

Extended Curriculum Vitae

Prof. Pietro Michiardi, Ph.D.

Appointments

2019 – now **Full Professor**, EURECOM, France

2016 – now **Director of the Data Science Department**, EURECOM, France
[50+ staff members, of which 7 Professors, 3 Research Staff, ~2 M annual budget]

2014 – 2019 **Associate Professor**, EURECOM, France

2009 – 2014 **Assistant Professor**, EURECOM, France

2005 – 2009 **Junior Assistant Professor**, EURECOM, France

2001 – 2004 **Ph.D. Student**, Telecom ParisTech, France

2000 – 2001 **Research Engineer**, EURECOM, France

1999 – 2000 **System Engineer**, Dolby Labs., USA

Education

2013: **HDR**, Habilitation a Diriger des Recherches, UCA, France

2004: **Ph.D.**, Computer Science (honors), Telecom Paris Tech, France

2001: **M.Sc.**, Electrical Engineering (honors), Politecnico di Torino, Italy

Coordination Activities / Board Membership

- **Member of Scientific Evaluation Committee:** 2023 – High Council for Evaluation of Research and Higher Education. Scientific evaluation of research units.
- **Head of the Tenure Committee:** 2023 – now, EURECOM. Chair of the tenure track committee, responsible for the evaluation of tenure track faculty.
- **Steering Committee:** 2019 – now, Interdisciplinary Institute of Artificial Intelligence, 3IA Cote d'Azur. Chair and Ph.D. selection committees

- **Scientific Board:** 2019 – now, Interdisciplinary Institutes of Artificial Intelligence, 3IA Cote d’Azur. Scientific coordination and animation.
- **Scientific Advisory Board:** 2018 – now, Data Science and AI Institute, Politecnico di Torino. Scientific coordination and consulting.
- **Scientific Evaluator:** 2016 – European Commission. Scientific evaluation of research proposals.
- **Scientific Advisory Board:** 2015 – now, GridPocket, SA. Scientific consulting.
- **Founder and Coordinator:** 2015 – 2019, of the *Data Science and Engineering Track / M.Sc.*. Created one of the first Data Science Engineering track and Masters’ in France. The track and the Master program are now one of the most popular at EURECOM, with more than 70% of all incoming students.
- **Coordinator:** 2005 – now, of the Double-Degree M.Sc. Programme Politecnico di Torino. Coordination at the University Rector level for the double-degree student exchange program. Incoming students from Politecnico di Torino represent 20% of all incoming students at EURECOM.
- **Coordinator:** 2012 – now, of the Joint Doctoral School EURECOM / Politecnico di Torino. Scientific coordination and Ph.D. selection committees.

Funded Research Projects

Industry Funded Projects

- **Huawei** [2022 – 2024]: Zero-shot Learning for Application Discovery (CIFRE)
- **Renault** [2022 – 2024]: Multimodal Representation Learning (CIFRE)
- **Huawei** [2022 – 2026]: Representation Learning
- **Huawei** [2021 – 2023]: Network Anomaly Detection (CIFRE)
- **Monaco Government** [2021 – 2022]: Bayesian Optimization for Tsunami Simulation and Evacuation Planning
- **Oracle Labs Zurich** [2021]: Probabilistic Financial Transaction Fraud Detection
- **Orange** [2019 – 2021]: Unsupervised Anomaly Detection (CIFRE)
- **Renault** [2019 – 2021]: Hierarchical, Probabilistic Reinforcement Learning (CIFRE)
- **Renault** [2019 – 2021]: Probabilistic RAW Data Fusion (CIFRE)
- **Amadeus** [2018 – 2021]: Probabilistic Price Forecasting
- **SAP** [2018 – 2021]: Interpretable Machine Learning (CIFRE)
- **Amadeus** [2016 – 2018]: Probabilistic Novelty Detection (CIFRE)

- **KPMG Frankfurt** [2017 – 2019]: Foundations of Probabilistic Machine Learning
- **Huawei Sophia-Antipolis** [2018]: Bayesian Optimization for ISPs
- **Monaco Government** [2016 – 2018]: Data Management for Smart Cities
- **KPMG Berlin** [2016 – 2017]: Distributed Cloud Infrastructures
- **KPMG Frankfurt** [2017]: Structured Extraction of Financial Data
- **AirFrance / KLM** [2015 – 2020]: Distributed Cloud Infrastructures for Machine Learning Workflows
- **Orange Labs** [2014 – 2016]: Performance Evaluation of Cloud Virtualization

EU Funded Projects

- CHIST-ERA Europe, *Muse-Com*, (2024-2027): **Principal Investigator**, Multi-modal Generative Modeling
- Horizon Europe, *SCYCLOPS*, (2023-2025): **Co-Principal Investigator**, Generative Modeling
- Horizon Europe, *Adroit6G*, (2023-2025): **Principal Investigator**, Diffusion-based Generative Modeling
- Horizon Europe, *Syclops*, (2023-2025): **Co-Principal Investigator**, Open, full-stack AI hardware/software co-design
- H2020, *IOSTACK*, (2015-2018): **Principal Investigator**, Software Defined Storage for Big Data Applications, <http://www.iostack.eu>
- IST-STREP FP7, *BIGFOOT*, (2012-2015): **Global Coordinator**, Big Data Analytics of Digital Footprints, <http://bigfootproject.eu>
- IST-IP FP7, *MPLANE*, (2012-2015): **Principal Investigator**, Intelligent Measurement Plane for Future Network and Application Management, <http://www.ict-mplane.eu>
- IST-STREP FP7, *NADA*, (2009-2011): **Principal Investigator**, Edge capacity hosting overlays of NANO DATA centers, <http://www.nanodatacenters.eu>
- IST-FET FP6, *CASCADAS* (2006-2008): **Principal Investigator**, Componentware for Autonomic Situation-aware Communications, and Dynamically Adaptable Services, <http://www.cascadas-project.org>
- IST-FET FP5, *MOBILEMAN* (2002-2004): Mobile Metropolitan Ad hoc Networks, <http://cnd.iit.cnr.it/mobileMAN>

French-Government Funded Projects:

- FUI, *eBOB*, (2015-2018): **Principal Investigator**, Scalable Machine Learning Algorithms for Price Prediction
- ANR, *PROSE*, (2010-2013): **Principal Investigator**, Content Sharing Through Peer-to-Peer Recommendation & Opportunistic Social Environment
- ANR, *VIPEER*, (2010-2013): **Principal Investigator**, Video Traffic Engineering in an Intra-Domain Context using Peer-to-Peer Paradigms
- ANR-A*STAR, MERLION project (SINGAPORE) (2009-2010): **Principal Investigator**, Peer-to-peer Content Storage
- ANR, *SPREADS*, (2008-2010): Safe P2p Reliable Architecture for Data Storage, <http://www.spreads.fr>
- GET, *DISPAIRSE* (2007-2008): Disque Pair a Pair Securise, <https://dispairse.point6.net/index.php/Accueil>
- ACI, *SPLASH* (2003-2005): Securing mobile ad hoc networks, <http://www.inrialpes.fr/planete/splash.html>

Teaching Activities

- **Deep Learning**, 2020 – now
- **Algorithmic Machine Learning**, 2015 – now
- **Large-scale Distributed Systems and Cloud Computing**, 2011–2019
- **Algorithm Design**, 2007–2014
- **Applied Game Theory**, 2009–2012
- **Web Technologies**, 2005–2011

Academic Supervision

PhD Students

- **Chao Wang**, CIFRE Sorbonne Univ., France, 2022-
- **Mustafa Bounoua**, CIFRE Sorbonne Univ., France, 2022-
- **Raphael Azorin**, CIFRE Sorbonne Univ., France, 2021-
- **Matthieu DaSilva**, CIFRE Sorbonne Univ., France, 2019-
- **Ugo Lecerf**, CIFRE Sorbonne Univ., France, 2019-2022
- **Lucas Pascal**, CIFRE Sorbonne Univ., France, 2019-2021

- **Julien Audibert**, Sorbonne Univ., France, 2019-2021
- **Rosa Candela**, Sorbonne Univ., France, 2018-2020
- **Graziano Mita**, CIFRE Sorbonne Univ., France, 2018-2020
- **Remi Dominguez**, CIFRE Telecom ParisTech, France, 2016–2018
- **Duc Trung Nguyen**, Telecom ParisTech, France, 2015–
- **Francesco Pace**, Telecom ParisTech, France, 2015–2018
- **Fabio Pulvirenti**, Politecnico di Torino, Italy, 2014–2017
- **Thibault Debatty**, Royal Military Academy, Belgium, 2014–2018
- **Son Hai Ha**, Telecom ParisTech, France, 2014–
- **Yongchao Tian**, Telecom ParisTech, France, 2014–2016
- **Hung-Duy Phan**, Telecom ParisTech, France 2013–2016
- **Mario Pastorelli**, Telecom ParisTech, France, 2011-2014
- **Xiaolan Sha**, CIFRE Telecom ParisTech, France, 2010-2013
- **Francesco Albanese**, Telecom ParisTech, France, 2008-
- **Laszlo Toka**, Telecom ParisTech, France, 2007-2011
- **Chi-Anh La**, Telecom ParisTech, France, 2006-2010

Visiting PhD Students

- **Giulio Franzese**, Visiting Ph.D. student (Politecnico di Torino), 2019-2020
- **Alessandro Lulli**, Visiting Ph.D. student (University of Pisa / CNR, Italy), 2015
- **Antonio Barbuzzi**, Visiting Ph.D. student (Politecnico di Bari, Italy), 2010
- **Flavio Esposito**, Visiting Ph.D. student (Boston University, USA), 2007
- **Krishna Ramachandran**, Visiting Ph.D. student (RPI, USA), 2006

Post Doctoral Researchers

- **Antonio Mastropietro**, Post-Doctoral Fellow (Ph.D. Politecnico di Torino)
- **Giulio Franzese**, Post-Doctoral Fellow (Ph.D. Politecnico di Torino)
- **Dimitrios Milios**, Post-Doctoral Fellow (Ph.D. University of Edinburgh)
- **Marco Milanese**, Research engineer (Ph.D. University of Torino)
- **Paul Pidou**, Research engineer (M.S. EURECOM)

- **Daniele Venzano**, Research engineer (M.S. University of Genova)
- **Xiaolan Sha**, Post-Doctoral Fellow (Ph.D. Telecom ParisTech)
- **Antonio Barbuzzi**, Post-Doctoral Fellow (Ph.D. Politecnico di Bari, Italy)
- **Matteo Dell’Amico**, Post-Doctoral Fellow (Ph.D. Univ. of Genova, Italy)
- **Damiano Carra**, Post-Doctoral Fellow (Ph.D. Univ. of Trento, Italy)

Honors and Awards

- *Notable Reviewer* award at ICLR, 2023
- *Top Reviewer* award at NEURIPS, 2022
- *Europe Raising Star* award, for coordinating the FP7 BigFoot project, 2016
- *Best paper* award at the IEEE ITC Conference, Wurzburg, Germany, 2016
- *Best paper* award at IEEE WiOpt, Sophia Antipolis, France, 2003
- *Telecom Valley Innovation Prize*, Sophia Antipolis, France, 2001
- *Erasmus Award Scholarship* for the academic year 1998-1999

Academic Activities

Chair Role: IEEE BigDap Workshop, 2015; IEEE WOWMOM / TSPUC, 2006

Local Conference Organizer: IEEE SecureComm, Nice, 2007; IEEE WCW, Sophia Antipolis, 2005

TPC Membership: NeurIPS, ICML, ICLR, AISTATS, IEEE KDIR, ACM DIDL, ACM WOSC, ACM SIGCOMM, ACM IMC, ACM MobiHoc, IEEE MASCOTS, IEEE CLOUD, IEEE CloudCom, IEEE EuroPar

Reviewer: NeurIPS, ICML, ACM SIGCOMM, IEEE INFOCOM, ACM CoNEXT, ACM CCS, ACM MobiHoc, ACM TISSEC, IEEE/ACM TON, IEEE ICNP, IEEE JSAC, IEEE TOMC, IEEE TPDS, IEEE TDSC, IEEE ICASSP

Patents

- **Renault** [2023] : “Reconstruction of impaired sensor data through multi-modal latent diffusion model”
- **Renault** [2020] : Concealed
- **Renault** [2020] : Concealed
- **Renault** [2019] : “Multi-Modal Auotecoding Vehicle Sensing”

- **Amadeus** [2018] : “Predicting Error Metrics of Time-Series Forecasting Models” (co-inventor with Prof. Filippone)
- **Huawei** [2018] : “Pareto Front Picture Quality Slider - Subjective interface to Objective Picture Quality Tuning” (co-inventor with Prof. Filippone)
- **Eurecom** [2003] : “Process for providing non repudiation of receipt (NRR) in an electronic transaction environment” (co-inventor with Prof. Molva)

References

- [1] Giulio Franzese, Simone Rossi, Lixuan Yang, Alessandro Finamore, Dario Rossi, Maurizio Filippone, and Pietro Michiardi. How much is enough? A study on diffusion times in score-based generative models. *Entropy*, 25(4):633, 2023.
- [2] Francesco Pace, Daniele Venzano, Damiano Carra, and Pietro Michiardi. A flexible heuristic to schedule distributed analytic applications in compute clusters. *IEEE Trans. Cloud Comput.*, 11(3):2217–2230, 2023.
- [3] Raphaël Azorin, Massimo Gallo, Alessandro Finamore, Dario Rossi, and Pietro Michiardi. "it's a match!" - A benchmark of task affinity scores for joint learning. *CoRR*, abs/2301.02873, 2023.
- [4] Giulio Franzese, Simone Rossi, Dario Rossi, Markus Heinonen, Maurizio Filippone, and Pietro Michiardi. Continuous-time functional diffusion processes. *CoRR*, abs/2303.00800, 2023.
- [5] Ba-Hien Tran, Giulio Franzese, Pietro Michiardi, and Maurizio Filippone. One-line-of-code data mollification improves optimization of likelihood-based generative models. *CoRR*, abs/2305.18900, 2023.
- [6] Mustapha Bounoua, Giulio Franzese, and Pietro Michiardi. Multi-modal latent diffusion. *CoRR*, abs/2306.04445, 2023.
- [7] Julien Audibert, Pietro Michiardi, Frédéric Guyard, Sébastien Marti, and Maria A. Zuluaga. Do deep neural networks contribute to multivariate time series anomaly detection? *Pattern Recognit.*, 132:108945, 2022.
- [8] Giulio Franzese, Dimitrios Milios, Maurizio Filippone, and Pietro Michiardi. Revisiting the effects of stochasticity for hamiltonian samplers. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 6744–6778. PMLR, 2022.
- [9] Ugo Lecerf, Christelle Yemdji Tchassi, Sébastien Aubert, and Pietro Michiardi. Automatically learning fallback strategies with model-free reinforcement learning in safety-critical driving scenarios. In *ICMLT 2022: 7th International Conference on Machine Learning Technologies, Rome, Italy, March 11 - 13, 2022*, pages 209–215. ACM, 2022.
- [10] Julien Audibert, Pietro Michiardi, Frédéric Guyard, Sébastien Marti, and Maria A. Zuluaga. Do deep neural networks contribute to multivariate time series anomaly detection? *CoRR*, abs/2204.01637, 2022.
- [11] Ugo Lecerf, Christelle Yemdji Tchassi, Sébastien Aubert, and Pietro Michiardi.

- Automatically learning fallback strategies with model-free reinforcement learning in safety-critical driving scenarios. *CoRR*, abs/2204.05196, 2022.
- [12] Ugo Lecerf, Christelle Yemdji Tchassi, and Pietro Michiardi. Safer autonomous driving in a stochastic, partially-observable environment by hierarchical contingency planning. *CoRR*, abs/2204.06509, 2022.
- [13] Giulio Franzese, Simone Rossi, Lixuan Yang, Alessandro Finamore, Dario Rossi, Maurizio Filippone, and Pietro Michiardi. How much is enough? A study on diffusion times in score-based generative models. *CoRR*, abs/2206.05173, 2022.
- [14] Giulio Franzese, Dimitrios Milios, Maurizio Filippone, and Pietro Michiardi. A scalable bayesian sampling method based on stochastic gradient descent isotropization. *Entropy*, 23(11):1426, 2021.
- [15] Pietro Michiardi, Damiano Carra, and Sara Migliorini. Cache-based multi-query optimization for data-intensive scalable computing frameworks. *Inf. Syst. Frontiers*, 23(1):35–51, 2021.
- [16] Lucas Pascal, Pietro Michiardi, Xavier Bost, Benoit Huet, and Maria A. Zuluaga. Maximum roaming multi-task learning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 9331–9341. AAAI Press, 2021.
- [17] Raphaël Azorin, Massimo Gallo, Alessandro Finamore, Maurizio Filippone, Pietro Michiardi, and Dario Rossi. Towards a generic deep learning pipeline for traffic measurements. In Gareth Tyson, Hannaneh Barahouei Pasandi, and Lars C. Wolf, editors, *CoNEXT-SW '21: Proceedings of the CoNEXT Student Workshop, Virtual Event / Munich, Germany, 7 December 2021*, pages 5–6. ACM, 2021.
- [18] Graziano Mita, Maurizio Filippone, and Pietro Michiardi. An identifiable double VAE for disentangled representations. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 7769–7779. PMLR, 2021.
- [19] Gia-Lac Tran, Dimitrios Milios, Pietro Michiardi, and Maurizio Filippone. Sparse within sparse gaussian processes using neighbor information. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 10369–10378. PMLR, 2021.
- [20] Matthieu Da Silva-Filarder, Andrea Ancora, Maurizio Filippone, and Pietro Michiardi. Multimodal variational autoencoders for sensor fusion and cross generation. In M. Arif Wani, Ishwar K. Sethi, Weisong Shi, Guangzhi Qu, Daniela Stan Raicu, and Ruoming Jin, editors, *20th IEEE International Conference on Machine*

Learning and Applications, ICMLA 2021, Pasadena, CA, USA, December 13-16, 2021, pages 1069–1076. IEEE, 2021.

- [21] Ba-Hien Tran, Simone Rossi, Dimitrios Miliotis, Pietro Michiardi, Edwin V. Bonilla, and Maurizio Filippone. Model selection for bayesian autoencoders. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 19730–19742, 2021.
- [22] Giulio Franzese, Yiqing Yan, Giuseppe Serra, Ivan D’Onofrio, Raja Appuswamy, and Pietro Michiardi. Generative DNA: representation learning for dna-based approximate image storage. In *International Conference on Visual Communications and Image Processing, VCIP 2021, Munich, Germany, December 5-8, 2021*, pages 1–5. IEEE, 2021.
- [23] Ba-Hien Tran, Simone Rossi, Dimitrios Miliotis, Pietro Michiardi, Edwin V. Bonilla, and Maurizio Filippone. Model selection for bayesian autoencoders. *CoRR*, abs/2106.06245, 2021.
- [24] Giulio Franzese, Dimitrios Miliotis, Maurizio Filippone, and Pietro Michiardi. A unified view of stochastic hamiltonian sampling. *CoRR*, abs/2106.16200, 2021.
- [25] Lucas Pascal, Pietro Michiardi, Xavier Bost, Benoit Huet, and Maria A. Zuluaga. Optimization strategies in multi-task learning: Averaged or separated losses? *CoRR*, abs/2109.11678, 2021.
- [26] Rémi Domingues, Pietro Michiardi, Jérémie Barlet, and Maurizio Filippone. A comparative evaluation of novelty detection algorithms for discrete sequences. *Artif. Intell. Rev.*, 53(5):3787–3812, 2020.
- [27] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. Elastic provisioning of cloud caches: A cost-aware TTL approach. *IEEE/ACM Trans. Netw.*, 28(3):1283–1296, 2020.
- [28] Graziano Mita, Paolo Papotti, Maurizio Filippone, and Pietro Michiardi. LIBRE: learning interpretable boolean rule ensembles. In Silvia Chiappa and Roberto Calandra, editors, *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 245–255. PMLR, 2020.
- [29] Julien Audibert, Pietro Michiardi, Frédéric Guyard, Sébastien Marti, and Maria A. Zuluaga. USAD: unsupervised anomaly detection on multivariate time series. In Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash, editors, *KDD ’20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 3395–3404. ACM, 2020.

- [30] Rosa Candela, Pietro Michiardi, Maurizio Filippone, and Maria A. Zuluaga. Model monitoring and dynamic model selection in travel time-series forecasting. In Yuxiao Dong, Dunja Mladenic, and Craig Saunders, editors, *Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track - European Conference, ECML PKDD 2020, Ghent, Belgium, September 14-18, 2020, Proceedings, Part IV*, volume 12460 of *Lecture Notes in Computer Science*, pages 513–529. Springer, 2020.
- [31] Dimitrios Milios, Pietro Michiardi, and Maurizio Filippone. A variational view on bootstrap ensembles as bayesian inference. *CoRR*, abs/2006.04548, 2020.
- [32] Giulio Franzese, Rosa Candela, Dimitrios Milios, Maurizio Filippone, and Pietro Michiardi. Isotropic SGD: a practical approach to bayesian posterior sampling. *CoRR*, abs/2006.05087, 2020.
- [33] Lucas Pascal, Pietro Michiardi, Xavier Bost, Benoit Huet, and Maria A. Zuluaga. Maximum roaming multi-task learning. *CoRR*, abs/2006.09762, 2020.
- [34] Graziano Mita, Maurizio Filippone, and Pietro Michiardi. Learning optimal conditional priors for disentangled representations. *CoRR*, abs/2010.09360, 2020.
- [35] Gia-Lac Tran, Dimitrios Milios, Pietro Michiardi, and Maurizio Filippone. Sparse within sparse gaussian processes using neighbor information. *CoRR*, abs/2011.05041, 2020.
- [36] Marco Milanesio, Christian Callegari, and Pietro Michiardi. Network level perspective in web sessions troubleshooting. *Int. J. Commun. Syst.*, 32(6), 2019.
- [37] Damiano Carra and Pietro Michiardi. Memory partitioning and management in memcached. *IEEE Trans. Serv. Comput.*, 12(4):564–576, 2019.
- [38] Gia-Lac Tran, Edwin V. Bonilla, John P. Cunningham, Pietro Michiardi, and Maurizio Filippone. Calibrating deep convolutional gaussian processes. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pages 1554–1563. PMLR, 2019.
- [39] Francesco Pace, Dimitrios Milios, Damiano Carra, and Pietro Michiardi. Dynamic resource shaping for compute clusters. In Elisa Bertino, Carl K. Chang, Peter Chen, Ernesto Damiani, Michael Goul, and Katsunori Oyama, editors, *2019 IEEE International Congress on Big Data, BigData Congress 2019, Milan, Italy, July 8-13, 2019*, pages 45–54. IEEE, 2019.
- [40] Pietro Michiardi, Damiano Carra, and Sara Migliorini. In-memory caching for multi-query optimization of data-intensive scalable computing workloads. In Paolo Papotti, editor, *Proceedings of the Workshops of the EDBT/ICDT 2019 Joint*

Conference, EDBT/ICDT 2019, Lisbon, Portugal, March 26, 2019, volume 2322 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2019.

- [41] Simone Rossi, Pietro Michiardi, and Maurizio Filippone. Good initializations of variational bayes for deep models. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5487–5497. PMLR, 2019.
- [42] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. Ttl-based cloud caches. In *2019 IEEE Conference on Computer Communications, INFOCOM 2019, Paris, France, April 29 - May 2, 2019*, pages 685–693. IEEE, 2019.
- [43] Duc-Trung Nguyen, Maurizio Filippone, and Pietro Michiardi. Exact gaussian process regression with distributed computations. In Chih-Cheng Hung and George A. Papadopoulos, editors, *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019, Limassol, Cyprus, April 8-12, 2019*, pages 1286–1295. ACM, 2019.
- [44] Remi Domingues, Pietro Michiardi, Jérémie Barlet, and Maurizio Filippone. A comparative evaluation of novelty detection algorithms for discrete sequences. *CoRR*, abs/1902.10940, 2019.
- [45] Rosa Candela, Giulio Franzese, Maurizio Filippone, and Pietro Michiardi. Sparsification as a remedy for staleness in distributed asynchronous SGD. *CoRR*, abs/1910.09466, 2019.
- [46] Graziano Mita, Paolo Papotti, Maurizio Filippone, and Pietro Michiardi. LIBRE: learning interpretable boolean rule ensembles. *CoRR*, abs/1911.06537, 2019.
- [47] Remi Domingues, Pietro Michiardi, Jihane Zouaoui, and Maurizio Filippone. Deep gaussian process autoencoders for novelty detection. *Mach. Learn.*, 107(8-10):1363–1383, 2018.
- [48] Remi Domingues, Maurizio Filippone, Pietro Michiardi, and Jihane Zouaoui. A comparative evaluation of outlier detection algorithms: Experiments and analyses. *Pattern Recognit.*, 74:406–421, 2018.
- [49] Giovanni Neglia, Damiano Carra, and Pietro Michiardi. Cache policies for linear utility maximization. *IEEE/ACM Trans. Netw.*, 26(1):302–313, 2018.
- [50] Gil Vernik, Michael Factor, Elliot K. Kolodner, Pietro Michiardi, Effi Ofer, and Francesco Pace. Stocator: Providing high performance and fault tolerance for apache spark over object storage. In Esam El-Araby, Dhableswar K. Panda, Sandra Gesing, Amy W. Apon, Volodymyr V. Kindratenko, Massimo Cafaro, and Alfredo Cuzzocrea, editors, *18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGRID 2018, Washington, DC, USA, May 1-4, 2018*, pages 462–471. IEEE Computer Society, 2018.

- [51] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. Elastic provisioning of cloud caches: a cost-aware TTL approach. In *Proceedings of the ACM Symposium on Cloud Computing, SoCC 2018, Carlsbad, CA, USA, October 11-13, 2018*, page 526. ACM, 2018.
- [52] Francesco Pace, Dimitrios Milios, Damiano Carra, Daniele Venzano, and Pietro Michiardi. Data-driven resource shaping for compute clusters. In *Proceedings of the ACM Symposium on Cloud Computing, SoCC 2018, Carlsbad, CA, USA, October 11-13, 2018*, page 527. ACM, 2018.
- [53] Dimitrios Milios, Raffaello Camoriano, Pietro Michiardi, Lorenzo Rosasco, and Maurizio Filippone. Dirichlet-based gaussian processes for large-scale calibrated classification. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6008–6018, 2018.
- [54] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. Elastic provisioning of cloud caches: a cost-aware TTL approach. *CoRR*, abs/1802.04696, 2018.
- [55] Pietro Michiardi, Damiano Carra, and Sara Migliorini. Cache-based multi-query optimization for data-intensive scalable computing frameworks. *CoRR*, abs/1805.08650, 2018.
- [56] Gia-Lac Tran, Edwin V. Bonilla, John P. Cunningham, Pietro Michiardi, and Maurizio Filippone. Calibrating deep convolutional gaussian processes. *CoRR*, abs/1805.10522, 2018.
- [57] Dimitrios Milios, Raffaello Camoriano, Pietro Michiardi, Lorenzo Rosasco, and Maurizio Filippone. Dirichlet-based gaussian processes for large-scale calibrated classification. *CoRR*, abs/1805.10915, 2018.
- [58] Francesco Pace, Dimitrios Milios, Damiano Carra, Daniele Venzano, and Pietro Michiardi. A data-driven approach to dynamically adjust resource allocation for compute clusters. *CoRR*, abs/1807.00368, 2018.
- [59] Simone Rossi, Pietro Michiardi, and Maurizio Filippone. Good initializations of variational bayes for deep models. *CoRR*, abs/1810.08083, 2018.
- [60] Daniele Apiletti, Elena Baralis, Tania Cerquitelli, Paolo Garza, Fabio Pulvirenti, and Pietro Michiardi. A parallel mapreduce algorithm to efficiently support itemset mining on high dimensional data. *Big Data Res.*, 10:53–69, 2017.
- [61] Alessandro Lulli, Lorenzo Gabrielli, Patrizio Dazzi, Matteo Dell’Amico, Pietro Michiardi, Mirco Nanni, and Laura Ricci. Scalable and flexible clustering solutions for mobile phone-based population indicators. *Int. J. Data Sci. Anal.*, 4(4):285–299, 2017.

- [62] Yongchao Tian, Ioannis Alagiannis, Erietta Liarou, Anastasia Ailamaki, Pietro Michiardi, and Marko Vukolic. Dinodb: An interactive-speed query engine for ad-hoc queries on temporary data. *IEEE Trans. Big Data*, 3(3):320–333, 2017.
- [63] Mario Pastorelli, Damiano Carra, Matteo Dell’Amico, and Pietro Michiardi. HFSP: bringing size-based scheduling to hadoop. *IEEE Trans. Cloud Comput.*, 5(1):43–56, 2017.
- [64] Giovanni Neglia, Damiano Carra, Mingdong Feng, Vaishnav Janardhan, Pietro Michiardi, and Dimitra Tsiggari. Access-time-aware cache algorithms. *ACM Trans. Model. Perform. Evaluation Comput. Syst.*, 2(4):21:1–21:29, 2017.
- [65] Yongchao Tian, Pietro Michiardi, and Marko Vukolic. Bleach: A distributed stream data cleaning system. In George Karypis and Jia Zhang, editors, *2017 IEEE International Congress on Big Data, BigData Congress 2017, Honolulu, HI, USA, June 25-30, 2017*, pages 113–120. IEEE Computer Society, 2017.
- [66] Son-Hai Ha, Patrick Brown, and Pietro Michiardi. Resource management for parallel processing frameworks with load awareness at worker side. In George Karypis and Jia Zhang, editors, *2017 IEEE International Congress on Big Data, BigData Congress 2017, Honolulu, HI, USA, June 25-30, 2017*, pages 161–168. IEEE Computer Society, 2017.
- [67] Francesco Pace, Daniele Venzano, Damiano Carra, and Pietro Michiardi. Flexible scheduling of distributed analytic applications. In *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGRID 2017, Madrid, Spain, May 14-17, 2017*, pages 100–109. IEEE Computer Society / ACM, 2017.
- [68] Gil Vernik, Michael Factor, Elliot K. Kolodner, Effi Ofer, Pietro Michiardi, and Francesco Pace. Stocator: an object store aware connector for apache spark. In *Proceedings of the 2017 Symposium on Cloud Computing, SoCC 2017, Santa Clara, CA, USA, September 24-27, 2017*, page 653. ACM, 2017.
- [69] Yosef Moatti, Eran Rom, Raúl Gracia Tinedo, Dalit Naor, Doron Chen, Josep Sampé, Marc Sánchez Artigas, Pedro García López, Filip Gluszak, Eric Deschdt, Francesco Pace, Daniele Venzano, and Pietro Michiardi. Too big to eat: Boosting analytics data ingestion from object stores with scoop. In *33rd IEEE International Conference on Data Engineering, ICDE 2017, San Diego, CA, USA, April 19-22, 2017*, pages 309–320. IEEE Computer Society, 2017.
- [70] Kurt Cutajar, Edwin V. Bonilla, Pietro Michiardi, and Maurizio Filippone. Random feature expansions for deep gaussian processes. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 884–893. PMLR, 2017.
- [71] Giovanni Neglia, Damiano Carra, and Pietro Michiardi. Cache policies for linear

- utility maximization. In *2017 IEEE Conference on Computer Communications, INFOCOM 2017, Atlanta, GA, USA, May 1-4, 2017*, pages 1–9. IEEE, 2017.
- [72] Gil Vernik, Michael Factor, Elliot K. Kolodner, Effi Ofer, Pietro Michiardi, and Francesco Pace. Stocator: a high performance object store connector for spark. In Doron Chen, Peter Desnoyers, and Eyal de Lara, editors, *Proceedings of the 10th ACM International Systems and Storage Conference, SYSTOR 2017, Haifa, Israel, May 22-24, 2017*, page 27:1. ACM, 2017.
- [73] Gil Vernik, Michael Factor, Elliot K. Kolodner, Pietro Michiardi, Effi Ofer, and Francesco Pace. Stocator: A high performance object store connector for spark. *CoRR*, abs/1709.01812, 2017.
- [74] Raúl Gracia Tinedo, Pedro García López, Marc Sánchez Artigas, Josep Sampé, Yosef Moatti, Eran Rom, Dalit Naor, Ramon Nou, Toni Cortes, William Oppermann, and Pietro Michiardi. Iostack: Software-defined object storage. *IEEE Internet Comput.*, 20(3):10–18, 2016.
- [75] Alessandro Lulli, Matteo Dell’Amico, Pietro Michiardi, and Laura Ricci. NG-DBSCAN: scalable density-based clustering for arbitrary data. *Proc. VLDB Endow.*, 10(3):157–168, 2016.
- [76] Matteo Dell’Amico, Damiano Carra, and Pietro Michiardi. PSBS: practical size-based scheduling. *IEEE Trans. Computers*, 65(7):2199–2212, 2016.
- [77] Francesco Pace, Marco Milanese, Daniele Venzano, Damiano Carra, and Pietro Michiardi. Experimental performance evaluation of cloud-based analytics-as-a-service. In *9th IEEE International Conference on Cloud Computing, CLOUD 2016, San Francisco, CA, USA, June 27 - July 2, 2016*, pages 196–203. IEEE Computer Society, 2016.
- [78] Thibault Debatty, Fabio Pulvirenti, Pietro Michiardi, and Wim Mees. Fast distributed k-nn graph update. In James Joshi, George Karypis, Ling Liu, Xiaohua Hu, Ronay Ak, Yinglong Xia, Weijia Xu, Aki-Hiro Sato, Sudarsan Rachuri, Lyle H. Ungar, Philip S. Yu, Rama Govindaraju, and Toyotaro Suzumura, editors, *2016 IEEE International Conference on Big Data (IEEE BigData 2016), Washington DC, USA, December 5-8, 2016*, pages 3308–3317. IEEE Computer Society, 2016.
- [79] Duy-Hung Phan and Pietro Michiardi. A novel, low-latency algorithm for multiple group-by query optimization. In *32nd IEEE International Conference on Data Engineering, ICDE 2016, Helsinki, Finland, May 16-20, 2016*, pages 301–312. IEEE Computer Society, 2016.
- [80] Alessandro Lulli, Lorenzo Gabrielli, Patrizio Dazzi, Matteo Dell’Amico, Pietro Michiardi, Mirco Nanni, and Laura Ricci. Improving population estimation from mobile calls: A clustering approach. In *IEEE Symposium on Computers and Communication, ISCC 2016, Messina, Italy, June 27-30, 2016*, pages 1097–1102. IEEE Computer Society, 2016.

- [81] Giovanni Neglia, Damiano Carra, Mingdong Feng, Vaishnav Janardhan, Pietro Michiardi, and Dimitra Tsigkari. Access-time aware cache algorithms. In Tobias Hoßfeld, Brian L. Mark, S.-H. Gary Chan, and Andreas Timm-Giel, editors, *28th International Teletraffic Congress, ITC 2016, Würzburg, Germany, September 12-16, 2016*, pages 148–156. IEEE, 2016.
- [82] Mirjana Ivanovic, Bernhard Thalheim, Barbara Catania, Klaus-Dieter Schewe, Marite Kirikova, Petr Saloun, Ajantha Dahanayake, Tania Cerquitelli, Elena Baralis, and Pietro Michiardi, editors. *New Trends in Databases and Information Systems - ADBIS 2016 Short Papers and Workshops, BigDap, DCSA, DC, Prague, Czech Republic, August 28-31, 2016, Proceedings*, volume 637 of *Communications in Computer and Information Science*. Springer, 2016.
- [83] Thibault Debatty, Pietro Michiardi, and Wim Mees. Fast online k-nn graph building. *CoRR*, abs/1602.06819, 2016.
- [84] Francesco Pace, Marco Milanese, Daniele Venzano, Damiano Carra, and Pietro Michiardi. Experimental performance evaluation of cloud-based analytics-as-a-service. *CoRR*, abs/1602.07919, 2016.
- [85] Yongchao Tian, Ioannis Alagiannis, Erietta Liarou, Anastasia Ailamaki, Pietro Michiardi, and Marko Vukolic. Dinodb: an interactive-speed query engine for ad-hoc queries on temporary data. *CoRR*, abs/1609.05096, 2016.
- [86] Yongchao Tian, Pietro Michiardi, and Marko Vukolic. Bleach: A distributed stream data cleaning system. *CoRR*, abs/1609.05113, 2016.
- [87] Francesco Pace, Daniele Venzano, Damiano Carra, and Pietro Michiardi. Flexible scheduling of distributed analytic applications. *CoRR*, abs/1611.09528, 2016.
- [88] Matteo Dell’Amico, Pietro Michiardi, László Toka, and Pasquale Cataldi. Adaptive redundancy management for durable P2P backup. *Comput. Networks*, 83:136–148, 2015.
- [89] Matteo Dell’Amico, Maurizio Filippone, Pietro Michiardi, and Yves Roudier. On user availability prediction and network applications. *IEEE/ACM Trans. Netw.*, 23(4):1300–1313, 2015.
- [90] Duy-Hung Phan, Quang-Nhat Hoang-Xuan, Matteo Dell’Amico, and Pietro Michiardi. Efficient and self-balanced ROLLUP aggregates for large-scale data summarization. In Barbara Carminati and Latifur Khan, editors, *2015 IEEE International Congress on Big Data, New York City, NY, USA, June 27 - July 2, 2015*, pages 158–165. IEEE Computer Society, 2015.
- [91] Alessandro Lulli, Thibault Debatty, Matteo Dell’Amico, Pietro Michiardi, and Laura Ricci. Scalable k-nn based text clustering. In *2015 IEEE International Conference on Big Data (IEEE BigData 2015), Santa Clara, CA, USA, October 29 - November 1, 2015*, pages 958–963. IEEE Computer Society, 2015.

- [92] Daniele Apiletti, Elena Baralis, Tania Cerquitelli, Paolo Garza, Pietro Michiardi, and Fabio Pulvirenti. Pampa-hd: A parallel mapreduce-based frequent pattern miner for high-dimensional data. In *IEEE International Conference on Data Mining Workshop, ICDMW 2015, Atlantic City, NJ, USA, November 14-17, 2015*, pages 839–846. IEEE Computer Society, 2015.
- [93] Christian Callegari, Marco Milanesio, and Pietro Michiardi. Troubleshooting web sessions with CUSUM. In *International Wireless Communications and Mobile Computing Conference, IWCMC 2015, Dubrovnik, Croatia, August 24-28, 2015*, pages 385–390. IEEE, 2015.
- [94] Damiano Carra and Pietro Michiardi. Cost-based memory partitioning and management in memcached. In *Proceedings of the 3rd VLDB Workshop on In-Memory Data Management and Analytics, IMDM@VLDB 2015, Kohala Coast, HI, USA, August 31, 2015*, pages 6:1–6:8. ACM, 2015.
- [95] Matteo Dell’Amico, Damiano Carra, and Pietro Michiardi. On fair size-based scheduling. *CoRR*, abs/1506.09158, 2015.
- [96] Thibault Debatty, Pietro Michiardi, Olivier Thonnard, and Wim Mees. Building k-nn graphs from large text data. In Jimmy Lin, Jian Pei, Xiaohua Hu, Wo Chang, Raghunath Nambiar, Charu C. Aggarwal, Nick Cercone, Vasant G. Honavar, Jun Huan, Bamshad Mobasher, and Saumyadipta Pyne, editors, *2014 IEEE International Conference on Big Data (IEEE BigData 2014), Washington, DC, USA, October 27-30, 2014*, pages 573–578. IEEE Computer Society, 2014.
- [97] Yufei Han, Xiaolan Sha, Etta Grover-Silva, and Pietro Michiardi. On the impact of socio-economic factors on power load forecasting. In Jimmy Lin, Jian Pei, Xiaohua Hu, Wo Chang, Raghunath Nambiar, Charu C. Aggarwal, Nick Cercone, Vasant G. Honavar, Jun Huan, Bamshad Mobasher, and Saumyadipta Pyne, editors, *2014 IEEE International Conference on Big Data (IEEE BigData 2014), Washington, DC, USA, October 27-30, 2014*, pages 742–747. IEEE Computer Society, 2014.
- [98] Duy-Hung Phan, Matteo Dell’Amico, and Pietro Michiardi. On the design space of mapreduce ROLLUP aggregates. In K. Selçuk Candan, Sihem Amer-Yahia, Nicole Schweikardt, Vassilis Christophides, and Vincent Leroy, editors, *Proceedings of the Workshops of the EDBT/ICDT 2014 Joint Conference (EDBT/ICDT 2014), Athens, Greece, March 28, 2014*, volume 1133 of *CEUR Workshop Proceedings*, pages 10–18. CEUR-WS.org, 2014.
- [99] Thibault Debatty, Pietro Michiardi, Wim Mees, and Olivier Thonnard. Determining the k in k-means with mapreduce. In K. Selçuk Candan, Sihem Amer-Yahia, Nicole Schweikardt, Vassilis Christophides, and Vincent Leroy, editors, *Proceedings of the Workshops of the EDBT/ICDT 2014 Joint Conference (EDBT/ICDT 2014), Athens, Greece, March 28, 2014*, volume 1133 of *CEUR Workshop Proceedings*, pages 19–28. CEUR-WS.org, 2014.

- [100] Damiano Carra and Pietro Michiardi. Memory partitioning in memcached: An experimental performance analysis. In *IEEE International Conference on Communications, ICC 2014, Sydney, Australia, June 10-14, 2014*, pages 1154–1159. IEEE, 2014.
- [101] Mario Pastorelli, Matteo Dell’Amico, and Pietro Michiardi. Os-assisted task pre-emption for hadoop. In *34th International Conference on Distributed Computing Systems Workshops (ICDCS 2014 Workshops), Madrid, Spain, June 30 - July 3, 2014*, pages 94–99. IEEE Computer Society, 2014.
- [102] Thibault Debatty, Pietro Michiardi, Olivier Thonnard, and Wim Mees. Scalable graph building from text data. In Wei Fan, Albert Bifet, Qiang Yang, and Philip S. Yu, editors, *Proceedings of the 3rd International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications, BigMine 2014, New York City, USA, August 24, 2014*, volume 36 of *JMLR Workshop and Conference Proceedings*, pages 120–132. JMLR.org, 2014.
- [103] Matteo Dell’Amico, Damiano Carra, Mario Pastorelli, and Pietro Michiardi. Revisiting size-based scheduling with estimated job sizes. In *IEEE 22nd International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems, MASCOTS 2014, Paris, France, September 9-11, 2014*, pages 411–420. IEEE Computer Society, 2014.
- [104] Yongchao Tian, Ioannis Alagiannis, Erietta Liarou, Anastasia Ailamaki, Pietro Michiardi, and Marko Vukolic. Dinodb: Efficient large-scale raw data analytics. In Rada Chirkova and Jun Yang, editors, *Proceedings of the First International Workshop on Bringing the Value of "Big Data" to Users, Data4U@VLDB 2014, Hangzhou, China, September 1, 2014*, page 1. ACM, 2014.
- [105] Mario Pastorelli, Matteo Dell’Amico, and Pietro Michiardi. Os-assisted task pre-emption for hadoop. *CoRR*, abs/1402.2107, 2014.
- [106] Matteo Dell’Amico, Damiano Carra, Mario Pastorelli, and Pietro Michiardi. Revisiting size-based scheduling with estimated job sizes. *CoRR*, abs/1403.5996, 2014.
- [107] Matteo Dell’Amico, Maurizio Filippone, Pietro Michiardi, and Yves Roudier. On user availability prediction and network applications. *CoRR*, abs/1404.7688, 2014.
- [108] Matteo Dell’Amico, Damiano Carra, and Pietro Michiardi. PSBS: practical size-based scheduling. *CoRR*, abs/1410.6122, 2014.
- [109] Damiano Carra, Pietro Michiardi, Hani Salah, and Thorsten Strufe. On the impact of incentives in emule {Analysis and Measurements of a Popular File-Sharing Application}. *IEEE J. Sel. Areas Commun.*, 31(9-Supplement):94–104, 2013.
- [110] Damiano Carra, Moritz Steiner, Pietro Michiardi, Ernst W. Biersack, Wolfgang

- Effelsberg, and Taoufik En-Najjary. Characterization and management of popular content in KAD. *IEEE Trans. Parallel Distributed Syst.*, 24(4):662–671, 2013.
- [111] Mario Pastorelli, Antonio Barbuzzi, Damiano Carra, Matteo Dell’Amico, and Pietro Michiardi. HFSP: size-based scheduling for hadoop. In Xiaohua Hu, Tsau Young Lin, Vijay V. Raghavan, Benjamin W. Wah, Ricardo Baeza-Yates, Geoffrey C. Fox, Cyrus Shahabi, Matthew Smith, Qiang Yang, Rayid Ghani, Wei Fan, Ronny Lempel, and Raghunath Nambiar, editors, *2013 IEEE International Conference on Big Data (IEEE BigData 2013), 6-9 October 2013, Santa Clara, CA, USA*, pages 51–59. IEEE Computer Society, 2013.
- [112] Xiaolan Sha, Daniele Quercia, Matteo Dell’Amico, and Pietro Michiardi. Trend makers and trend spotters in a mobile application. In Amy S. Bruckman, Scott Counts, Cliff Lampe, and Loren G. Terveen, editors, *Computer Supported Cooperative Work, CSCW 2013, San Antonio, TX, USA, February 23-27, 2013*, pages 1365–1374. ACM, 2013.
- [113] Daniele Venzano and Pietro Michiardi. A measurement study of data-intensive network traffic patterns in a private cloud. In *IEEE/ACM 6th International Conference on Utility and Cloud Computing, UCC 2013, Dresden, Germany, December 9-12, 2013*, pages 476–481. IEEE Computer Society, 2013.
- [114] Mario Pastorelli, Antonio Barbuzzi, Damiano Carra, and Pietro Michiardi. HFSP: the hadoop fair sojourn protocol. *CoRR*, abs/1302.2749, 2013.
- [115] Pietro Michiardi, Damiano Carra, Francesco Albanese, and Azer Bestavros. Peer-assisted content distribution on a budget. *Comput. Networks*, 56(7):2038–2048, 2012.
- [116] Chi-Anh La, Pietro Michiardi, Claudio Casetti, Carla-Fabiana Chiasserini, and Marco Fiore. Content replication in mobile networks. *IEEE J. Sel. Areas Commun.*, 30(9):1762–1770, 2012.
- [117] László Toka, Pasquale Cataldi, Matteo Dell’Amico, and Pietro Michiardi. Redundancy management for P2P backup. In Albert G. Greenberg and Kazem Sohraby, editors, *Proceedings of the IEEE INFOCOM 2012, Orlando, FL, USA, March 25-30, 2012*, pages 2986–2990. IEEE, 2012.
- [118] Thomas Mager, Ernst W. Biersack, and Pietro Michiardi. A measurement study of the wuala on-line storage service. In *12th IEEE International Conference on Peer-to-Peer Computing, P2P 2012, Tarragona, Spain, September 3-5, 2012*, pages 237–248. IEEE, 2012.
- [119] Tobias Lauinger, Engin Kirda, and Pietro Michiardi. Paying for piracy? an analysis of one-click hosters’ controversial reward schemes. In Davide Balzarotti, Salvatore J. Stolfo, and Marco Cova, editors, *Research in Attacks, Intrusions, and Defenses - 15th International Symposium, RAID 2012, Amsterdam, The Netherlands*,

September 12-14, 2012. *Proceedings*, volume 7462 of *Lecture Notes in Computer Science*, pages 169–189. Springer, 2012.

- [120] Xiaolan Sha, Daniele Quercia, Pietro Michiardi, and Matteo Dell’Amico. Spotting trends: the wisdom of the few. In Pdraig Cunningham, Neil J. Hurley, Ido Guy, and Sarabjot Singh Anand, editors, *Sixth ACM Conference on Recommender Systems, RecSys ’12, Dublin, Ireland, September 9-13, 2012*, pages 51–58. ACM, 2012.
- [121] László Toka, Pasquale Cataldi, Matteo Dell’Amico, and Pietro Michiardi. Redundancy management for P2P backup. *CoRR*, abs/1201.2360, 2012.
- [122] Flavio Esposito, Ibrahim Matta, Debajyoti Bera, and Pietro Michiardi. On the impact of seed scheduling in peer-to-peer networks. *Comput. Networks*, 55(15):3303–3317, 2011.
- [123] Daishi Kato, Kaoutar Elkhyaoui, Kazuo Kunieda, Keiji Yamada, and Pietro Michiardi. A scalable interest-oriented peer-to-peer pub/sub network. *Peer-to-Peer Netw. Appl.*, 4(2):165–177, 2011.
- [124] László Toka and Pietro Michiardi. Analysis of user-driven peer selection in peer-to-peer backup and storage systems. *Telecommun. Syst.*, 47(1-2):49–63, 2011.
- [125] Damiano Carra, Giovanni Neglia, Pietro Michiardi, and Francesco Albanese. On the robustness of bittorrent swarms to greedy peers. *IEEE Trans. Parallel Distributed Syst.*, 22(12):2071–2078, 2011.
- [126] Damiano Carra, Moritz Steiner, and Pietro Michiardi. Adaptive load balancing in KAD. In Tohru Asami and Teruo Higashino, editors, *2011 IEEE International Conference on Peer-to-Peer Computing, P2P 2011, Kyoto, Japan, August 31 - September 2, 2011*, pages 92–101. IEEE, 2011.
- [127] László Toka, Matteo Dell’Amico, and Pietro Michiardi. Data transfer scheduling for P2P storage. In Tohru Asami and Teruo Higashino, editors, *2011 IEEE International Conference on Peer-to-Peer Computing, P2P 2011, Kyoto, Japan, August 31 - September 2, 2011*, pages 132–141. IEEE, 2011.
- [128] Rajesh Sharma, Anwitaman Datta, Matteo Dell’Amico, and Pietro Michiardi. An empirical study of availability in friend-to-friend storage systems. In Tohru Asami and Teruo Higashino, editors, *2011 IEEE International Conference on Peer-to-Peer Computing, P2P 2011, Kyoto, Japan, August 31 - September 2, 2011*, pages 348–351. IEEE, 2011.
- [129] Chi-Anh La, Pietro Michiardi, Claudio Casetti, Carla-Fabiana Chiasserini, and Marco Fiore. Content replication and placement in mobile networks. *CoRR*, abs/1102.3013, 2011.
- [130] Georgios Smaragdakis, Nikolaos Laoutaris, Pietro Michiardi, Azer Bestavros, John W. Byers, and Mema Roussopoulos. Distributed network formation for n-way

- broadcast applications. *IEEE Trans. Parallel Distributed Syst.*, 21(10):1427–1441, 2010.
- [131] Matteo Dell’Amico, Pietro Michiardi, and Yves Roudier. Password strength: An empirical analysis. In *INFOCOM 2010. 29th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 15-19 March 2010, San Diego, CA, USA*, pages 983–991. IEEE, 2010.
- [132] Antonio Barbuzzi, Pietro Michiardi, Ernst W. Biersack, and Gennaro Boggia. Parallel bulk insertion for large-scale analytics applications. In Gregory V. Chockler, Ymir Vigfusson, Marcos K. Aguilera, and Marc Shapiro, editors, *Proceedings of the 4th International Workshop on Large Scale Distributed Systems and Middleware, LADIS ’10, Zürich, Switzerland, July 28-29, 2010*, pages 27–31. ACM, 2010.
- [133] László Toka, Matteo Dell’Amico, and Pietro Michiardi. Online data backup: A peer-assisted approach. In *IEEE Tenth International Conference on Peer-to-Peer Computing, P2P 2010, Delft, The Netherlands, 25-27 August 2010*, pages 1–10. IEEE, 2010.
- [134] Chi-Anh La, Pietro Michiardi, Claudio Casetti, Carla-Fabiana Chiasserini, and Marco Fiore. A lightweight distributed solution to content replication in mobile networks. In *2010 IEEE Wireless Communications and Networking Conference, WCNC 2010, Proceedings, Sydney, Australia, 18-21 April 2010*, pages 1–6. IEEE, 2010.
- [135] László Toka, Matteo Dell’Amico, and Pietro Michiardi. On scheduling and redundancy for P2P backup. *CoRR*, abs/1009.1344, 2010.
- [136] Matteo Dell’Amico, Pietro Michiardi, and Yves Roudier. Back to the future: On predicting user uptime. *CoRR*, abs/1010.0626, 2010.
- [137] Roberto Di Pietro, Pietro Michiardi, and Refik Molva. Confidentiality and integrity for data aggregation in WSN using peer monitoring. *Secur. Commun. Networks*, 2(2):181–194, 2009.
- [138] Pietro Michiardi and László Toka. Selfish neighbor selection in peer-to-peer backup and storage applications. In Henk J. Sips, Dick H. J. Epema, and Hai-Xiang Lin, editors, *Euro-Par 2009 Parallel Processing, 15th International Euro-Par Conference, Delft, The Netherlands, August 25-28, 2009. Proceedings*, volume 5704 of *Lecture Notes in Computer Science*, pages 548–560. Springer, 2009.
- [139] Claudio Casetti, Carla-Fabiana Chiasserini, Marco Fiore, Chi-Anh La, and Pietro Michiardi. P2P cache-and-forward mechanisms for mobile ad hoc networks. In *Proceedings of the 14th IEEE Symposium on Computers and Communications (ISCC 2009), July 5-8, Sousse, Tunisia*, pages 386–392. IEEE Computer Society, 2009.
- [140] Flavio Esposito, Ibrahim Matta, Pietro Michiardi, Nobuyuki Mitsutake, and Dami-

- ano Carra. Seed scheduling for peer-to-peer networks. In *Proceedings of The Eighth IEEE International Symposium on Networking Computing and Applications, NCA 2009, July 9-11, 2009, Cambridge, Massachusetts, USA*, pages 28–35. IEEE Computer Society, 2009.
- [141] Kaoutar Elkhiyaoui, Daishi Kato, Kazuo Kunieda, Keiji Yamada, and Pietro Michiardi. A scalable interest-oriented peer-to-peer pub/sub network. In Henning Schulzrinne, Karl Aberer, and Anwitaman Datta, editors, *Proceedings P2P 2009, Ninth International Conference on Peer-to-Peer Computing, 9-11 September 2009, Seattle, Washington, USA*, pages 204–211. IEEE, 2009.
- [142] Pietro Michiardi, Carla-Fabiana Chiasserini, Claudio Casetti, Chi-Anh La, and Marco Fiore. On a selfish caching game. In Srikanta Tirthapura and Lorenzo Alvisi, editors, *Proceedings of the 28th Annual ACM Symposium on Principles of Distributed Computing, PODC 2009, Calgary, Alberta, Canada, August 10-12, 2009*, pages 284–285. ACM, 2009.
- [143] Claudio Casetti, Carla-Fabiana Chiasserini, Marco Fiore, Chi-Anh La, and Pietro Michiardi. A holistic approach to information distribution in ad hoc networks. *CoRR*, abs/0901.1782, 2009.
- [144] Matteo Dell’Amico, Pietro Michiardi, and Yves Roudier. Measuring password strength: An empirical analysis. *CoRR*, abs/0907.3402, 2009.
- [145] Chi-Anh La, Pietro Michiardi, Claudio Casetti, Carla-Fabiana Chiasserini, and Marco Fiore. A lightweight distributed solution to content replication in mobile networks. *CoRR*, abs/0909.2024, 2009.
- [146] Nikolaos Laoutaris, Damiano Carra, and Pietro Michiardi. Uplink allocation beyond choke/unchoke: or how to divide and conquer best. In Arturo Azcorra, Gustavo de Veciana, Keith W. Ross, and Leandros Tassiulas, editors, *Proceedings of the 2008 ACM Conference on Emerging Network Experiment and Technology, CoNEXT 2008, Madrid, Spain, December 9-12, 2008*, page 18. ACM, 2008.
- [147] Georgios Smaragdakis, Azer Bestavros, Nikolaos Laoutaris, John W. Byers, Pietro Michiardi, and Mema Roussopoulos. Swarming on optimized graphs for n-way broadcast. In *INFOCOM 2008. 27th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 13-18 April 2008, Phoenix, AZ, USA*, pages 141–145. IEEE, 2008.
- [148] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. On the impact of greedy strategies in bittorrent networks: The case of bittyrant. In Klaus Wehrle, Wolfgang Kellerer, Sandeep K. Singhal, and Ralf Steinmetz, editors, *Proceedings P2P’08, Eighth International Conference on Peer-to-Peer Computing, 8-11 September 2008, Aachen, Germany*, pages 311–320. IEEE Computer Society, 2008.
- [149] László Toka and Pietro Michiardi. A dynamic exchange game. In Rida A. Bazzi and Boaz Patt-Shamir, editors, *Proceedings of the Twenty-Seventh Annual ACM Sym-*

- posium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, page 428. ACM, 2008.
- [150] Roberto Di Pietro and Pietro Michiardi. Gossip-based aggregate computation: computing faster with non address-oblivious schemes. In Rida A. Bazzi and Boaz Patt-Shamir, editors, *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, page 442. ACM, 2008.
- [151] László Toka and Pietro Michiardi. Analysis of user-driven peer selection in peer-to-peer backup and storage systems. In John S. Baras and Costas Courcoubetis, editors, *3rd International ICST Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2008, Athens, Greece, October 20-24, 2008*, page 7. ICST/ACM, 2008.
- [152] Chi-Anh La and Pietro Michiardi. Characterizing user mobility in second life. In Christos Faloutsos, Thomas Karagiannis, and Pablo Rodriguez, editors, *Proceedings of the first Workshop on Online Social Networks, WOSN 2008, Seattle, WA, USA, August 17-22, 2008*, pages 79–84. ACM, 2008.
- [153] Chi-Anh La and Pietro Michiardi. Characterizing user mobility in second life. *CoRR*, abs/0804.1928, 2008.
- [154] Pietro Michiardi, Krishna K. Ramachandran, and Biplab Sikdar. Modeling seed scheduling strategies in bittorrent. In Ian F. Akyildiz, Raghupathy Sivakumar, Eylem Ekici, Jaudelice Cavalcante de Oliveira, and Janise McNair, editors, *NETWORKING 2007. Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet, 6th International IFIP-TC6 Networking Conference, Atlanta, GA, USA, May 14-18, 2007, Proceedings*, volume 4479 of *Lecture Notes in Computer Science*, pages 606–616. Springer, 2007.
- [155] Pietro Michiardi, Paul Marrow, Richard Tateson, and Fabrice Saffre. Aggregation dynamics in service overlay networks. In *Proceedings of the First International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2007, Boston, MA, USA, July 9-11, 2007*, pages 129–140. IEEE Computer Society, 2007.
- [156] Nouha Oualha, Pietro Michiardi, and Yves Roudier. A game theoretic model of a protocol for data possession verification. In *2007 International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2007), 18-21 June 2007, Helsinki, Finland, Proceedings*, pages 1–6. IEEE Computer Society, 2007.
- [157] Arnaud Legout, Guillaume Urvoy-Keller, and Pietro Michiardi. Rarest first and choke algorithms are enough. In Jussara M. Almeida, Virgílio A. F. Almeida, and Paul Barford, editors, *Proceedings of the 6th ACM SIGCOMM Internet Measurement Conference, IMC 2006, Rio de Janeiro, Brazil, October 25-27, 2006*, pages 203–216. ACM, 2006.

- [158] Guillaume Urvoy-Keller and Pietro Michiardi. Impact of inner parameters and overlay structure on the performance of bittorrent. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 23-29 April 2006, Barcelona, Catalunya, Spain*. IEEE, 2006.
- [159] Pietro Michiardi and Refik Molva. Identity based message authentication for dynamic networks. In Simone Fischer-Hübner, Kai Rannenberg, Louise Yngström, and Stefan Lindskog, editors, *Security and Privacy in Dynamic Environments, Proceedings of the IFIP TC-11 21st International Information Security Conference (SEC 2006), 22-24 May 2006, Karlstad, Sweden*, volume 201 of *IFIP*, pages 50–61. Springer, 2006.
- [160] Walid Bagga, Stefano Crosta, Pietro Michiardi, and Refik Molva. Establishment of ad-hoc communities through policy-based cryptography. In Giovanni Di Crescenzo and Luigi V. Mancini, editors, *Proceedings of the Second Workshop on Cryptography for Ad-hoc Networks, WCAN@ICALP 2006, Venice, Italy, July 16, 2006*, volume 171 of *Electronic Notes in Theoretical Computer Science*, pages 107–120. Elsevier, 2006.
- [161] Arnaud Legout, Guillaume Urvoy-Keller, and Pietro Michiardi. Rarest first and choke algorithms are enough. *CoRR*, abs/cs/0609026, 2006.
- [162] Pietro Michiardi and Refik Molva. Analysis of coalition formation and cooperation strategies in mobile ad hoc networks. *Ad Hoc Networks*, 3(2):193–219, 2005.
- [163] Eitan Altman, Arzad Alam Kherani, Pietro Michiardi, and Refik Molva. Non-cooperative forwarding in ad-hoc networks. In Raouf Boutaba, Kevin C. Almeroth, Ramón Puigjaner, Sherman X. Shen, and James P. Black, editors, *NETWORKING 2005: Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communication Systems, 4th International IFIP-TC6 Networking Conference, Waterloo, Canada, May 2-6, 2005, Proceedings*, volume 3462 of *Lecture Notes in Computer Science*, pages 486–498. Springer, 2005.
- [164] Eitan Altman, Vivek S. Borkar, Arzad Alam Kherani, Pietro Michiardi, and Refik Molva. Some game-theoretic problems in wireless ad-hoc networks. In Gabriele Kotsis and Otto Spaniol, editors, *Wireless Systems and Mobility in Next Generation Internet, First International Workshop of the EURO-NGI Network of Excellence, Dagstuhl Castle, Germany, June 7-9, 2004, Revised Selected Papers*, volume 3427 of *Lecture Notes in Computer Science*, pages 82–104. Springer, 2004.
- [165] Refik Molva and Pietro Michiardi. Security in ad hoc networks. In Marco Conti, Silvia Giordano, Enrico Gregori, and Stephan Olariu, editors, *Personal Wireless Communications, IFIP-TC6 8th International Conference, PWC 2003, Venice, Italy, September 23-25, 2003, Proceedings*, volume 2775 of *Lecture Notes in Computer Science*, pages 756–775. Springer, 2003.

- [166] Pietro Michiardi and Refik Molva. Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In Borka Jerman-Blazic and Tomaz Klobucar, editors, *Advanced Communications and Multimedia Security, IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, September 26-27, 2002, Portoroz, Slovenia*, volume 228 of *IFIP Conference Proceedings*, pages 107–121. Kluwer, 2002.