Recognition for Path-Planning (Page 750)

Peta Masters (*RMIT University*)
Sebastian Sardina (*RMIT University*)

BDI Agent Reasoning with Guidance from HTN Recipes (Page 759)

Lavindra de Silva (*University of Nottingham*)

Progressing Intention Progression: A Call for a Goal-Plan Tree Contest (Page 768)

Brian Logan (*University of Nottingham*)
John Thangarajah (*RMIT University*)
Neil Yorke-Smith (*American University of Beirut*)

Session 3E: Virtual Agents

Using Virtual Narratives to Explore Children's Story Understanding (Page 773)

Julie Porteous (*Teesside University*)
Fred Charles (*Bournemouth University*)
Cameron Smith (*Teesside University*)
Marc Cavazza (*University of Kent*)
Jolien Mouw (*Leiden University*)
Paul van den Broek (*Leiden University*)

MISER: Mise-En-Scène Region Support for Staging Narrative Actions in Interactive Storytelling (Page 782)

Jamie Matthews (*Teesside University*)
Fred Charles (*Bournemouth University*)
Julie Porteous (*Teesside University*)
Alexandra Mendes (*Teesside University*)

Pedagogical Agents as Team Members: Impact of Proactive and Pedagogical Behavior on the User (Page 791)

Mukesh Barange (*INSA de Rouen*)
Julien Saunier (*INSA de Rouen*)
Alexandre Pauchet (*INSA de Rouen*)

Incorporating Emotion Perception into Opponent Modeling for Social Dilemmas (Page 801)

Rens Hoegen (*University of Southern California*)
Giota Stratou (*University of Southern California*)
Jonathan Gratch (*University of Southern California*)

"Is It Just Me?" : Evaluating Attribution of Negative Feedback as a Function of Virtual Instructor's Gender and Proxemics (Page 810)

Dan Feng (*Northeastern University*)
David C. Jeong (*University of Southern California*)
Nicole C. Krämer (*University Duisburg-Essen*)
Lynn C. Miller (*University of Southern California*)
Stacy Marsella (*Northeastern University*)

Session 3F: Transportation and Routing

Data-Driven Simulation and Optimization for Incident Response in Urban Railway Networks (Page 819)

Jácint Szabó (*IBM Research Lab, Zurich*)
Sebastien Blandin (*IBM Research, Singapore*)
Charles Brett (*IBM Research, Singapore*)

Real-time Adaptive Tolling Scheme for Optimized Social Welfare in Traffic Networks (Page 828)

Guni Sharon (*University of Texas at Austin*)
Josiah P. Hanna (*University of Texas at Austin*)
Tarun Rambha (*Cornell University*)
Michael W. Levin (*University of Texas at Austin*)
Michael Albert (*Duke University*)
Stephen D. Boyles (*University of Texas at Austin*)
Peter Stone (*University of Texas at Austin*)

Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks (Page 837)

Hang Ma (*University of Southern California*)
Jiaoyang Li (*Tsinghua University*)
T. K. Satish Kumar (*University of Southern California*)
Sven Koenig (*University of Southern California*)

Learning to Minimise Regret in Route Choice (Page 846)

Gabriel de O. Ramos (*Universidade Federal do Rio Grande do Sul*)
Bruno C. da Silva (*Universidade Federal do Rio Grande do Sul*)
Ana L. C. Bazzan (*Universidade Federal do Rio Grande do Sul*)

Joint Movement of Pebbles in Solving the (N^2-1)-Puzzle and its Applications in Cooperative Path-Finding: (JAAMAS Extended Abstract) (Page 856)

Pavel Surynek (*National Institute of Advanced Industrial Science and Technology (AIST)*)
Petr Michalík (*Accenture Central Europe B.V.Consulting*)

Session 4A: Computational Social Choice 3

Proxy Voting for Better Outcomes (Page 858)

Gal Cohensius (*Technion-Israel Institute of Technology*)
Shie Mannor (*Technion-Israel Institute of Technology*)
Reshef Meir (*Technion-Israel Institute of Technology*)
Eli Meir (*Technion-Israel Institute of Technology*)
Ariel Orda (*Technion-Israel Institute of Technology*)

Real Candidacy Games: A New Model for Strategic Candidacy (Page 867)

Itay Sabato (*Hebrew University of Jerusalem*)
Svetlana Obraztsova (*Nanyang Technological University*)
Zinovi Rabinovich (*Nanyang Technological University*)
Jeffrey S. Rosenschein (*Hebrew University of Jerusalem*)

Doodle Poll Games (Page 876)

Svetlana Obraztsova (*Nanyang Technological University*)
Maria Polukarov (*University of Southampton*)
Zinovi Rabinovich (*Nanyang Technological University*)
Edith Elkind (*University of Oxford*)

Distant Truth: Bias Under Vote Distortion Costs (Page 885)

Svetlana Obraztsova (*Nanyang Technological University*)
Omer Lev (*University of Toronto*)
Evangelos Markakis (*Athens University of Economics and Business*)
Zinovi Rabinovich (*Nanyang Technological University*)
Jeffrey S. Rosenschein (*Hebrew University of Jerusalem*)

A Restricted Markov Tree Model for Inference and Generation in Social Choice with Incomplete Preferences (Page 893)

John A. Doucette (*New College of Florida*)
Robin Cohen (*University of Waterloo*)

Session 4B: Computational Game Theory & Mechanism Design 3

Combining Incremental Strategy Generation and Branch and Bound Search for Computing Maxmin Strategies

in Imperfect Recall Games (Page 902)

Jiri Cermák (*Czech Technical University in Prague*)
Branislav Bošanský (*Czech Technical University in Prague*)
Michal Pechoucek (*Czech Technical University in Prague*)

Computing Approximate Pure Nash Equilibria in Digraph k-Coloring Games (Page 911)

Raffaello Carosi (*Gran Sasso Science Institute*)
Michele Flammini (*Gran Sasso Science Institute & University of L'Aquila*)
Gianpiero Monaco (*University of L'Aquila*)

Detecting Switches Against Non-Stationary Opponents (JAAMAS Extended Abstract) (Page 920)

Pablo Hernandez-Leal (*Centrum Wiskunde & Informatica*)
Yusen Zhan (*Washington State University*)
Matthew E. Taylor (*Washington State University*)
L. Enrique Sucar (*INAOE*)
Enrique Munoz de Cote (*INAOE & PROWLER.io Ltd.*)

An Exploration Strategy Facing Non-Stationary Agents (JAAMAS Extended Abstract) (Page 922)

Pablo Hernandez-Leal (*Centrum Wiskunde & Informatica*)
Yusen Zhan (*Washington State University*)
Matthew E. Taylor (*Washington State University*)
L. Enrique Sucar (*INAOE*)
Enrique Munoz de Cote (*INAOE & PROWLER.io Ltd.*)

Safely Using Predictions in General-Sum Normal Form Games (Page 924)

Steven Damer (*University of Minnesota*)
Maria Gini (*University of Minnesota*)

Session 4C: Judgment Aggregation and Argumentation

Planning for Persuasion (Page 933)

Elizabeth Black (*King's College London*)
Amanda J. Coles (*King's College London*)
Christopher Hampson (*King's College London*)

Are Ranking Semantics Sensitive to the Notion of Core? (Page 943)

Bruno Yun (*University of Montpellier*)
Madalina Croitoru (*University of Montpellier*)
Pierre Bisquert (*INRA Montpellier*)

Complexity Results for Aggregating Judgments using Scoring or Distance-Based Procedures (Page 952)

Ronald de Haan (*Technische Universität Wien*)
Marija Slavkovic (*University of Bergen*)

Reinforcement Learning for Multi-Step Expert Advice (Page 962)

Patrick Philipp (*Karlsruhe Institute of Technology*)
Achim Rettinger (*Karlsruhe Institute of Technology*)

Applying Copeland Voting to Design an Agent-Based Hyper-Heuristic (Page 972)

Vinicius Renan de Carvalho (*University of São Paulo (USP)*)
Jaime Simão Sichman (*University of São Paulo (USP)*)

Session 4D: Novel Applications in Smart Grids and Mobility

A Multiagent System Approach to Scheduling Devices in Smart Homes (Page 981)

Ferdinando Fioretto (*University of Michigan*)
William Yeoh (*New Mexico State University*)
Enrico Pontelli (*New Mexico State University*)

Label Correction and Event Detection for Electricity Disaggregation (Page 990)

Mark Valovage (*University of Minnesota*)
Maria Gini (*University of Minnesota*)

A Distributed Constraint Optimization (DCOP) Approach to the Economic Dispatch with Demand Response (Page 999)

Ferdinando Fioretto (*University of Michigan*)
William Yeoh (*New Mexico State University*)
Enrico Pontelli (*New Mexico State University*)
Ye Ma (*Siemens Industry Inc.*)
Satishkumar J. Ranade (*New Mexico State University*)

Save Money or Feel Cozy? A Field Experiment Evaluation of a Smart Thermostat that Learns Heating Preferences (Page 1008)

Mike Shann (*University of Zurich*)
Alper Alan (*University of Southampton*)
Sven Seuken (*University of Zurich*)
Enrico Costanza (*University College London*)

Sarvapali D. Ramchurn (*University of Southampton*)

Adaptive Pricing Mechanisms for On-Demand Mobility (Page 1017)

Maciej Drwal (*University of Southampton*)

Enrico Gerding (*University of Southampton*)

Sebastian Stein (*University of Southampton*)

Keiichiro Hayakawa (*Toyota Central R&D Labs., Inc.*)

Hironobu Kitaoka (*Toyota Motor Corporation*)

Session 4E: Networking and Communication

Uniform Information Exchange in Multi-channel Wireless Ad Hoc Networks (Page 1026)

Dongxiao Yu (*Huazhong University of Science and Technology*)

Li Ning (*Chinese Academy of Sciences*)

Yong Zhang (*Chinese Academy of Sciences*)

Hai Jin (*Huazhong University of Science and Technology*)

Yuxuan Wang (*University of Hong Kong*)

Francis C.M. Lau (*University of Hong Kong*)

Shengzhong Feng (*Shenzhen Institutes of Advanced Technology, CAS*)

Multi-Agent Nonlinear Negotiation for Wi-Fi Channel Assignment (Page 1035)

Enrique de la Hoz (*University of Alcala*)

Ivan Marsa-Maestre (*University of Alcala*)

Jose Manuel Gimenez-Guzman (*University of Alcala*)

David Orden (*University of Alcala*)

Mark Klein (*Massachusetts Institute of Technology*)

A Distributed, Multi-Agent Approach to Reactive Network Resilience (Page 1044)

Enrique de la Hoz (*University of Alcala*)

Jose Manuel Gimenez-Guzman (*University of Alcala*)

Ivan Marsa-Maestre (*University of Alcala*)

Luis Cruz-Piris (*University of Alcala*)

David Orden (*University of Alcala*)

Splee: A Declarative Information-Based Language for Multiagent Interaction Protocols (Page 1054)

Amit K. Chopra (*Lancaster University*)

Samuel H. Christie V. (*North Carolina State University*)

Munindar P. Singh (*North Carolina State University*)

Vocabulary Alignment in Openly Specified Interactions (Page 1064)

Paula Chocron (*Universitat Autònoma de Barcelona*)

Marco Schorlemmer (*Universitat Autònoma de Barcelona*)

Session 4F: Teamwork, Coordination, and Cooperation

Theoretical Foundations of Team Matchmaking (Page 1073)

Josh Alman (*Massachusetts Institute of Technology*)

Dylan McKay (*Massachusetts Institute of Technology*)

Incentivizing Cooperation between Heterogeneous Agents in Dynamic Task Allocation (Page 1082)

Sofia Amador Nelke (*Ben-Gurion University of the Negev*)

Roie Zivan (*Ben-Gurion University of the Negev*)

Causality, Responsibility and Blame in Team Plans (Page 1091)

Natasha Alechina (*University of Nottingham*)

Joseph Y. Halpern (*Cornell University*)

Brian Logan (*University of Nottingham*)

Simultaneously Learning and Advising in Multiagent Reinforcement Learning (Page 1100)

Felipe Leno da Silva (*Escola Politécnica of the University of São Paulo*)

Ruben Glatt (*Escola Politécnica of the University of São Paulo*)

Anna Helena Reali Costa (*Escola Politécnica of the University of São Paulo*)

Cooperative Set Function Optimization Without Communication or Coordination (Page 1109)

Gustavo Malkomes (*Washington University in St. Louis*)

Kefu Lu (*Washington University in St. Louis*)

Blakeley Hoffman (*University of South Carolina*)

Roman Garnett (*Washington University in St. Louis*)

Benjamin Moseley (*Washington University in St. Louis*)

Richard Mann (*University of Leeds*)

Session 5A: Computational Social Choice 4

Anyone But Them: The Complexity Challenge for A Resolute Election Controller (Page 1133)

Yongjie Yang (*Central South University & Saarland University*)

Jianxin Wang (*Central South University*)

The Complexity of Bribery and Control in Group Identification (Page 1142)

Gábor Erdélyi (*University of Siegen*)
Christian Reger (*University of Siegen*)
Yongjie Yang (*Saarland University*)

Complexity Results for Manipulation, Bribery and Control of the Kemeny Judgment Aggregation

Procedure (Page 1151)

Ronald de Haan (*Technische Universität Wien*)

Distributed Monitoring of Election Winners (Page 1160)

Arnold Filtser (*Ben-Gurion University*)
Nimrod Talmon (*Weizmann Institute of Science*)

The Complexity of Control and Bribery in Majority Judgment (Page 1169)

Yongjie Yang (*Central South University & Saarland University*)

On the Complexity of Borda Control in Single-Peaked Elections (Page 1178)

Yongjie Yang (*Central South University & Saarland University*)

Session 5B: Logic & Game Theory

Analyzing Games with Ambiguous Player Types Using the MINthenMAX Decision Model (Page 1187)

Ilan Nehama (*Hebrew University of Jerusalem & Tel-Aviv University*)

Strategic Disclosure of Opinions on a Social Network (Page 1196)

Umberto Grandi (*University of Toulouse*)
Emiliano Lorini (*University of Toulouse*)
Arianna Novaro (*University of Toulouse*)
Laurent Perrussel (*University of Toulouse*)

Hiding Actions in Multi-Player Games (Page 1205)

Vadim Malvone (*Università degli Studi di Napoli Federico II*)
Aniello Murano (*Università degli Studi di Napoli Federico II*)
Loredana Sorrentino (*Università degli Studi di Napoli Federico II*)

Synthesizing Optimal Social Laws for Strategic Agents via Bayesian Mechanism Design (Page 1214)

Jun Wu (*Nanjing University*)
Lei Zhang (*Nanjing University*)
Chongjun Wang (*Nanjing University*)
Junyuan Xie (*Nanjing University*)

Attenuate Locally, Win Globally: An Attenuation-based Framework for Online Stochastic Matching with

Timeouts (Page 1223)

Brian Brubach (*University of Maryland*)
Karthik Abinav Sankararaman (*University of Maryland*)
Aravind Srinivasan (*University of Maryland*)
Pan Xu (*University of Maryland*)

Multi-Channel Marketing with Budget Complementarities (Page 1232)

Haifeng Zhang (*Vanderbilt University*)
Yevgeniy Vorobeychik (*Vanderbilt University*)
Ariel D. Procaccia (*Carnegie Mellon University*)

Session 5C: Logics 3

Fixpoint Approximation of Strategic Abilities under Imperfect Information (Page 1241)

Wojciech Jamroga (*Polish Academy of Sciences*)
Michał Knapik (*Polish Academy of Sciences*)
Damian Kurpiewski (*Polish Academy of Sciences*)

Decidability Results for ATL* with Imperfect Information and Perfect Recall (Page 1250)

Raphaël Berthon (*ENS Rennes*)
Bastien Maubert (*University of Naples Federico II*)
Aniello Murano (*University of Naples Federico II*)

Agent-based Abstractions for Verifying Alternating-time Temporal Logic with Imperfect Information (Page 1259)

Francesco Belardinelli (*UEVE & IRIT Toulouse*)
Alessio Lomuscio (*Imperial College London*)

Verification of Multi-agent Systems with Imperfect Information and Public Actions (Page 1268)

Francesco Belardinelli (*Université d'Evry & IRIT Toulouse*)
Alessio Lomuscio (*Imperial College London*)
Aniello Murano (*Università degli Studi di Napoli*)
Sasha Rubin (*Università degli Studi di Napoli*)

Game-Theoretic Semantics for ATL+ with Applications to Model Checking (Page 1277)

Valentin Goranko (*Stockholm University & University of Johannesburg*)
Antti Kuusisto (*University of Bremen*)

Raine Rönholm (*University of Tampere*)

Bisimulations for Verifying Strategic Abilities with an Application to ThreeBallot (Page 1286)

Francesco Belardinelli (*Université d'Evry & IRIT Toulouse*)

Rodica Condurache (*Université Paris-Est Créteil*)

Catalin Dima (*Université Paris-Est Créteil*)

Wojciech Jamroga (*Polish Academy of Sciences*)

Andrew V. Jones (*Vector Software Inc.*)

Session 5D: Social Networks

Group Reasoning in Social Environments (Page 1296)

Erman Acar (*University of Mannheim*)

Gianluigi Greco (*University of Calabria*)

Marco Manna (*University of Calabria*)

Uncharted but not Uninfluenced: Influence Maximization with an Uncertain Network (Page 1305)

Bryan Wilder (*University of Southern California*)

Amulya Yadav (*University of Southern California*)

Nicole Immorlica (*Microsoft Research*)

Eric Rice (*University of Southern California*)

Milind Tambe (*University of Southern California*)

Robustness in Discrete Preference Games (Page 1314)

Vincenzo Auletta (*University of Salerno*)

Ioannis Caragiannis (*University of Patras & CTI "Diophantus"*)

Diodato Ferraioli (*University of Salerno*)

Clemente Galdi (*University of Napoli "Federico II"*)

Giuseppe Persiano (*University of Salerno*)

Contrasting the Spread of Misinformation in Online Social Networks (Page 1323)

Marco Amoruso (*University of Salerno*)

Daniele Anello (*University of Salerno*)

Vincenzo Auletta (*University of Salerno*)

Diodato Ferraioli (*University of Salerno*)

Limiting Concept Spread in Environments with Interacting Concepts (Page 1332)

James Archbold (*University of Warwick*)

Nathan Griffiths (*University of Warwick*)

On the Construction of Covert Networks (Page 1341)

Marcin Waniek (*University of Warsaw*)

Tomasz P. Michalak (*University of Oxford & University of Warsaw*)

Talal Rahwan (*Masdar Institute of Science and Technology*)

Michael Wooldridge (*University of Oxford*)

Session 5E: Human-Robot Interaction and Learning in Robotics

Temporal Models for Robot Classification of Human Interruptibility (Page 1350)

Siddhartha Banerjee (*Georgia Institute of Technology*)

Sonia Chernova (*Georgia Institute of Technology*)

A Tale of Two Architectures: A Dual-Citizenship Integration of Natural Language and the Cognitive Map (Page 1360)

Tom Williams (*Tufts University*)

Collin Johnson (*University of Michigan*)

Matthias Scheutz (*Tufts University*)

Benjamin Kuipers (*University of Michigan*)

Multi-Robot Human Guidance: Human Experiments and Multiple Concurrent Requests (Page 1369)

Piyush Khandelwal (*Cogitai Inc. & University of Texas at Austin*)

Peter Stone (*University of Texas at Austin*)

Spoken Instruction-Based One-Shot Object and Action Learning in a Cognitive Robotic Architecture (Page 1378)

Matthias Scheutz (*Tufts University*)

Evan Krause (*Tufts University*)

Brad Oosterveld (*Tufts University*)

Tyler Frasca (*Tufts University*)

Robert Platt (*Northeastern University*)

Effect of Leader Placement on Robotic Swarm Control (Page 1387)

Rohan Tiwari (*IIT-Delhi*)

Puneet Jain (*IIT-Delhi*)

Sachit Butail (*Northern Illinois University*)

Sujit P. Baliyarisimhuni (*IIT-Delhi*)

Michael A. Goodrich (*Brigham Young University*)

Probabilistic Supervisory Control Theory (pSCT) Applied to Swarm Robotics (Page 1395)

Yuri Kaszubowski Lopes (*University of Sheffield*)
Stefan M. Trenkwalder (*University of Sheffield*)
André B. Leal (*Santa Catarina State University*)
Tony J. Dodd (*University of Sheffield*)
Roderich Groß (*University of Sheffield*)

Session 5F: Emergence

Error Cascades in Collective Behavior: A Case Study of the Gradient Algorithm on 1000 Physical Agents (Page 1404)

Melvin Gauci (*Harvard University*)
Monica E. Ortiz (*Harvard University*)
Michael Rubenstein (*Northwestern University*)
Radhika Nagpal (*Harvard University*)

Inverse Reinforcement Learning in Swarm Systems (Page 1413)

Adrian Šošić (*Technische Universität Darmstadt*)
Wasiur R. KhudaBukhsh (*Technische Universität Darmstadt*)
Abdelhak M. Zoubir (*Technische Universität Darmstadt*)
Heinz Koepl (*Technische Universität Darmstadt*)

Social Manifestation of Guilt Leads to Stable Cooperation in Multi-Agent Systems (Page 1422)

Luis Moniz Pereira (*Universidade Nova de Lisboa*)
Tom Lenaerts (*Université Libre de Bruxelles & Vrije Universiteit Brussel*)
Luis A. Martinez-Vaquero (*National Research Council of Italy*)
The Anh Han (*Teesside University*)

Commitment and Participation in Public Goods Games (JAAMAS Extended Abstract) (Page 1431)

The Anh Han (*Teesside University*)
Luis Moniz Pereira (*Universidade Nova de Lisboa*)
Tom Lenaerts (*Université Libre de Bruxelles &*)

Understanding Norm Change: An Evolutionary Game-Theoretic Approach (Page 1433)

Soham De (*University of Maryland*)
Dana S. Nau (*University of Maryland*)
Michele J. Gelfand (*University of Maryland*)

Multi-Agent Flag Coordination Games (Page 1442)

David Kohan Marzagão (*King's College London*)
Nicolás Rivera (*King's College London*)
Colin Cooper (*King's College London*)
Peter McBurney (*King's College London*)
Kathleen Steinhöfel (*King's College London*)

Extended Abstracts

A Deep Learning Approach for Norm Conflict Identification (Page 1451)

João Paulo Aires (*Pontifical Catholic University of Rio Grande do Sul*)
Felipe Meneguzzi (*Pontifical Catholic University of Rio Grande do Sul*)

An Agent-Based Self-Organizing Traffic Environment for Urban Evacuations (Page 1454)

Mohammad Al-Zinati (*Jordan University of Science and Technology*)
Rym Z. Wenkstern (*University of Texas at Dallas*)

Parametric Runtime Verification of Multiagent Systems (Page 1457)

Davide Ancona (*University of Genova*)
Angelo Ferrando (*University of Genova*)
Viviana Mascardi (*University of Genova*)

Timed ATL: Forget Memory, Just Count (Page 1460)

Étienne André (*Université Paris 13*)
Laure Petrucci (*Université Paris 13*)
Wojciech Jamroga (*Polish Academy of Sciences*)
Michał Knapik (*Polish Academy of Sciences*)
Wojciech Penczek (*Polish Academy of Sciences & University of Natural Sciences and Humanities*)

Synergistic Team Composition (Page 1463)

Ewa Andrejczuk (*IIIA-CSIC*)
Juan A. Rodríguez-Aguilar (*IIIA-CSIC*)
Carme Roig (*Institut Torras i Bages*)
Carles Sierra (*IIIA-CSIC*)

On Approximate Welfare- and Revenue-Maximizing Equilibria for Size-Interchangeable Bidders (Page 1466)

Enrique Areyan Viqueira (*Brown University*)
Amy Greenwald (*Brown University*)
Victor Naroditskiy (*OneMarketData*)

Coalition Formability Semantics with Conflict-Eliminable Sets of Arguments (Page 1469)

Ryuta Arisaka (*National Institute of Informatics*)
Ken Satoh (*National Institute of Informatics*)

Pareto Optimal Allocation under Uncertain Preferences (Page 1472)

Haris Aziz (*Data61, CSIRO and UNSW Australia*)
Ronald de Haan (*Technische Universität Wien*)
Baharak Rastegari (*University of Glasgow*)

Coalitional Exchange Stable Matchings in Marriage and Roommate Markets (Page 1475)

Haris Aziz (*Data61, CSIRO and UNSW Australia*)
Adrian Goldwaser (*UNSW Australia*)

Large-Scale Complex Adaptive Systems using Multi-Agent Modeling and Simulation (Page 1478)

Lachlan Birdsey (*University of Adelaide*)
Claudia Szabo (*University of Adelaide*)
Katrina Falkner (*University of Adelaide*)

Profit Optimization in Commercial Ridesharing (Page 1481)

Arpita Biswas (*Indian Institute of Science*)
Ragavendran Gopalakrishnan (*Conduent Labs India*)
Theja Tulabandhula (*University of Illinois, Chicago*)
Koyel Mukherjee (*IBM Research-India*)
Asmita Metrewar (*Apple India*)
Raja Subramaniam Thangaraj (*Conduent Labs India*)

Gesture-Based Control of Autonomous UAVs (Page 1484)

Jonathon Bolin (*University of Tulsa*)
Chad Crawford (*University of Tulsa*)
William Macke (*University of Tulsa*)
Jon Hoffman (*University of Tulsa*)
Sam Beckmann (*University of Tulsa*)
Sandip Sen (*University of Tulsa*)

A Modular Framework for Decentralised Multi-Agent Planning (Page 1487)

Rafael C. Cardoso (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)
Rafael H. Bordini (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)

Affect of Robot's Competencies on Children's Perception (Page 1490)

Shruti Chandra (*École Polytechnique Fédérale de Lausanne & University of Lisboa*)
Raul Paradedá (*University of Lisbon & State University of Rio Grande do Norte*)
Hang Yin (*École Polytechnique Fédérale de Lausanne & University of Lisboa*)
Pierre Dillenbourg (*École Polytechnique Fédérale de Lausanne*)
Rui Prada (*University of Lisbon*)
Ana Paiva (*University of Lisbon*)

K-Memory Strategies in Repeated Games (Page 1493)

Lijie Chen (*Tsinghua University*)
Fangzhen Lin (*HKUST*)
Pingzhong Tang (*Tsinghua University*)
Kangning Wang (*Tsinghua University*)
Ruosong Wang (*Tsinghua University*)
Shiheng Wang (*HKUST*)

WITHDRAWN Post Processing & Paginating (4/6, LT) Evolutionary Game-theoretic Modeling of Past Societies'

Organization Structure (Page 1496)

Angelos Chliaoutakis (*Technical University of Crete*)
Georgios Chalkiadakis (*Technical University of Crete*)

A Comparison of Targeted Layered Containment Strategies for a Flu Pandemic in Three US Cities (Page 1499)

Shuyu Chu (*Virginia Tech*)
Samarth Swarup (*Virginia Tech*)
Jiangzhuo Chen (*Virginia Tech*)
Achla Marathe (*Virginia Tech*)

Give Agents Some REST: A Resource-oriented Abstraction Layer for Internet-scale Agent Environments (Page 1502)

Andrei Ciortea (*Université Lyon, MINES Saint-Étienne*)
Olivier Boissier (*Université Lyon, MINES Saint-Étienne*)
Antoine Zimmermann (*Université Lyon, MINES Saint-Étienne*)
Adina Magda Florea (*University Politehnica of Bucharest*)

Max-sum Revisited; The Real Power of Damping (Page 1505)

Liel Cohen (*Ben-Gurion University*)
Roie Zivan (*Ben-Gurion University*)

A Practical Resource-Constrained Norm Monitor (Page 1508)

Natalia Criado (*King's College London*)

Multi-player Approximate Nash Equilibria (Page 1511)

Artur Czumaj (*University of Warwick*)
Michail Fasoulakis (*University of Warwick and ICS-FORTH*)
Marcin Jurdzinski (*University of Warwick*)

Zero-Sum Game Techniques for Approximate Nash Equilibria (Page 1514)

Artur Czumaj (*University of Warwick*)
Michail Fasoulakis (*University of Warwick and ICS-FORTH*)
Marcin Jurdzinski (*University of Warwick*)

Reusing Skills for First-Time Solution of Navigation Tasks in Platform Videogames (Page 1517)

Michael Dann (*RMIT University*)
Fabio Zambetta (*RMIT University*)
John Thangarajah (*RMIT University*)

Other-Condemning Anger = Blaming Accountable Agents for Unattainable Desires (Page 1520)

Mehdi Dastani (*Utrecht University*)
Emiliano Lorini (*IRIT-CNRS, Toulouse University*)
John-Jules Meyer (*Utrecht University*)
Alexander Pankov (*Utrecht University*)

Memory-Based Mechanisms for Economic Agents (Page 1523)

Gil Dollberg (*Hebrew University of Jerusalem*)
Aviv Zohar (*Hebrew University of Jerusalem*)

Modelling and Reasoning about Remediation Actions in BDI Agents (Page 1526)

João Faccin (*Universidade Federal do Rio Grande do Sul*)
Ingrid Nunes (*Universidade Federal do Rio Grande do Sul & TU Dortmund*)

Doxastic Reasoning with Multi-Source Justifications based on Second Order Propositional Modal Logic (Page 1529)

Tuan-Fang Fan (*National Penghu University of Science and Technology*)
Churn-Jung Liau (*Academia Sinica*)

Two Forms of Explanations in Computational Assumption-based Argumentation (Page 1532)

Xiuyi Fan (*Nanyang Technological University*)
Siyuan Liu (*Nanyang Technological University*)
Huiquo Zhang (*Nanyang Technological University*)
Chunyan Miao (*Nanyang Technological University*)
Cyril Leung (*University of British Columbia*)

Fair Allocation of Indivisible Goods with Different Entitlements (Page 1535)

Alireza Farhadi (*University of Maryland*)
MohammadTaghi Hajiaghayi (*University of Maryland*)
Mohammad Ghodsi (*Sharif University of Technology*)
Sebastien Lahaie (*Google Research*)
David Pennock (*Microsoft Research*)
Masoud Seddighin (*Sharif University of Technology*)
Saeed Seddighin (*University of Maryland*)
Hadi Yami (*University of Maryland*)

Flexible POMDP Framework for Human-Robot Cooperation in Escort Tasks (Page 1538)

Fabio-Valerio Ferrari (*University of Caen Normandy*)
Laurent Jeanpierre (*University of Caen Normandy*)
Abdel-illah Mouaddib (*University of Caen Normandy*)

Finding a Needle in a Haystack: Satellite Detection of Moving Objects in Marine Environments (Page 1541)

Natalie Fridman (*ImageSat International (ISI)*)
Doron Amir (*ImageSat International (ISI)*)
Ilan Schwartzman (*ImageSat International (ISI)*)
Oded Stawitzky (*ImageSat International (ISI)*)
Igor Kleinerman (*ImageSat International (ISI)*)
Sharon Kligsberg (*ImageSat International (ISI)*)
Noa Agmon (*Bar-Ilan University*)

Context-Based Concurrent Experience Sharing in Multiagent Systems (Page 1544)

Dan Garant (*University of Massachusetts*)
Bruno C. da Silva (*Federal University of Rio Grande do Sul*)
Victor Lesser (*University of Massachusetts*)
Chongjie Zhang (*Tsinghua University*)

Learning to Assemble Objects with a Robot Swarm (Page 1547)

Gregor H.W. Gebhardt (*Technische Universität Darmstadt*)
Kevin Daun (*Technische Universität Darmstadt*)
Marius Schnaubelt (*Technische Universität Darmstadt*)
Alexander Hendrich (*Technische Universität Darmstadt*)
Daniel Kauth (*Technische Universität Darmstadt*)
Gerhard Neumann (*University of Lincoln*)

interActors: A Model for Separating Complex Communication Concerns in Multiagent Computations (Page 1550)

Hongxing Geng (*University of Saskatchewan*)
Nadeem Jamali (*University of Saskatchewan*)

Agent Behaviors for Joining and Leaving a Flock (Page 1553)

Katie Genter (*University of Texas at Austin*)
Peter Stone (*University of Texas at Austin*)

DYCOM: A Dynamic Truthful Budget Balanced Double-sided Combinatorial Market (Page 1556)

Rica Gonen (*Open University of Israel*)
Ozi Egri (*Open University of Israel*)

A Flexible Approach for Designing Optimal Reward Functions (Page 1559)

Ricardo Grunitzki (*Universidade Federal do Rio Grande do Sul*)
Bruno C. da Silva (*Universidade Federal do Rio Grande do Sul*)
Ana L. C. Bazzan (*Universidade Federal do Rio Grande do Sul*)

Prom Week Meets Skyrim: Developing a Social Agent Architecture in a Commercial Game (Page 1562)

Manuel Guimarães (*Universidade de Lisboa*)
Pedro Santos (*Universidade de Lisboa*)
Arnav Jhala (*North Carolina State University*)

Adapting with Honeypot Configurations to Detect Evolving Exploits (Page 1565)

Marcus Gutierrez (*University of Texas at El Paso*)
Christopher Kiekintveld (*University of Texas at El Paso*)

Argumentation-Based Defeasible Reasoning For Existential Rules (Page 1568)

Abdelraouf Hecham (*University of Montpellier*)
Madalina Croitoru (*University of Montpellier*)
Pierre Bisquert (*IATE, INRA Montpellier*)

Local Norm Phenomena in Multi-Agent Systems under Community Networks (Page 1570)

Shuyue Hu (*Chinese University of Hong Kong*)
Ho-fung Leung (*Chinese University of Hong Kong*)

Evolving Collective Driving Behaviors (Page 1573)

Chien-Lun (Allen) Huang (*University of Cape Town*)
Geoff Nitschke (*University of Cape Town*)

On Parameterized Complexity of Group Activity Selection Problems on Social Networks (Page 1575)

Ayumi Igarashi (*University of Oxford*)
Robert Bredereck (*University of Oxford*)
Edith Elkind (*University of Oxford*)

Working Together: Committee Selection and the Supermodular Degree (Page 1578)

Rani Izsak (*Weizmann Institute of Science*)

Safety First: Strategic Navigation in Adversarial Environments (Page 1581)

Ofri Keidar (*Bar Ilan University*)
Noa Agmon (*Bar Ilan University*)

Refinements and Randomised Versions of Some Tournament Solutions (Page 1584)

Justin Kruger (*Universite Paris-Dauphine & PSL Research University*)
Stéphane Airiau (*Universite Paris-Dauphine & PSL Research University*)

Mechanism Design for Ontology Alignment (Page 1587)

Piotr Krysta (*University of Liverpool*)
Minming Li (*City University of Hong Kong*)
Terry R. Payne (*University of Liverpool*)
Nan Zhi (*University of Liverpool*)

Generalised Discount Functions applied to a Monte-Carlo AI/mu Implementation (Page 1589)

Sean Lamont (*Australian National University*)
John Aslanides (*Australian National University*)
Jan Leike (*Oxford University & Google Deepmind*)
Marcus Hutter (*Australian National University*)

Agent-based Influence Maintenance in Social Networks (Page 1592)

Weihua Li (*Auckland University of Technology*)
Quan Bai (*Auckland University of Technology*)
Tung Doan Nguyen (*Auckland University of Technology*)
Minjie Zhang (*University of Wollongong*)

Protecting Elections with Minimal Resource Consumption (Page 1595)

Yunpeng Li (*Southeast University*)
Yichuan Jiang (*Southeast University*)
Weiwei Wu (*Southeast University*)

Dynamic Generalization Kanerva Coding in Reinforcement Learning for TCP Congestion Control Design (Page 1598)

Wei Li (*Northeastern University*)
Fan Zhou (*Northeastern University*)
Waleed Meleis (*Northeastern University*)
Kaushik Chowdhury (*Northeastern University*)

Autonomous Model Management via Reinforcement Learning - Extended Abstract (Page 1601)

Elad Liebman (*University of Texas at Austin*)
Eric Zavesky (*AT&T Research*)
Peter Stone (*University of Texas at Austin*)

Disparity-Aware Group Formation for Recommendation (Page 1604)

Lin Xiao (*Tsinghua University*)
Zhang Min (*Tsinghua University*)
Zhang Yongfeng (*University of Massachusetts, Amherst*)
Gu Zhaoquan (*Hong Kong University*)

Multi-armed Bandit Mechanism with Private Histories (Page 1607)

Chang Liu (*Alibaba Group*)
Qingpeng Cai (*Tsinghua University*)
Yukui Zhang (*Alibaba Group*)

Making and Improving Predictions of Interest Using an MDP Model (Page 1610)

Yunlong Liu (*Xiamen University*)
Yifeng Zeng (*Xiamen University & Teesside University*)
Hexing Zhu (*Xiamen University*)
Yun Tang (*Xiamen University*)

Automating Decision Making to Help Establish Norm-Based Regulations (Page 1613)

Maitte Lopez-Sanchez (*Universitat de Barcelona*)
Marc Serramia (*Universitat de Barcelona*)
Juan A. Rodriguez-Aguilar (*Artificial Intelligence Research Institute*)
Javier Morales (*University of Oxford*)
Michael Wooldridge (*University of Oxford*)

Towards a Framework for Predicting Opportunism in Multi-agent Systems (Page 1616)

Jieting Luo (*Utrecht University*)
John-Jules Meyer (*Utrecht University*)
Max Knobbout (*Triple*)

Symmetry Detection and Exploitation for Function Approximation in Deep RL (Page 1619)

Anuj Mahajan (*Conduent Labs India*)
Theja Tulabandhula (*University of Illinois, Chicago*)

Designing Learning Algorithms over the Sequence Form of an Extensive-Form Game (Page 1622)

Edoardo Manino (*University of Southampton*)
Nicola Gatti (*Politecnico di Milano*)
Marcello Restelli (*Politecnico di Milano*)

A Theoretical and Empirical Analysis of Reward Transformations in Multi-Objective Stochastic Games (Page 1625)

Patrick Mannion (*Galway-Mayo Institute of Technology*)
Jim Duggan (*National University of Ireland Galway*)
Enda Howley (*National University of Ireland Galway*)

Limited Observations and Local Information in Convention Emergence (Page 1628)

James Marchant (*University of Warwick*)
Nathan Griffiths (*University of Warwick*)

On the Gap between Outcomes of Voting Rules (Page 1631)

Anurita Mathur (*Indian Institute of Science*)
Arnab Bhattacharyya (*Indian Institute of Science*)

An Empathic Agent that Alleviates Stress by Providing Support via Social Media (Page 1634)

Lenin Medeiros (*Vrije Universiteit Amsterdam*)
Tibor Bosse (*Vrije Universiteit Amsterdam*)

A Conversational Agent Powered by Online Learning (Page 1637)

Vânia Mendonça (*Instituto Superior Técnico / INESC-ID*)
Francisco S. Melo (*Instituto Superior Técnico / INESC-ID*)
Luísa Coheur (*Instituto Superior Técnico / INESC-ID*)
Alberto Sardinha (*Instituto Superior Técnico / INESC-ID*)

SLA-Mechanisms for Electricity Trading under Volatile Supply and Varying Criticality of Demand (Page 1640)

Georgios Methenitis (*Centrum Wiskunde & Informatica, & Delft University of Technology*)
Michael Kaisers (*Centrum Wiskunde & Informatica*)
Han La Poutré (*Centrum Wiskunde & Informatica, & Delft University of Technology*)

Firefly-based Facial Expression Recognition: Extended Abstract (Page 1643)

Kamlesh Mistry (*Teesside University*)
Li Zhang (*Northumbria University*)
Yifeng Zeng (*Teesside University*)
Mengda He (*Teesside University*)

Evolutionary Synthesis of Stable Normative Systems (Page 1646)

Javier Morales (*University of Oxford*)
Michael Wooldridge (*University of Oxford*)
Juan A. Rodríguez-Aguilar (*Artificial Intelligence Research Institute (III-CSIC)*)
Maite López-Sánchez (*University of Barcelona*)

Dealing with Incompatibilities among Goals (Page 1649)

Mariela Morveli Espinoza (*Federal University of Technology - Paraná*)
Ayslan T. Possebom (*Federal University of Technology - Paraná*)
Josep Puyol-Gruart (*Artificial Intelligence Research Institute (III-CSIC)*)
Cesar A. Tacla (*Federal University of Technology - Paraná*)

Model-Based Testing of an Industrial Multi-Robot Navigation System (Page 1652)

Clemens Mühlbacher (*Graz University of Technology*)
Gerald Steinbauer (*Graz University of Technology*)
Stefan Gspandl (*incubedIT*)
Micheal Reip (*incubedIT*)

AQUAMan: An Adaptive QoE-Aware Negotiation Mechanism for SaaS Elasticity Management (Page 1655)

Amro Najjar (*Mines Saint-Etienne*)
Olivier Boissier (*Mines Saint-Etienne*)
Gauthier Picard (*Mines Saint-Etienne*)

Extending the Range of Delivery Drones by Exploratory Learning of Energy Models (Page 1658)

Ty Nguyen (*Ulsan National Institute of Science and Technology*)
Tsz-Chiu Au (*Ulsan National Institute of Science and Technology*)

Generating Virtual Characters from Personality Traits via Reverse Correlation and Linear Programming (Page 1661)

Fabrizio Nunnari (*DFKI / MMCI*)
Alexis Heloir (*DFKI / MMCI / LAMIH UMR CNRS/UVHC 8201*)

Towards a Spatio-Temporal Agent-Based Recommender System (Page 1664)

Amel Ben Othmane (*Université Côte d'Azur*)
Andrea Tettamanzi (*Université Côte d'Azur*)
Serena Villata (*Université Côte d'Azur*)
Nhan Le Thanh (*Université Côte d'Azur*)

A Dominant Strategy Truthful, Deterministic Multi-Armed Bandit Mechanism with Logarithmic Regret (Page 1667)

Divya Padmanabhan (*Indian Institute of Science (IISc)*)
Satyanath Bhat (*Indian Institute of Science (IISc)*)
Prabuchandran K. J. (*IBM IRL*)
Shirish Shevade (*Indian Institute of Science (IISc)*)
Y. Narahari (*Indian Institute of Science (IISc)*)

Uttering Only What is Needed: Enthymemes in Multi-Agent Systems (Page 1670)

Alison R. Panisson (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)
Rafael H. Bordini (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)

FIDES: How Emotions and Small Talks May Influence Trust in an Embodied vs. Non-embodied Robot (Page 1673)

Raul Paradeda (*University of Lisbon & State University of Rio Grande do Norte*)
Mojgan Hashemian (*University of Lisbon*)
Carla Guerra (*University of Lisbon*)
Rui Prada (*University of Lisbon*)
João Dias (*University of Lisbon*)
Ana Paiva (*University of Lisbon*)

Contextual Ranking of Behaviors for Large-scale Multiagent Simulations (Page 1676)

Nidhi Parikh (*Virginia Tech*)
Madhav V. Marathe (*Virginia Tech*)
Samarth Swarup (*Virginia Tech*)

Thwarting Vote Buying Through Decoy Ballots (Page 1679)

David C. Parkes (*Harvard University*)
Paul Tylkin (*Harvard University*)
Lirong Xia (*Rensselaer Polytechnic Institute*)

Curriculum Design for Machine Learners in Sequential Decision Tasks (Page 1682)

Bei Peng (*Washington State University*)
James MacGlashan (*Brown University*)

Robert Lofin (*North Carolina State University*)
Michael L. Littman (*Brown University*)
David L. Roberts (*North Carolina State University*)
Matthew E. Taylor (*Washington State University*)

Detecting Commitment Abandonment by Monitoring Sub-Optimal Steps during Plan Execution (Page 1685)

Ramon Fraga Pereira (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)
Nir Oren (*University of Aberdeen*)
Felipe Meneguzzi (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)

Multiple-Profile Prediction-of-Use Games (Page 1688)

Andrew Perrault (*University of Toronto*)
Craig Boutilier (*University of Toronto*)

Using Tags To Bootstrap Stereotypes And Trust (Page 1691)

Caroline Player (*University of Warwick*)
Nathan Griffiths (*University of Warwick*)

Deeper Understanding of Vague Instructions through Simulated Execution (Page 1694)

Mihai Pomarlan (*University of Bremen*)
Daniel Nyga (*University of Bremen*)
Mareike Picklum (*University of Bremen*)
Sebastian Koralewski (*University of Bremen*)
Michael Beetz (*University of Bremen*)

Data Driven Strategies for Active Monocular SLAM using Inverse Reinforcement Learning (Page 1697)

Vignesh Prasad (*International Institute of Information Technology*)
Rishabh Jangir (*Indian Institute of Technology*)
Ravindran Balaraman (*Indian Institute of Technology*)
K. Madhava Krishna (*International Institute of Information Technology*)

On-the-fly Learning and Monitoring of Partially Observed Navigation Plan (Page 1700)

J. Vince Pulido (*University of Virginia*)
MaryAnne Fields (*U.S. Army Research Laboratory*)
Laura Barnes (*University of Virginia*)

Pure Exploration in Episodic Fixed-Horizon Markov Decision Processes (Page 1703)

Sudeep Raja Putta (*Conduent Labs India*)
Theja Tulabandhula (*University of Illinois at Chicago*)

Analysing Congestion Problems in Multi-agent Reinforcement Learning (Page 1705)

Roxana Radulescu (*Vrije Universiteit Brussel*)
Peter Vranx (*Vrije Universiteit Brussel*)
Ann Nowé (*Vrije Universiteit Brussel*)

Active Perception at the Architecture Level (Page 1708)

Niv Rafaeli (*Bar Ilan University*)
Gal A. Kaminka (*Bar Ilan University*)

Strategically Misleading the User: Building a Deceptive Virtual Suspect (Page 1711)

Diogo Rato (*Universidade de Lisboa*)
Brian Ravenet (*INESC-ID*)
Rui Prada (*Universidade de Lisboa*)
Ana Paiva (*Universidade de Lisboa*)

Extending JaCallIVE Framework to Create Virtual Worlds by Means of an OWL Ontology (Page 1714)

Jaime A. Rincon (*Universitat Politècnica de València*)
Vicente Botti (*Universitat Politècnica de València*)
Vicente Julian (*Universitat Politècnica de València*)
Carlos Carrascosa (*Universitat Politècnica de València*)

MATe: Multiagent Architecture for Taming e-Devices (Page 1716)

Vladimir Rocha (*University of São Paulo*)
Anarosa Alves Franco Brandão (*University of São Paulo*)

Pricing Optimization for Selling Reusable Resources (Page 1719)

Jiang Rong (*Chinese Academy of Sciences*)
Tao Qin (*Microsoft Research*)
Bo An (*Nanyang Technological University*)
Tie-Yan Liu (*Microsoft Research*)

Speeding up Tabular Reinforcement Learning Using State-Action Similarities (Page 1722)

Ariel Rosenfeld (*Bar-Ilan University*)
Matthew E. Taylor (*Washington State University*)
Sarit Kraus (*Bar-Ilan University*)

Model-Driven Engineering in Agent-based Modeling and Simulation: a Case Study in the Traffic Signal Control Domain (Page 1725)

Fernando Santos (*UFRGS & UDESC*)

Ingrid Nunes (*UFRGS & TU Dortmund*)

Ana L. C. Bazzan (*Universidade Federal do Rio Grande do Sul (UFRGS)*)

Analysis of Meta-level Communication for Distributed Resource Allocation Problems (Page 1728)

Matthew Saponaro (*University of Delaware*)

Keith Decker (*University of Delaware*)

Consensus on Social Graphs under Increasing Peer Pressure (Page 1731)

Justin Semonsen (*Rutgers University*)

Christopher Griffin (*United States Naval Academy*)

Anna Squicciarini (*Penn State University*)

Sarah Rajtmajer (*Quantitative Scientific Solutions*)

Inverse Reinforcement Learning Under Noisy Observations (Page 1733)

Shervin Shahryari (*University of Georgia*)

Prashant Doshi (*University of Georgia*)

Optimal Decision Making with CP-nets and PCP-nets (Page 1736)

Sujoy Sikdar (*Rensselaer Polytechnic Institute*)

Sibel Adali (*Rensselaer Polytechnic Institute*)

Lirong Xia (*Rensselaer Polytechnic Institute*)

Multiagent Coordination Using Graph Structured Mathematical Optimization (Page 1739)

Arambam James Singh (*Singapore Management University*)

Akshat Kumar (*Singapore Management University*)

Optimizing Resource Allocation with Intelligent Agents (Page 1742)

Lucas O. Souza (*University of Brasilia*)

Celia G. Ralha (*University of Brasilia*)

Bruno W. P. Hoelz (*University of Brasilia*)

Bootstrapping Trust with Partial and Subjective Observability (Page 1745)

Phillip Taylor (*University of Warwick*)

Nathan Griffiths (*University of Warwick*)

Lina Barakat (*Kings College London*)

Simon Miles (*Kings College London*)

Online Recognition of Navigation Goals Through Goal Mirroring (Page 1748)

Mor Vered (*Bar Ilan University*)

Gal A. Kaminka (*Bar Ilan University*)

On Repeated Stackelberg Security Game with the Cooperative Human Behavior Model for Wildlife Protection (Page 1751)

Binru Wang (*Nanjing University*)

Yuan Zhang (*Nanjing University*)

Sheng Zhong (*Nanjing University*)

Task Parametrization through Multi-modal Analysis of Robot Experiences (Page 1754)

Jan Winkler (*Universität Bremen*)

Asil Kaan Bozcuoglu (*Universität Bremen*)

Mihai Pomarlan (*Universität Bremen*)

Michael Beetz (*Universität Bremen*)

Mechanism Design for Social Law Synthesis under Incomplete Information (Page 1757)

Jun Wu (*Nanjing University*)

Lei Zhang (*Nanjing University*)

Chongjun Wang (*Nanjing University*)

Junyuan Xie (*Nanjing University*)

A Trust-based Mixture of Gaussian Processes Model for Robust Participatory Sensing (Page 1760)

Qikun Xiang (*Nanyang Technological University*)

Jie Zhang (*Nanyang Technological University*)

Ido Nevat (*TUMCREATE*)

Pengfei Zhang (*University of Oxford*)

Budgeted Online Assignment in Crowdsourcing Markets: Theory and Practice (Page 1763)

Pan Xu (*University of Maryland*)

Aravind Srinivasan (*University of Maryland*)

Kanithi K. Sarpatwar (*IBM Thomas J. Watson Research Center*)

Kun-Lung Wu (*IBM Thomas J. Watson Research Center*)

Pursuing a Faster Evader Based on an Agent Team with Unstable Speeds (Page 1766)

Fuhan Yan (*Southeast University*)

Yichuan Jiang (*Southeast University*)

Approval Voting with Intransitive Preferences (Page 1769)

Yongjie Yang (*Saarland University*)

Cognition-enabled Task Interpretation for Human-Robot Teams in a Simulation-based Search and Rescue

Mission (Page 1772)

Fereshta Yazdani (*University of Bremen*)

Matthias Scheutz (*Tufts University*)

Michael Beetz (*University of Bremen*)

A Dynamic Resource Allocation Approach for Concurrent Emergency Events in Metropolitan Regions (Page 1775)

Jihang Zhang (*University of Wollongong*)

Minjie Zhang (*University of Wollongong*)

Fenghui Ren (*University of Wollongong*)

Transfer Learning in Multi-Armed Bandit: A Causal Approach (Page 1778)

Junzhe Zhang (*Purdue University*)

Elias Bareinboim (*Purdue University*)

Demonstration Abstracts

Bang: A System for Training and Visualization in Multi-agent Team Formation: (Demonstration) (Page 1781)

Saulo Antunes Silva (*CEFET-MG*)

Sandro Renato Dias (*CEFET-MG*)

Leandro Soriano Marcolino (*Lancaster University*)

KINGFISHER: Total Maritime Awareness System (Page 1784)

Natalie Fridman (*ImageSat International (ISI)*)

Doron Amir (*ImageSat International (ISI)*)

Ilan Schwartzman (*ImageSat International (ISI)*)

Igor Kleinerman (*ImageSat International (ISI)*)

Yulia Sulkovski (*ImageSat International (ISI)*)

Zenia Gorlik (*ImageSat International (ISI)*)

Sharon Kligisberg (*ImageSat International (ISI)*)

Databroker System: A New Online Advertising System That Keeps Users Engaged While Preserving Their Privacy (Page 1787)

Rica Gonen (*The Open University of Israel*)

Prom Week Meets Skyrim (Page 1790)

Manuel Guimarães (*Universidade de Lisboa*)

Pedro Santos (*Universidade de Lisboa*)

Arnav Jhala (*North Carolina State University*)

Automatic Construction of Agent-based Simulation Using Business Process Diagrams and Ontology-based Models (Page 1793)

Donghun Kang (*Nanyang Technological University*)

Zhenchao C. Bing (*Nanyang Technological University*)

Wen Song (*Nanyang Technological University*)

Zehong Hu (*Nanyang Technological University*)

Shuo Chen (*Nanyang Technological University*)

Jie Zhang (*Nanyang Technological University*)

Hui Xi (*Rolls-Royce Singapore Pte Ltd*)

A Multi-Agent Platform for Augmented Reality based Product-Service Systems (Page 1796)

Nils Masuch (*Technische Universität Berlin*)

Tobias Küster (*Technische Universität Berlin*)

Johannes Fähndrich (*Technische Universität Berlin*)

Marco Lützenberger (*Technische Universität Berlin*)

Sahin Albayrak (*Technische Universität Berlin*)

Mise-En-Scène of Narrative Action in Interactive Storytelling (Page 1799)

Jamie Matthews (*Teesside University*)

Fred Charles (*Bournemouth University*)

Julie Porteous (*Teesside University*)

Alexandra Mendes (*Teesside University*)

ABSTRACTme: Modularized Environment Modeling in Agent-based Simulations (Page 1802)

Deividi Moreira (*Universidade Federal do Rio Grande do Sul (UFRGS)*)

Fernando Santos (*Universidade Federal do Rio Grande do Sul (UFRGS) & UDESC*)

Matheus Barbieri (*Universidade Federal do Rio Grande do Sul (UFRGS)*)

Ingrid Nunes (*Universidade Federal do Rio Grande do Sul (UFRGS) & Technical University of Dortmund*)

Ana L. C. Bazzan (*Universidade Federal do Rio Grande do Sul (UFRGS)*)

A New Solution to the Traffic Managing System for Autonomous Vehicles (Page 1805)

Rodrigo Rodrigues Novaes Júnior (*CEFET-MG*)

Daniel de Sousa Santos (*CEFET-MG*)

Gabriel Martins Franco Santiago (*CEFET-MG*)

Sandro Renato Dias (*CEFET-MG*)

An Interactive Narrative Platform for Story Understanding Experiments (Page 1808)

Julie Porteous (*Teesside University*)
Fred Charles (*Bournemouth University*)
Cameron Smith (*Teesside University*)
Marc Cavazza (*University of Kent*)
Jolien Mouw (*Leiden University*)
Paul van den Broek (*Leiden University*)

Synthesising Industry-Standard Manufacturing Process Controllers (Page 1811)

Lavindra de Silva (*University of Nottingham*)
Paolo Felli (*University of Nottingham*)
Jack C. Chaplin (*University of Nottingham*)
Brian Logan (*University of Nottingham*)
David Sanderson (*University of Nottingham*)
Svetan Ratchev (*University of Nottingham*)

A Multi-Agent System for Coordinating Vessel Traffic (Page 1814)

Teck-Hou Teng (*Singapore Management University*)
Hoong Chuin Lau (*Singapore Management University*)
Akshat Kumar (*Singapore Management University*)

Using Intention Recognition in a Simulation Platform to Assess Physical Activity Levels of an Office Building (Page 1817)

Yifeng Zeng (*Teesside University*)
Zhang Zhang (*Teesside University*)
The Anh Han (*Teesside University*)
Iain R. Spears (*Teesside University*)
Shengchao Qin (*Teesside University*)

Doctoral Consortium Abstracts

Solving Imperfect Recall Games (Page 1820)

Jiri Cermák (*Czech Technical University in Prague*)

Engineering Socially Intelligent Personal Agents via Norms (Page 1822)

Nirav Ajmeri (*North Carolina State University*)

A Framework for Large Scale Complex Adaptive Systems Modeling, Simulation, and Analysis (Page 1824)

Lachlan Birdsey (*University of Adelaide*)

Identifying and Responding to Cooperative Actions in General-sum Normal Form Games (Page 1826)

Steven Damer (*University of Minnesota*)

Investigating Travel Choice Behaviour: a New Approach Using Interactive Experiments with Driving Simulators (Page 1828)

Muhammad Fayyaz (*University of Sydney*)
Michiel Bliemer (*University Of Sydney*)
Matthew Beck (*University Of Sydney*)

Real-World Evaluation and Deployment of Adversary Attack Prediction Models (Page 1830)

Benjamin Ford (*University of Southern California*)

Asking Human Help in Contingent Planning (Page 1832)

Ignasi Andrés Franch (*IME - USP*)

Bridging the Gap Between Simulation and Reality (Page 1834)

Josiah P. Hanna (*The University of Texas at Austin*)

Coalition Formation in Structured Environments (Page 1836)

Ayumi Igarashi (*University of Oxford*)

Coping with Hardness of Welfare Maximization by Introducing Useful Complexity Measures (Page 1838)

Rani Izsak (*Weizmann Institute of Science*)

Agent-based Security and Efficiency Estimation in Airport Terminals (Page 1840)

Stef Janssen (*Delft University of Technology*)

Contingent Payment Mechanisms for Future Coordination (Page 1842)

Hongyao Ma (*Harvard University*)

Adaptive and Dynamic Allocation of Security Resources (Page 1843)

Sara Marie Mc Carthy (*University of Southern California*)

Persuasive Negotiation Dialogues using Rhetorical Arguments (Page 1845)

Mariela Morveli-Espinoza (*Federal University of Technology - Paraná*)

Theoretical Models for Learning from Multiple, Heterogenous and Strategic Agents (Page 1847)

Divya Padmanabhan (*Indian Institute of Science*)

Argumentation Schemes and Enthymemes in Multi-agent Systems (Page 1849)

Alison R. Panisson (*Pontifical Catholic University of Rio Grande do Sul (PUCRS)*)

How Do Humans Teach: On Curriculum Design for Machine Learners (Page 1851)

Bei Peng (*Washington State University*)

Consensus Decision-Making Using Argumentation (Page 1853)

Ayslan Trevisan Possebom (*Federal University of Technology - Paraná*)

Minimising Regret in Route Choice (Page 1855)

Gabriel de O. Ramos (*Universidade Federal do Rio Grande do Sul*)

On Demonstrating the Impact of Defeasible Reasoning in Practice via a Multi-layer Argument-based Framework (Page 1857)

Lucas Rizzo (*Dublin Institute of Technology*)

Social Norms of Cooperation in Multiagent Systems (Page 1859)

Fernando P. Santos (*Universidade de Lisboa*)

Moving Target Defense: A Symbiotic Framework for AI & Security (Page 1861)

Sailik Sengupta (*Arizona State University*)

Strategic Reasoning in Digital Zero-Sum Games (Page 1863)

Anderson Rocha Tavares (*Universidade Federal de Minas Gerais*)

Overcoming Existing Limitations in Electricity-based Artificial Intelligence Applications (Page 1865)

Mark Valovage (*University of Minnesota*)

Intent Recognition Through Goal Mirroring (Page 1867)

Mor Vered (*Bar Ilan University*)

Principled Autonomous Decision Making for Markets (Page 1869)

Enrique Areyan Viqueira (*Brown University*)

Computing Approximate Pure Nash Equilibria in Digraph k -Coloring Games

Raffaello Carosi
Gran Sasso Science Institute -
L'Aquila, Italy
raffaello.carosi@gssi.infn.it

Michele Flammini
Gran Sasso Science Institute,
and DISIM - University of
L'Aquila, Italy
michele.flammini@univaq.it

Gianpiero Monaco
DISIM - University of L'Aquila,
Italy
gianpiero.monaco@univaq.it

ABSTRACT

We investigate approximate pure Nash equilibria in digraph k -coloring games, where we are given an unweighted directed graph together with a set of k colors. Vertices represent agents and arcs capture their mutual unidirectional interests. The strategy set of each agent v consists of the k colors and the payoff of v in a given state or coloring is given by the number of outgoing neighbors with a color different from the one of v . Such games form some of the basic payoff structures in game theory, model lots of real-world scenarios with selfish agents and extend or are related to several fundamental class of games.

It is known that the problem of understanding whether the game admits a pure Nash equilibrium is NP-complete. Therefore we focus on designing polynomial time algorithms that return approximate Nash equilibria. Informally, we say that a coloring is a γ -Nash equilibrium (for some $\gamma \geq 1$) if no agent can strictly improve her payoff by a multiplicative factor of γ by changing color. We first propose a deterministic polynomial time algorithm that, for any $k \geq 3$, returns a k -coloring that is a $\Delta_o(G)$ -Nash equilibrium, where $\Delta_o(G)$ is the maximum outdegree of the digraph.

We then provide our two main results: i) By exploiting the constructive version of the well known Lovász Local Lemma, we show a randomized algorithm with polynomial expected running time that, given any constant $k \geq 2$, computes a constant-Nash equilibrium for a broad class of digraphs, i.e., for digraphs where, for any $v \in V$, $\delta_v^v(G) = \Omega(\ln \Delta_o(G) + \ln \Delta_i(G))$ where $\Delta_o(G)$ (resp. $\Delta_i(G)$) is the maximum outgoing (resp. maximum ingoing) degree of G , and $\delta_v^v(G)$ is the outgoing degree of agent v . ii) For generic digraphs, we show a deterministic polynomial time algorithm that computes a $(1 + \epsilon)$ -Nash equilibrium, for any $\epsilon > 0$, by using $O(\frac{\log n}{\epsilon})$ colors.

Keywords

Noncooperative games; computation; Game theory for practical applications.

1. INTRODUCTION

In this paper, we consider anti-coordination games where

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AA-MAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.

Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

agents need to mutually anti coordinate their strategies in order to maximize their payoffs. Such games form some of the basic payoff structures in game theory and indeed they can model lots of real-world scenarios. For instance, agents can be miners deciding which land to drill for resources. A miner maximizes her happiness when the number of other miners that choose the same land is minimized. Or company employees trying to learn diverse skills. If an employee learns a skill that few other employees also learn, she will have more chance of exploiting it. Moreover, suppose that companies or countries have to produce some product. Typically they are interested in maximizing their profit, that most of the time corresponds to minimize the number of their trading competitors that produce the same commodity. Another example in a social networks setting could be the one in which people have to choose an outfit so that a minimum number of friends wears the same dress. Finally in a radio setting, radio towers are agents and their goal is selecting a frequency such that neighbouring radio-towers have a different one in order to minimize the interference. In all these scenarios agents have an interest in to mutually anti coordinate their strategies.

In general, each situation in which agents have to anti-coordinate their strategies so that the “penalty” (that can be of any type) on their reward is minimized can be described by anti-coordination games. Generally an agent does not care about all other agents’ decisions, but wants to anti-coordinate her strategy only with a subset of agents of interest. Moreover, this interest may not be necessarily mutual but only one-sided.

A way of modeling this kind of setting with selfish agents is the following *graph k -coloring game*. There is a graph where a) agents are vertices, b) edges between agents indicate social connections, namely whether there is a mutual or unidirectional interest (depending on whether the graph is undirected or directed, respectively) about an agent’s decision, and c) the set of the possible strategies is given by a natural number $k \geq 2$ representing the number of alternative available choices. Each agent picks a color (i.e., a number from 1 to k) and the resulting strategy vector induces a k -coloring of the nodes. The objective of every agent is to select the color that maximizes her payoff, which is equal to the number of neighbours with different color from her own.

In such a setting with selfish agents, a stable solution of the game is represented by a (pure) Nash equilibrium, that is a k -coloring where no agent can improve her payoff by changing color. The Nash equilibrium is one of the most important concepts in game theory, forming the basis of much

recent work in multiagent decision making and electronic marketplaces. As such, efficiently computing Nash equilibria is one of the most important problems in computational game theory.

When the graph is undirected, the graph k -coloring game is a potential game [19] and therefore a Nash equilibrium always exists. Moreover, when the graph is unweighted, the game converges (i.e., the dynamic where at each step one agent performs an improving move converges) to a Nash equilibrium in polynomial time [15, 18].

Unfortunately, in case of directed graphs, for any constant $k \geq 2$, the problem of understanding whether the game admits a Nash equilibrium is NP-complete [18]. Therefore, like in a variety of classes of games falling in this class, where Nash equilibria do not exist or cannot be computed in polynomial time, we focus on the milder form of approximate Nash equilibria. Namely, a state is called a γ -Nash equilibrium (for some $\gamma \geq 1$) if no agent can strictly improve her payoff by a factor of γ by changing her strategy. These equilibria can suitably model the cases in which agents incur a proportional cost in changing their strategies.

Motivated by the above results, we focus on the problem of computing γ -Nash equilibria for digraph k -coloring games. Specifically, we design polynomial time algorithms that return k -coloring where no agent can strictly improve her payoff by a factor of γ , for some fixed $\gamma \geq 1$.

We provide some positive results. In particular, we partially answer the following important questions: given a fixed number of available colors, for which family of digraphs we are able to compute good approximate Nash equilibria in polynomial time? Moreover, given a generic digraph, what is the minimum number of colors such that we are able to compute good approximate Nash equilibria in polynomial time?

To the best of our knowledge, this is the first paper providing polynomial time algorithms that compute approximate Nash equilibria for the digraph k -coloring game.

1.1 Our results

All our results refer to digraph k -coloring games induced by a directed graph G with n nodes, m edges and k available colors.

We first notice that, as proved in Section 2, a pure Nash equilibrium is not guaranteed to exist for any number of players n and colors $k < n$. In fact, even for $k = n - 1$, it is possible to provide an instance for which no Nash equilibrium exists, while $k = n$ clearly suffices as it can be easily achieved by assigning a different color to each node.

Moreover we remark that, for any $k \geq 2$, a Nash equilibrium (i.e., $\gamma = 1$) in bipartite digraphs and directed acyclic graphs (DAGs) exists and can be found in polynomial time (see Section 2). However, in general, even for $k = 2$ a γ -Nash equilibria might not exist for any bounded value of γ . In fact, in any 2-coloring of a directed odd cycle (i.e., a cycle with an odd number of nodes) there is always at least one node having utility zero. This property does not hold for $k > 2$. In fact, in Section 3.1 we first provide a polynomial time algorithm that, for any $k \geq 3$, returns a k -coloring that is a $\Delta_o(G)$ -Nash equilibrium, where $\Delta_o(G)$ (resp. $\Delta_i(G)$) is the maximum outgoing (resp. maximum ingoing) degree of G .

Our main results are the following:

- i) By exploiting the constructive version of the well known

Lovász Local Lemma [20], in Section 3.2 we show a randomized algorithm with polynomial expected running time that, given any constant $k \geq 2$, computes a constant approximate Nash equilibrium¹ for a broad class of directed unweighted graphs. In particular, our algorithm works for digraphs $G = (V, E)$ where, for any $v \in V$, the outgoing degree of v is such that $\delta_o^v(G) = \Omega(\ln \Delta_o(G) + \ln \Delta_i(G))$. For instance this holds for digraphs where the minimum outgoing degree is $\Omega(\log n)$.

ii) Then, for generic digraphs, in Section 3.3 we show a deterministic polynomial time algorithm that computes $(1 + \epsilon)$ -Nash equilibria, for any $\epsilon > 0$, by using $O(\frac{\log n}{\epsilon})$ colors. Such a construction works also for a higher number of colors, showing that $(1 + \epsilon)$ -Nash equilibria exist and can be computed in polynomial time for $k = \Omega(\frac{\log n}{\epsilon})$. This contrasts to the fact that, for any $k < n$, pure Nash equilibria are not guaranteed to exist and in general cannot be computed in polynomial time, unless $P = NP$.

1.2 Related Work

The graph k -coloring game has been first investigated in [15, 18], where the authors show that, when the graph is unweighted and undirected, it is possible to compute a Nash Equilibrium in polynomial time. When the graph is weighted undirected, the problem of computing an equilibrium is PLS-complete even for $k = 2$ [27]. In fact, for such a value of k , it coincides with the classical Max Cut game. Poljak et al. [25] prove that Nash equilibria can be computed in polynomial time for the cut game if the maximum degree of the graph is at most 3. Moreover, for graphs with maximum degree $d > 3$, they notice that it is easy to compute an approximate d -Nash equilibrium. [5, 9] give an algorithm that, for any $\epsilon > 0$, computes in polynomial time a $(3 + \epsilon)$ -equilibrium for the cut game. All the above results exploit the potential function method. However, digraph k -coloring games in general do not admit a potential function. Indeed, when the graph is unweighted directed, even the problem of understanding whether they admit a Nash equilibrium is NP-complete for any fixed $k \geq 2$ [18].

Besides cut games, digraph k -coloring games are related to many other fundamental games considered in the scientific literature.

One example is given by the graphical games introduced in [16]. In these games the payoff of each agent depends only on the strategies of her neighbours in a given social knowledge graph defined over the set of the agents, where an arc (i, j) means that j influences i 's payoff. An interesting class of graphical games is the one of the graphical congestion games [7], where each agent has to choose a set of resources while taking into account that each resource e has a latency function f_e depending on the number of agents using e . For the case where each agent can choose only one of the available resources, also called load balancing, and the latency function is linear (i.e., $f_e(x) = x$, where x the number of agents using e), a Nash equilibrium for the arising digraph k -coloring game corresponds to a Nash equilibrium for an equivalent instance of the graphical congestion game and vice versa. In [7] it is shown that each graphical congestion game defined over a directed acyclic graph admits a Nash equilibrium that can also be found in polynomial time

¹The constant depends on different parameters and, for the sake of readability, its value is not specified here. It can be found at the end of the analysis in Section 3.2.

(we describe the related algorithm in Section 2). We are not aware of papers providing polynomial time algorithms that compute approximate Nash equilibria in graphical congestion games for generic directed graphs. However, approximation algorithms are known for many graphical games (see for example [29, 17]).

The game studied in this paper can also be seen as a particular hedonic game (see [4] for a nice introduction to hedonic games) with an upper bound (i.e., k) to the number of coalitions. Specifically, given a k -coloring, the agents with the same color can be seen as members of the same coalition of the hedonic game. In order to get the equivalence among the two games, the hedonic utility of an agent v can be defined as the overall number of her neighbors minus the number of agents of her neighbourhood that are in the same coalition. Nash equilibria issues in hedonic games have been largely investigated under several different assumptions [6, 8, 12, 13, 14, 24] (just to cite a few). An interesting recent study on Nash equilibria for graphical hedonic games can be found in [23]. To the best of our knowledge, no known result concerns approximate Nash equilibria.

Another related stream of research considers coordination games. The idea is that agents are rewarded for choosing common strategies in order to capture the influences. Apt et al. [2] propose a coordination game modeled as an undirected graph where vertices are agents and each agent has a list of allowed colors. Given a coloring, an agent has a payoff equal to the number of adjacent nodes with her same color. The authors show that the game converges to a NE in polynomial time. Apt et al. [3] consider the coordination game on directed weighted graphs. They prove that the problem of determining the existence of a Nash equilibrium on directed graphs is NP-complete. [3, 28] show that there exist some special digraph topologies for which Nash equilibria are guaranteed and can be found in polynomial time. Rahn et al. [26] extend the results of [2] by considering weighted undirected graphs and the problem of finding α -approximate k -equilibria, namely colorings where no coalition of at most k agents can deviate such that each member increases his payoff by at least a factor α . The authors show the existence of α -approximate k -equilibria for $\alpha > 2$. Anshelevich et al. [1] propose a coordination game where agents get incentive for coordinating. They show a polynomial time algorithm that computes a constant approximate Nash equilibrium with good social welfare.

Panagopoulou and Spirakis [22] study games where Nash equilibria are proper vertex coloring in undirected unweighted graphs setting. In particular, they consider the game where each agent v has to choose a color among k available ones and her payoff is equal to the number of vertices in the graph that have chosen her same color, unless some neighbour of v has chosen the same color, and in this case the payoff of v is 0. They prove that this is a potential game and that a Nash equilibrium can be found in polynomial time.

2. PRELIMINARIES

Given an unweighted directed simple graph (or simple digraph) $G = (V, E)$, we suppose that $|V| = n$ and $|E| = m$. An arc $(v, w) \in E$ is directed from node v to w ; w is called the head and v is called the tail of the arc. An outgoing arc from node v is any arc $(v, w) \in E$. We denote by $\delta_o^v(G)$ the outgoing degree of node v , that is the number




Figure 1: Instance for which there is no NE even with a number of colors linear in the number of players. A bold arrow means that there is complete incidence from the source subgraph to the target subgraph.

of outgoing arcs from node v in the graph G . An ingoing arc to node v is any arc $(w, v) \in E$. We denote by $\delta_i^v(G)$ the ingoing degree of node v . Moreover we denote with $d_o(G) = \min_{v=1, \dots, n} \delta_o^v(G)$ the minimum outgoing degree of G . Let $\Delta_o(G) = \max_{v=1, \dots, n} \delta_o^v(G)$ be the maximum outgoing degree of G . Let $\Delta_i(G)$ be the maximum ingoing degree of G . We will omit to specify (G) when clear from the context. Finally we denote with $G[V']$, where $V' \subseteq V$, the subgraph of G induced by V' .

In the *digraph k -coloring game* we are given a directed unweighted graph $G = (V, E)$ without self loops in which each node $v \in V$ is a selfish agent (in the following we will use node and agent interchangeably), and a set of k available colors. Each agent has the same set of actions, which is the set of the k available colors. A state of the game $c = \{c_1, \dots, c_n\}$ is a k -coloring, where c_v is the color (i.e., a number from 1 to k) chosen by agent v . In a certain coloring c , the payoff (or the utility) of an agent is the number of neighbors choosing colors different from her own, where the neighborhood of an agent v is the set of vertices induced by all v 's outgoing edges. Formally, for a coloring c , an agent v 's payoff $\mu_c(v) = \sum_{(v, w) \in E: c_v \neq c_w} 1$.

Let (c_{-v}, c'_v) denote the coloring obtained from c by changing the strategy of agent v from c_v to c'_v . A coloring c is a pure Nash or stable equilibrium, in the following NE for short, if no agent v can improve her payoff by changing strategy, i.e., color. Formally $c = \{c_1, \dots, c_n\}$ is a NE if $\mu_c(v) \geq \mu_{(c_{-v}, c'_v)}(v)$ for any possible color c'_v and for any $v \in V$.

A NE is not guaranteed to exist even for a large number of available colors k . In particular, while there is always a NE with n colors just assigning to each node a different color, for every $k \leq n - 1$ a NE is not guaranteed to exist. In fact, it is possible to prove the following proposition.

Proposition 1. *For arbitrarily large values of $n \geq 3$, and any fixed k such that $1 < k \leq n - 1$, there exist instances of the digraph k -coloring game with n nodes not admitting any NE.*

Proof. Consider the following instance of digraph k -coloring game: there are two nodes x, y , where x has an arc directed toward y ; moreover, there is a complete directed clique K_{k-1} of size $k - 1$, and each node in the clique has an arc directed toward x ; finally, y has arcs directed toward all nodes in the clique.

Suppose we have k colors. Assume by contradiction that a stable coloring exists and let c be the color assigned to x . By the stability constraint, the nodes in the clique must have the remaining $k - 1$ colors, one per vertex. In fact, if some node in the clique is using the same color as x , then it could switch to some unused color that increases its utility. Moreover, also y should have a color different from the ones in the clique, that is it must have color c . But then x would have utility 0 and would improve by switching color: a contradiction to the fact that the coloring was stable.

Notice that in the above construction $k = n - 1$. In order to prove the claim for every fixed value of k such that $1 < k \leq n - 1$, it suffices to add to the above graph of $k + 1$ nodes $n - k - 1$ additional dummy nodes with an arc directed toward x . The instance is depicted in Figure 1. \square

We remark that Proposition 1 is not directly implied by the NP-completeness of determining the existence of a pure NE shown in [18], since the proof works only for $k \ll n$.

If a digraph is bipartite, a NE with $k \geq 2$ colors always exists and it can be easily found in polynomial time. In fact, given a digraph, it is well known that it is possible in polynomial time to test whether the graph is bipartite or not and, in the affirmative case, return a proper 2-coloring of it (i.e., if the set of nodes v is partitioned in to the sets V_1 and V_2 , then all the nodes in V_1 are colored with color 1 and all the nodes in V_2 are colored with color 2. Clearly such a coloring maximizes the utility of all the agents.

It is also easy to see that, if a digraph is without cycles (i.e., it is a DAG), then a NE with $k \geq 2$ colors always exists and it can be easily found in polynomial time. In fact, given a digraph, it is well known that it is possible in polynomial time to test whether the graph is a DAG or not and, in the affirmative case, to return a topological sorting of the vertices, i.e., a linear ordering of its vertices such that for every directed arc (v, w) from vertex v to vertex w , v comes before w in the ordering. Then it is enough to consider vertices in the reverse order and color each node with its best response, that is choosing the color that maximizes her payoff given the choices made by the previous agents.

Unfortunately, for general directed graphs, the problem of determining whether the digraph k -coloring game admits a NE is NP-Hard, for all $k \geq 2$ [18]. For this reason, we consider the notion of approximate Nash equilibrium. A coloring c is a γ -Nash equilibrium (γ -NE or γ -stable for short), for some $\gamma \geq 1$, if no agent can strictly improve her payoff by a multiplicative factor of γ by changing color. Formally a coloring $c = \{c_1, \dots, c_n\}$ is a γ -NE if $\gamma \mu_c(v) \geq \mu_{(c_{-v}, c'_v)}(v)$ for any possible color c'_v and for any $v \in V$.

The notion of approximate NE has been considered in many settings (see [21]). Typically, impossibility results (as in our case), together with the fact that in many real-life applications it may be the case that agents incur a proportional cost for changing their strategies, have naturally led researchers to consider this relaxed notion of Nash equilibrium.

3. APPROXIMATE-NE

In this section we show our polynomial time algorithms that compute approximate NE for digraph k -coloring games.

In order to get familiar with the model and before presenting our main results, we first show a polynomial time algorithm that for any $k \geq 3$ returns a k -coloring such that

agents are in a $\Delta_o(G)$ -NE. This solution is particularly suitable for graphs with small $\Delta_o(G)$.

We notice that when $k = 2$ any 2-coloring of a directed odd cycle (i.e., a cycle with an odd number of nodes) is not a γ -NE for any $\gamma \geq 1$. Indeed it is easy to see that for any possible coloring there exists a node having payoff zero.

3.1 Warm up: computing $\Delta_o(G)$ -NE

We present a simple polynomial time algorithm that given a digraph G returns a k -coloring, for any $k \geq 3$ (indeed the algorithm uses three colors), where every node v such that $\delta_o^v(G) \geq 1$ has payoff at least 1. Clearly this corresponds to a $\Delta_o(G)$ -NE because $\Delta_o(G)$ is the maximum payoff that any agent can achieve. The algorithm is iterative. At each iteration the algorithm visits the graph induced by the uncolored nodes and detects a cycle or a path (in the case of the visiting reaches a node without outgoing edge in the induced subgraph). Then it colors the nodes of the cycle or the path by alternating three colors (for instance colors 1, 2 and 3) in a way that every node gets payoff of at least 1. In particular if the subgraph is a cycle then the algorithm considers nodes of the cycle in clockwise order and assigns the colors in such order (starting by any node) by alternating the three colors. If the subgraph is a path from node v to node w , then it colors the node w by a different color with respect to the already colored node v , if the arc $(w, v) \in E$ (otherwise it means that $\delta_o^w = 0$ and we can assign any color to w), and then alternates colors (in this case two colors are enough) for the other nodes of the path considered in the reverse order starting from w . We notice that if the algorithm does not detect any odd cycle then two colors are enough. A formal description of the algorithm is Algorithm Approx1.

The proof of the following theorem is straightforward.

Theorem 2. *Given a digraph G , Algorithm Approx1 computes in polynomial time a coloring that uses at most three colors where each node v with $\delta_o^v \geq 1$ gets utility at least 1, i.e., Approx1 returns a $\Delta_o(G)$ -NE.*

3.2 Constant approximate NE obtained by exploiting the Lovász Local Lemma

In this section we design an algorithm that for any constant $k \geq 2$, computes a constant approximate NE for a large class of digraph. In particular, by exploiting the well known Lovász Local Lemma [11], we first show that if we color each node with one of the k available colors uniformly at random, there is positive probability (i.e., strictly greater than zero) that such random coloring returns a constant approximate NE for digraph where, for any $v \in V$, $\delta_o^v(G) = \Omega(\ln \Delta_o(G) + \ln \Delta_i(G))$. For instance this happens for digraphs where the minimum outgoing degree is $\Omega(\log n)$. It implies that constant approximate NE always exist for any constant value $k \geq 2$ in this class of digraphs. We point out that this result is already interesting, given that the problem of understanding whether the digraph k -coloring game admits a NE is NP-complete for any $k \geq 2$. Furthermore, we can use the results of Theorem 1.2 of [20] in order to get a randomized algorithm whose expected running time is polynomial, that computes such constant approximate NE (we will describe the algorithm we get from Theorem 1.2 of [20] at the end of this section).

The Lovász Local Lemma (LLL) is a powerful tool used for demonstrating that, given a large set of events with some dependencies among them, the probability that none of these

Algorithm: Algorithm Approx1

Input: a digraph $G = (V, E)$ and a set $1, \dots, k$ of colors

- 1 $L \leftarrow$ empty list
- 2 **while** $\exists v \in V$ still not colored **do**
- 3 $L = [v]$
- 4 $x = v$
- 5 **while** $\exists i \in V$ such that $\{x, i\} \in E$ and i is still not colored **do**
- 6 **if** $i \in L$ **then**
- 7 let $L' \subseteq L$ be the sublist of elements of L starting from node i up to the end
- 8 **if** L' is of even length **then**
- 9 color the nodes in L' by alternating two colors
- 10 $x = i$
- 11 break
- 12 **else**
- 13 color the nodes in $L' \setminus [i]$ by alternating two colors
- 14 color i with a third color
- 15 $x = i$
- 16 break
- 17 **else**
- 18 $x = i$
- 19 append i to the end of L
- 20 **if** $v == x$ and v is not colored **then**
- 21 **if** $\text{deg}^+(v) > 0$ **then**
- 22 assign to v the color that maximizes its payoff
- 23 **else**
- 24 assign to v a random color
- 25 **if** $v \neq x$ and x is not colored **then**
- 26 **if** $\text{deg}^+(x) > 0$ **then**
- 27 assign to x the color that maximizes its payoff
- 28 color the nodes in $L \setminus \{x\}$ from the last node down to v by alternating a color different from $c(x)$ and $c(x)$ itself
- 29 **else**
- 30 color the nodes in L by alternating two colors

events happens is strictly greater than 0 if some conditions are met. We use the following version of the (LLL):

Definition 1. (LLL). Let A_1, A_2, \dots, A_n be a set of bad events, and let $D_i \subseteq \{A_1, A_2, \dots, A_n\}$ denote the "dependency set" of A_i (namely A_i is mutually independent of all the events not in D_i). If there exists a set of real numbers $x_1, \dots, x_n \in [0, 1]$ such that $\Pr[A_i] \leq x_i \prod_{j \in D_i} (1 - x_j)$ for all i , then $\Pr[\bigwedge_{i=1}^n \bar{A}_i] \geq \prod_{i=1}^n (1 - x_i) > 0$.

We show that the LLL can be used for proving that, under certain conditions on the structure of the digraph, there is positive probability that the Random k -coloring algorithm returns a constant approximate NE for any values of $k \geq 2$. That is, we show the existence of constant approximate NE for a broad subclass of digraphs. We define for each $v \in V$ a "bad event" I_v , that is the case in which in a coloring returned by Algorithm Random k -coloring the node v is not

Algorithm: Random k -coloring

Input: a digraph $G = (V, E)$ and a set $1, \dots, k$ of colors

- 1 **for** $i = 1$ to n **do**
- 2 randomly color node i by color $j = 1, \dots, k$, with uniform probability $\frac{1}{k}$.




Figure 2: Graphical representation of all the types of event that can influence a node's behaviour

γ -stable for some constant γ (the value of γ will be computed at the end of the analysis):

Event $I_v =$ "node v is unhappy, i.e., it is not γ -stable."

In order to apply the LLL we first need to bound the maximum size of the dependency set of each bad event. A generic event I_v is dependent on all the bad events that consider at least one node that regards I_v , namely v 's neighbours and v itself. Thus, the possible event types in I_v 's dependency set are shown below together with their maximum possible multiplicity in brackets:

- The events in which v 's neighbours are involved in their neighbours' unhappiness ($\leq \delta_i^v \delta_o^v$);
- The events I_w , for any $w \neq v$ where v contributes to w 's unhappiness, namely there exists a directed edge (w, v) ($\leq \delta_i^v$);
- The events I_w , for any $w \neq v$, that make w unhappy, namely there exists a directed edge (v, w) ($\leq \delta_o^v$).

This means that a bad event I_v is dependent to at most $\delta_i^v \delta_o^v + \delta_i^v + \delta_o^v$ other bad events. The possible dependencies are shown graphically in Figure 2.

If we denote by dep_v the dependency set of I_v , then we have that:

$$|\text{dep}_v| \leq \delta_o^v + \delta_i^v + \delta_o^v \delta_i^v \quad \forall v \in V \quad (1)$$

Let $u_d = \Delta_o + \Delta_i + \Delta_o \Delta_i$ be an upper bound to the number of dependencies of any node, that is $\text{dep}_v \leq u_d$ for any $v \in V$. In the following we suppose that $\text{dep}_v \geq 2$. Indeed if $\text{dep}_v = 1$ then two cases are possible, that is either $\delta_o^v = 1$ and $\delta_i^v = 0$, or, $\delta_o^v = 0$ and $\delta_i^v = 1$. In the former case we can color node v at the end of the procedure with a color different than its unique (outgoing) neighbour. In fact its color does not affect the payoff of any other node. In the latter case, node v is always happy because it does not have

chance to get payoff greater than zero. If $d_v = 0$ then it is an isolated node and therefore we can forget about it. We thus consider that $u_d \geq 2$.

Theorem 3. *There exist values $x_v \in [0, 1)$, for any $v = 1, \dots, n$, such that, under certain conditions, $Pr[I_v] \leq x_v \prod_{w \in dep_v} (1 - x_w)$ for all $v \in V$. That is the LLL holds.*

Proof. In order to show that the LLL holds, let us define $x_v = \frac{1}{u_d} \forall v \in V$. We want to show that if the following inequality is true:

$$Pr[I_v] \leq \frac{1}{u_d} \left(1 - \frac{1}{u_d}\right)^{u_d} \quad \forall v \in V \quad (2)$$

Then, since $u_d \geq |dep_v| \forall v$, and $\left(1 - \frac{1}{u_d}\right) < 1$, it implies:

$$Pr[I_v] \leq \frac{1}{u_d} \left(1 - \frac{1}{u_d}\right)^{u_d} \leq \frac{1}{u_d} \left(1 - \frac{1}{u_d}\right)^{|dep_v|} \quad \forall v \in V \quad (3)$$

namely the LLL holds. In order to prove equation (2) we use the Chernoff bound [10] that gives a bound on the deviation of the sum of random variables from their expected value.

Definition 2. (*Chernoff bound*). let X_1, X_2, \dots, X_n be independent random variables, and let X be their sum and $\mu = E[X]$ the expected value. Then for any $\beta > 0$:

$$Pr[X \geq (1 + \beta)\mu] \leq e^{-\frac{\beta\mu}{3}} \quad (4)$$

So, in order to use the Chernoff bound we rewrite every bad event I_v as:

$$I_v = \text{“node } v \text{ has at least } (1 + \beta) \frac{\delta_o^v}{k} \text{ neighbours colored with the same color”}$$

Notice that in the solution returned by Algorithm Random k -coloring, $\frac{\delta_o^v}{k}$ is the expected number of nodes w having v 's color and such that there exists a directed edge (v, w) . This is true since, with k colors, the probability that a node has the same color as v is $\frac{1}{k}$ and every node v has outdegree δ_o^v . We remark that we suppose k to be a constant value.

We use the Chernoff bound in equation (2) in the following way:

$$Pr[I_v] \leq e^{-\frac{\beta\delta_o^v}{3k}} \leq \frac{1}{u_d} \left(1 - \frac{1}{u_d}\right)^{u_d} \quad (5)$$

We want to find for which values of β , δ_o^v and k the above inequality holds. Since the function $\left(1 - \frac{1}{u_d}\right)^{u_d}$ is increasing for any value $u_d > 1$ and moreover we suppose that $u_d \geq 2$, we obtain that $\left(1 - \frac{1}{u_d}\right)^{u_d} \geq \frac{1}{4}$. So we get from equation

(5) that, if $e^{-\frac{\beta\delta_o^v}{3k}} \leq \frac{1}{4u_d}$, then equation (5) holds, namely the Lovász lemma holds. Therefore we obtain that for each node $v \in V$:

$$\frac{\beta\delta_o^v}{3k} \geq \ln(u_d 4) = \ln[(\Delta_o \Delta_i + \Delta_o + \Delta_i)4] \geq \ln \Delta_o \Delta_i + \ln 4$$

It implies that:

$$\delta_o^v \geq \frac{3k}{\beta} (\ln \Delta_o \Delta_i + \ln 4) \quad (6)$$

Therefore, since β and k are constant values, then we get that when $\delta_o^v = \Omega(\ln \Delta_o + \ln \Delta_i)$, i.e., when the outgoing degree of each node v is sufficiently large then the LLL holds. For instance we get that the LLL holds for general unweighted digraphs where $d_o = \Omega(\ln n)$. Moreover, with a very similar analysis it is possible to show that if the digraph is such that $\delta_o^v = \delta_o$ for any $v \in V$, i.e, the outgoing degree of all the nodes is the same (this is a broader class of digraphs than regular ones) then LLL holds when $\delta_o = \Omega(\ln \Delta_i)$. Finally we notice that the greater d_o is, the smaller β is required to be.

We now show which is the smallest value of γ such that γ -NE exists. For doing so, we consider the minimum possible value β such that the LLL still holds, that is, according to equation (6) equal to $\frac{3k}{\delta_o^v} ((\ln \Delta_o + \ln \Delta_i) + \ln 4)$.

$$\begin{aligned} \gamma &= \max_{v \in V} \frac{\text{maximum possible utility}}{\text{minimum expected utility}} = \frac{\delta_o^v}{\delta_o^v - (1 + \beta) \frac{\delta_o^v}{k}} = \\ &= \frac{k}{k - (1 + \beta)} = \frac{k}{k - 1 - \frac{3k(\ln(\Delta_o \Delta_i) + \ln 4)}{\delta_o^v}} \\ &= \frac{\delta_o^v}{\delta_o^v - \frac{\delta_o^v}{k} - 3(\ln(\Delta_o \Delta_i) + \ln 4)} = \\ &= \frac{\delta_o^v - \frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4) + \frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4)}{\frac{k-1}{k} \delta_o^v - 3(\ln(\Delta_o \Delta_i) + \ln 4)} = \\ &= \frac{\delta_o^v - \frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4) + \frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4)}{\frac{k-1}{k} \left(\delta_o^v - \frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4) \right)} = \\ &= \frac{k}{k-1} + \frac{\frac{k}{k-1} 3(\ln(\Delta_o \Delta_i) + \ln 4)}{\frac{k-1}{k} \delta_o^v - 3(\ln(\Delta_o \Delta_i) + \ln 4)} \approx \\ &\approx \frac{k}{k-1} + \left(\frac{k}{k-1} \right)^2 O \left(\frac{1}{r - \frac{k}{k-1}} \right) \end{aligned}$$

$$\text{Where } r = \frac{\delta_o^v}{\ln(\Delta_o \Delta_i)}.$$

It is easy to see that when $\delta_o^v = \Omega(\ln \Delta_o \Delta_i)$ the above value of γ is constant. \square

By summarizing we have proved that, under certain conditions, according to the LLL, Algorithm Random k -coloring returns a constant approximate NE with probability greater than zero.

Moreover, by Theorem 1.2 of [20] we know that there exists a simple randomized algorithm that returns a stable coloring. The algorithm is very simple: start from a random coloring (not necessarily stable) and, if exists, we arbitrarily

pick an unhappy node v (i.e. the bad event I_v holds) and we randomly assign a new color to all the nodes in the dependency set dep_v (this is called a resampling of the event I_v). We continue resampling violated events until the algorithm reaches a coloring that is a γ -NE (i.e. the values of γ we get in the above analysis). If the conditions of the LLL are satisfied, then [20] proves that the algorithm terminates computing a γ -NE and the expected number of times an event is resampled is polynomial.

We point out that the question of whether the algorithm can be derandomized is an interesting issue. Unfortunately, the results of [20], Theorem 1.4, allow to obtain a deterministic algorithm with a running time that depends exponentially on the maximum number of dependencies per random variable. As a consequence, if the input graph is such that the in- and out-degree of each vertex are bounded by a constant, the conditions of Theorem 1.4 are satisfied and we directly obtain a deterministic polynomial time algorithm. However, for such a case, the simple polynomial time construction provided in Section 3.1 already gives a constant Nash equilibria with 3 colors. For general graphs, understanding whether the algorithm can be derandomized is a hard question, since it is directly connected to the more general question of improving the existing construction results of LLL.

3.3 Approximate NE for general digraphs with a logarithmic number of colors

In this section we present a polynomial time algorithm that, for any digraph G and $\epsilon > 0$, computes a $(1 + \epsilon)$ -NE by using $O\left(\frac{\log n}{\epsilon}\right)$ colors. The algorithm is iterative. At the beginning all the nodes are not colored. In each iteration the algorithm colors a subset of nodes as described below.

Let V' be the current set of uncolored nodes (at the beginning $V' = V$). In each iteration i , we consider the subdigraph $G' = (V', E')$ of G induced by V' , and $k' = \lceil \frac{3(1+\epsilon)}{\epsilon} \rceil$ new colors $c_{ik'-k'+1}, \dots, c_{ik'}$ not used in the previous iterations. We then define the undirected graph $G'' = (V', E'')$ such that an undirected edge $\{v, w\}$ is in E'' if at least one between (v, w) and (w, v) is in E' . Let δ^v be the number of neighbors of v in G'' .

Next, we use the algorithm described in Section 3 of [18] that computes a stable coloring for the unweighted undirected graph k' -coloring game (from now on we call this algorithm `UndirectColoring`). The underlying idea of the algorithm is the following: it starts from any arbitrary coloring and one node changes color if, by doing so, it strictly improves its utility (i.e. the number of neighbors with different color increases). The algorithm stops when it is no more possible to perform such moves. The solution is computed in polynomial time and it is a NE. Let $V_1, V_2, \dots, V_{k'}$ be the coloring of V' induced by `UndirectColoring`, namely $v \in V_j$ means that v is colored $c_{ik'-k'+j}$.

For each node $v \in V'$, if its outgoing degree in G' is at least $\lceil \frac{\delta^v}{3} \rceil$ then v is colored as in the equilibrium computed by `UndirectColoring`, namely v is colored $c_{ik'-k'+j}$ if $v \in V_j$. The algorithm updates the set of the uncolored nodes and iterates until all nodes are colored. A formal description of the algorithm is Algorithm `Approx3`.

We now show the performance and correctness of Algorithm `Approx3`.

Theorem 4. *Given any digraph $G = (V, E)$ and $\epsilon > 0$,*

Algorithm: `Approx3`

Input: a digraph $G = (V, E)$, $k' = \lceil \frac{3(1+\epsilon)}{\epsilon} \rceil$.
Output: a coloring of G such that all nodes are in a $(1 + \epsilon)$ -NE, for any $\epsilon > 0$.

```

1  $i = 1$ 
2  $V' = V$ 
3 while  $V' \neq \emptyset$  do
4    $c_{ik'-k'+1}, \dots, c_{ik'}$   $\leftarrow$   $k'$  new colors
5    $G' = (V', E')$   $\leftarrow$  digraph such that
      $E' := \{(v, w) : v, w \in V'\}$ 
6    $G'' = (V', E'')$   $\leftarrow$  undirected graph such that
      $E'' := \{\{v, w\} : (v, w) \in E' \vee (w, v) \in E'\}$ 
7    $\delta^v \leftarrow |\{y \in V' : \{v, y\} \in E''\}| \forall v \in V'$ 
8   apply UndirectColoring to  $G''$  and let  $V_1, V_2, \dots, V_{k'}$ 
     be the partition of  $V$  induced by the coloring
9   for every node  $v \in V'$  do
10    if  $\delta^v(G') \geq \lceil \delta^v/3 \rceil$  then
11      color  $v$   $c_{ik'-k'+j}$ , where  $v \in V_j$ 
12     $V' = V' \setminus \{v\}$ 
13  $i = i+1$ 

```

algorithm `Approx3` returns in polynomial time a $(1 + \epsilon)$ -NE by using at most $\frac{6(1+\epsilon)}{\epsilon} \log n = O\left(\frac{\log n}{\epsilon}\right)$ colors.

Proof. Consider an iteration i . If the coloring returned by `UndirectColoring` induces a partition $V_1, V_2, \dots, V_{k'}$ of V' in iteration i , then for each $v \in V'$ let $\delta_{V_j}^v$ be the number of neighbors that v has in subset V_j . We have:

$$\delta_{V_j}^v \geq \delta_{V_c}^v \quad \text{for all } v \in V_c, \text{ and for all } V_j, j \neq c \quad (7)$$

because otherwise v could improve its utility by changing its color, thus violating the fact that the solution returned is stable. Consider then a node $v \in V_c$. According to the coloring returned by `UndirectColoring`, its utility is:

$$\sum_{V_j: j \neq c} \delta_{V_j}^v \geq (k' - 1) \delta_{V_c}^v \quad (8)$$

This means that, since the degree of v in G'' is δ^v , the number of neighbors in the same cluster as v is at most:

$$\delta_{V_c}^v \leq \frac{\delta^v}{k'} = \frac{\delta^v}{\lceil \frac{3(1+\epsilon)}{\epsilon} \rceil} \leq \frac{\delta^v \epsilon}{3(1+\epsilon)} \quad (9)$$

Let us assume that v is colored at the end of iteration i . Then, by the construction of `Approx3`, $\delta^v(G') \geq \lceil \frac{\delta^v}{3} \rceil$, so that the approximation factor γ_v of node v is

$$\begin{aligned} \gamma_v &\leq \frac{\text{maximum utility of } v}{\text{utility of } v} \leq \frac{\delta^v(G')}{\delta^v(G') - \frac{\delta^v \epsilon}{3(1+\epsilon)}} \leq \\ &\leq \frac{\frac{\delta^v}{3}}{\frac{\delta^v}{3} - \frac{\delta^v \epsilon}{3(1+\epsilon)}} = \frac{1}{1 - \frac{\epsilon}{1+\epsilon}} = 1 + \epsilon. \end{aligned}$$

Thus, v is $(1 + \epsilon)$ -stable, and since every node is colored during some iteration, `Approx3` finally returns a $(1 + \epsilon)$ -stable coloring.

We now show that Approx3 ends in a polynomial number of iterations by proving that we reduce the number of edges in G'' by a constant multiplicative factor in each iteration, until the algorithm ends. Consider again a generic iteration i . Let A be the set of nodes v such that $\delta_v^v(G') \geq \lceil \frac{\delta_v^v}{3} \rceil$ (namely they are colored) and let B be the remaining nodes. If E_i is the set of edges when iteration i starts and E_{i+1} is the set of edges remaining after iteration i , then:

$$|E_{i+1}| \leq \sum_{v \in B} \delta_v^v(G') \leq \frac{\sum_{v \in B} \delta_v^v(G')}{2} \leq \frac{|E_i|}{2} \quad (10)$$

This is true because the first inequality holds by definition (we keep only arcs between uncolored nodes) and if the second inequality is not true, then there is some node in B that has been colored at the end of iteration i , namely it should be in set A . Thus, the maximum number of iterations is $\log|E| \leq 2 \log n$. Since we use k' new colors in each iteration, the overall number of used colors is at most $2k' \log n \leq \frac{6(1+\epsilon)}{\epsilon} \log n = O(\frac{\log n}{\epsilon})$. \square

We remark that the coloring returned by the algorithm remains stable even if we increase the number of colors, so that a $(1+\epsilon)$ -approximate NE exist for any $k = \Omega(\frac{\log n}{\epsilon})$.

4. CONCLUSIONS AND FUTURE WORK

In this paper we considered the digraph k -coloring games, that are able to model lots of real-world scenarios with selfish agents and are related to some of the most fundamental classes of games investigated in the scientific literature.

We focused on Nash stability and, motivated by the fact that NE might not exist for every number of colors $k < n$ and that in general their determination is a NP-hard problem, we focused on the design of polynomial time algorithms that return approximate Nash Equilibria for various values of k . Clearly such results are stronger than mere existential ones, since they imply directly the existence and the fact that such equilibria can be suitably reached in practical scenarios.

Our paper should be seen as opening this research direction. In fact, we provided the first non trivial results for unweighted digraphs. However there are many open problems suggested by our work.

On one hand, it is important to understand what is the best approximate NE that can be computed in polynomial time for every $k = O(\log n)$. On the other hand, even if it turns out that it is computationally hard to determine good approximate Nash equilibria, for such values of k it would be nice also to prove that good approximate Nash equilibria always exist. In particular, what is the minimum value of γ guaranteeing the existence of a γ -approximate equilibria?

Moreover, a further step is that of considering weighted digraphs. Even if we did not treat this case explicitly, we emphasize that our result for $k = \Omega(\frac{\log n}{\epsilon})$ colors can be generalized to show the existence of an equilibria with the same approximation ratio for $k = \Omega(\frac{\log W}{\epsilon})$, where W is the overall sum of the edge weights. However, its determination in polynomial time remains open. In this wording, it is also worth remarking that, even when the graph is weighted undirected (we recall that in such a case pure Nash equilibria always exists since we have a potential game,

however the problem of computing pure Nash Equilibria is PLS-complete), efficient algorithms that compute constant approximate Nash equilibria only exist for the case when $k = 2$, that is for the cut game, with $\gamma = 3 + \epsilon$. Moreover, even for such a particular setting, there is no evidence that the current proposed solutions are the best approximate Nash equilibria that can be computed in polynomial time, i.e., that the known results are tight.

In general, we believe that questions related to the computational complexity of approximate Nash equilibria for the digraph k -coloring games deserve further investigation.

REFERENCES

- [1] E. Anshelevich and S. Sekar. Approximate equilibrium and incentivizing social coordination. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.*, pages 508–514, 2014.
- [2] K. R. Apt, M. Rahn, G. Schäfer, and S. Simon. Coordination games on graphs (extended abstract). In *Web and Internet Economics - 10th International Conference, WINE 2014. Proceedings*, pages 441–446, 2014.
- [3] K. R. Apt, S. Simon, and D. Wojtczak. Coordination games on directed graphs. In *Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015.*, pages 67–80, 2015.
- [4] H. Aziz and R. Savani. Hedonic games. In *Handbook of Computational Social Choice*, chapter 15. Cambridge University Press, 2016.
- [5] A. Bhalgat, T. Chakraborty, and S. Khanna. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *Proceedings 11th ACM Conference on Electronic Commerce (EC), 2010*, pages 73–82, 2010.
- [6] V. Bilò, A. Fanelli, M. Flammini, G. Monaco, and L. Moscardelli. On the price of stability of fractional hedonic games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015*, pages 1239–1247, 2015.
- [7] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Graphical congestion games. *Algorithmica*, 61(2):274–297, 2011.
- [8] F. Bloch and E. Diamantoudi. Noncooperative formation of coalitions in hedonic games. *Int. J. Game Theory*, 40(2):263–280, 2011.
- [9] I. Caragiannis, A. Fanelli, and N. Gravin. Short sequences of improvement moves lead to approximate equilibria in constraint satisfaction games. In *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Proceedings*, pages 49–60, 2014.
- [10] H. Chernoff. Chernoff bound. In *International Encyclopedia of Statistical Science*, pages 242–243. 2011.
- [11] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *A. Hajnal, R. Rado, and V. T. Sós, eds. Infinite and Finite Sets (to Paul Erdős on his 60th birthday)*, pages 609–627. 1975.
- [12] M. Feldman, L. Lewin-Eytan, and J. Naor. Hedonic

- clustering games. *TOPC*, 2(1):4, 2015.
- [13] M. Gairing and R. Savani. Computing stable outcomes in hedonic games. In *Algorithmic Game Theory - Third International Symposium, SAGT 2010*, pages 174–185, 2010.
- [14] M. Gairing and R. Savani. Computing stable outcomes in hedonic games with voting-based deviations. In *10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 559–566, 2011.
- [15] M. Hoefer. *Cost sharing and clustering under distributed competition*. PhD thesis, University of Konstanz, 2007.
- [16] M. J. Kearns, M. L. Littman, and S. P. Singh. Graphical models for game theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, 2001*, pages 253–260, 2001.
- [17] M. J. Kearns, M. L. Littman, and S. P. Singh. Graphical models for game theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, University of Washington, Seattle, Washington, USA, August 2-5, 2001*, pages 253–260, 2001.
- [18] J. Kun, B. Powers, and L. Reyzin. Anti-coordination games and stable graph colorings. In *Algorithmic Game Theory - 6th International Symposium, SAGT 2013, Proceedings*, pages 122–133, 2013.
- [19] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [20] R. A. Moser and G. Tardos. A constructive proof of the general lovász local lemma. *J. ACM*, 57(2), 2010.
- [21] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [22] P. N. Panagopoulou and P. G. Spirakis. A game theoretic approach for efficient graph coloring. In *Algorithms and Computation, 19th International Symposium, ISAAC 2008, Proceedings*, pages 183–195, 2008.
- [23] D. Peters. Graphical hedonic games of bounded treewidth. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 2016*, pages 586–593, 2016.
- [24] D. Peters and E. Elkind. Simple causes of complexity in hedonic games. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, pages 617–623, 2015.
- [25] S. Poljak. Integer linear programs and local search for max-cut. *SIAM J. Comput.*, 24(4):822–839, 1995.
- [26] M. Rahn and G. Schäfer. Efficient equilibria in polymatrix coordination games. In *Mathematical Foundations of Computer Science 2015 - 40th International Symposium, MFCS 2015, Proceedings*, pages 529–541, 2015.
- [27] A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87, 1991.
- [28] S. Simon and D. Wojtczak. Efficient local search in coordination games on graphs. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, pages 482–488, 2016.
- [29] D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence, July 28 - August 1, 2002, Edmonton, Alberta, Canada.*, pages 345–351, 2002.