Giuliano Casale

Imperial College London
Department of Computing
180 Queen's Gate, London SW7 2AZ, UK

g.casale@imperial.ac.uk
http://wp.doc.ic.ac.uk/gcasale/
https://orcid.org/0000-0003-4548-7951

Academic Employment

2020-	*Reader*. Imperial College London, Department of Computing, UK.
2015-20	*Senior Lecturer*. Imperial College London, Department of Computing, UK.
2013-15	*Lecturer*. Imperial College London, Department of Computing, UK.
2010-12	*Imperial College Research Fellow*. Imperial College London, Department of Computing, UK.
2007-08	*Postdoctoral Research Associate*. College of William & Mary, VA, US. Group: Prof. E. Smirni.

Other employment

2009	*Research Staff Member*. SAP Research, UK.
2006	*Research Scientist*. Neptuny, Milan, Italy. (Startup, now part of BMC Software and Moviri.)

Education

2003-06	*Ph.D., Computing*, Politecnico di Milano, Italy. Advisor: Prof. G. Serazzi.
1997-02	*MEng, Computing*, Politecnico di Milano, Italy. Grade: 100/100. 5-year curriculum.
2000-01	*Erasmus Exchange Student,* University of Manchester, Department of Computing, UK.

Awards

2022	Best Paper Award, ICSOC (with R. Wang, A. Filieri).
2022	Best Paper Award, IEEE/IFIP DSN (with S. Tuli, N. Jennings).
2022	Best Paper Award, IEEE INFOCOM (with S. Tuli, N. Jennings).
2017	Best Paper Award, ACM SIGMETRICS.
2016	Best Paper Award, ACM/SPEC ICPE (with W. Wang, A. Kattepur, M. Nambiar).
2015	Best Paper Award, IEEE ICCAC (with D. J. Dubois).
2015	Best Paper Award, IEEE CLOUD (with J. Wen, L. Lu, E. Smirni).
2012	Best Paper Award, ACM/SPEC ICPE (with P.G. Harrison).
2008	Best Paper Award, ACM/IFIP/USENIX MIDDLEWARE (with N. Mi, L. Cherkasova, E. Smirni).
2017	Best Student Paper Award, IFIP/IEEE IM (with K. Molka).
2008	Best Student Paper Award, QEST (with Z. Zhang, E. Smirni).
2004	Best Student Paper Award, IEEE MASCOTS (with G. Serazzi).
2019	Best Demo Award, ACM/SPEC ICPE.
2018	Best-in-Session Presentation Award, IEEE INFOCOM.
2003	IBM Best European Graduates Recognition.

Professional Service

Officer roles and steering committees

2019-23	SIG Chair: ACM SIGMETRICS.
2015-19	Board of Directors: ACM SIGMETRICS.

2021-24	Steering Committee: QEST conference series.
2023	Steering Committee: Reliability of Softw.-Driven Next-Gen. Networks workshop series.
2017-21	Steering Committee: QUDOS workshop series.
2014-16	Steering Committee: ICAC conference series.

CONFERENCE CHAIRING

2013	General Co-Chair: ACM/SPEC ICPE 2013.
2021	Program Co-Chair: IEEE/IFIP DSN 2021.
2017	Program Co-Chair: VALUETOOLS 2017.
2015	Program Co-Chair: IEEE MASCOTS 2015.
2014	Program Co-Chair: USENIX ICAC 2014.
2012	Program Co-Chair: QEST 2012.
2012	Program Co-Chair: ACM SIGMETRICS/Performance 2012.

EDITORIAL WORK

2022-	Editor-in-Chief, Elsevier Performance Evaluation (PEVA).
2011-15	Newsletter Editor, ACM SIGMETRICS Performance Evaluation Review (PER).
2024-	Associate Editor, ACM Trans. on Autonomous and Adaptive Systems (TAAS).
2018-	Associate Editor, ACM Trans. on Modeling and Perf. Eval. of Computing Systems (TOMPECS).
2017-23	Associate Editor, IEEE Trans. on Network and Service Management (TNSM).
2022	Associate Editor, Elsevier Performance Evaluation (PEVA).
2022	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2021	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2020	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2019	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2018	Special Issue Guest Editor, Elsevier Performance Evaluation.
2017	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2016	Special Issue Guest Editor, IEEE Trans. on Network and Service Management.
2012	Special Issue Guest Editor, Elsevier Performance Evaluation.
2009	Special Issue Guest Editor, ACM Performance Evaluation Review.

MEMBERSHIP, ROLES, AND COMMITTEE ACTIVITIES IN PROFESSIONAL SOCIETIES

2024	ACM SIGMETRICS 2024 Doctoral Dissertation Award Committee.
2024	Membership Committee, IFIP WG 7.3.
2023	ACM SIGMETRICS 2023 Doctoral Dissertation Award Committee.
2020	British Computer Society (BCS) Fellow.
2020	ACM Senior member.
2019	EiC Search committee, ACM Trans. on Modeling and Perf. Evaluation of Computing Systems.
2017	Reviewer: BCS Distinguished Dissertation award.
2017	Reviewer: ACM Future of Computing Academy programme.
2013-	IFIP Working Group 7.3 member, Computer System Modeling.
2012-	Imperial College representative, SPEC Research Group.
2011-12	Secretary/Treasurer, IEEE Special Technical Community on Sustainable Computing.

CONFERENCE ORGANIZATION COMMITTEES

2020	EU projects track chair: ESOCC.
2016	Organizer: 2nd QUDOS Workshop (at ISSTA 2016).
2015	Organizer: 1st QUDOS Workshop (at ESEC/FSE 2015).

2015	Replication Packages Co-Chair: ESEC/FSE.
2014	Publicity Chair: BDC.
2014	Publicity Chair: QEST.
2011	Organizer: Joint HPDC/ SIGMETRICS Student Poster Event.
2011	Organizer: SIGMETRICS Student Industry Workshop.
2011	Student Activities Chair: SIGMETRICS.
2011	Track Chair: EURO-PAR.
2010	Publicity Chair: QEST.
2007	Publicity Chair: MINENET.

OTHER SERVICE

2025	Book Reviewer: Elsevier.
2022	Session Chair: QEST.
2021	Session Chair: TOSME workshop.
2019	Workshop Host: MAMA.
2017	Session Chair: SIGMETRICS.
2016	Session Chair: ICPE.
2015	Session Chair: VALUETOOLS.
2013	Session Chair: MASCOTS.
2013	Book Reviewer: Morgan Kaufmann Publishers.
2011	Session Chair: SIGMETRICS.
2011	Track Chair: EURO-PAR.
2010	Session Chair: HOTMETRICS.
2009	Session Chair: QEST.
2009	Session Chair: HOTMETRICS.
2009	Session Chair: MASCOTS.
2006	Session Chair: SYS-ML.

RESEARCH VISITS

2017	University of Melbourne, Australia. Host: R. Buyya.
2016	Carleton University, Ottawa, Canada. Host: M. Woodside.
2015	Tata Consulting Services Research, Mumbai, India. Host: M. Nambiar.
2015	Gran Sasso Science Institute, Italy. Host: C. Trubiani.
2014	College of William & Mary, VA, US. Host: E. Smirni.
2013	College of William & Mary, VA, US. Host: E. Smirni.
2012	BCAM, Bilbao, Spain. Host: J. Anselmi.
2004	UCLA, Computer Science Dept., CA, US. Host: R.R. Muntz.

TECHNICAL PROGRAM COMMITTEES

I served 150+ times in TPCs (conferences and workshops):

ACSOS (1)	AI4AS (1)	AIMC (1)	ANNET (2)
ANSERVAPP (2)	APSYS (1)	ASMTA (4)	BDC (2)
BigData (2)	CAC (1)	CCGRID (3)	CLOUD (6)
CLOUDMDE (1)	DCPERF (3)	DSN (7)	DSSO (1)
ECMS (1)	EPEW (4)	ESOCC (2)	EURO-PAR (1)
GREENMETRICS (4)	HICS (1)	HIM (1)	HOTMETRICS (1)
HPDC (1)	IC2E (1)	ICAC (3)	ICCCN (3)
ICDCS (10)	ICPE (4)	INFQ (2)	IPDPS (1)
ISIICT (2)	IWQoS (4)	JCC (1)	MAMA (10)
MAM (2)	MASCOTS (5)	MDHPCL (1)	MICAS (3)
MOD (1)	NOMS (3)	PABS (1)	PERFORMANCE (7)
QEST (11)	ROSSA (1)	SC BoF (1)	SDSSC (1)
SeAC (1)	SIGMETRICS (10)	SOCC (1)	SRDS (1)

Stochastic Models (1)	TOSME (1)	UCC (1)	VALUETOOLS (9)
WAIN (1)	WOSP-C (1)		

JOURNAL REVIEWS

I served as journal reviewer 90+ times:

Access (2)	ASOC (1)	Bernoulli J. (1)
CAI (1)	COMCOM (10)	Computer J. (3)
DKE (1)	EJOR (7)	FGCS (2)
IEEE Computer (1)	IEEE Internet of Things J. (1)	INFORMS JoC (1)
J. of Grid Computing (1)	JSS (1)	NRL (1)
OPRE (1)	PER (3)	PEVA (18)
QUESTA (3)	SIMPAT (3)	SOSYM (1)
Stochastic Models (1)	TAAS (1)	TACON (1)
TC (2)	TCC (3)	TETC (1)
TII (1)	TOMPECS (3)	TON (2)
TOP (1)	TPDPS (3)	TSC (1)
TSE (4)	TSMC A (1)	TNSM (10)
TOSEM (2)	TP (3)	TRE (1)

RESEARCH AND GRANT ASSESSMENT

2024	Grant Reviewer, EPSRC Discipline hopping in ICT, UK.
2023	Grant Reviewer, EPSRC Responsive mode, UK.
2023	Grant Reviewer, Horizon Europe, FET-OPEN scheme, 2023 call.
2022	Grant Reviewer, EPSRC Open Fellowship scheme, UK.
2021	Grant Reviewer, Italian PRIN programme.
2021	Grant Reviewer, EPSRC New Investigator Award scheme, UK.
2020	Grant Reviewer, ESRC, UK.
2020	Grant Reviewer, Horizon 2020, FET-OPEN scheme, 2020 call.
2020	Grant Reviewer, Christian Doppler research association.
2020	Grant Reviewer, Kuwait Foundation for the Advancement of Sciences.
2019	Grant Reviewer, Horizon 2020, FET-OPEN scheme, 2019 national call.
2019	Grant Reviewer, Romanian National Council for Scientific Research, 2019 call.
2019	Grant Reviewer, British Council Newton Institutional Links grants, 2019 call.
2018-	EPSRC Full Peer Review College Member, UK.
2018	Grant Reviewer, British Council Newton Institutional Links grants, 2018 call.
2017	Reviewer: Mock Research Assessment Exercise, Dept. Information Eng., CUHK, Hong Kong.
2017-18	EPSRC Associate Peer Review College Member, UK.
2017	Grant Reviewer, EPSRC Responsive Mode, UK.
2017	Grant Reviewer, NSERC Council, Canada.
2016	Grant Reviewer, Romanian National Council for Scientific Research.
2016	Grant Reviewer, Flemish Research Foundation (FWO), Belgium.
2016	Review Panel Member for Italian PRIN projects.
2015	Grant Reviewer for Swiss National Science Foundation.
2014	Grant Reviewer for the Daphne Jackson Trust, UK.
2014	Grant Reviewer for NSERC Council, Canada.
2012	Reviewer for Italian Research Assessment Exercise 2004-10 (VQR).
2012	Review Panel Member for the Romanian National Council for Scientific Research.
2012	Review Panel Member for the Italian Future in Research 2012 programme.

OUTREACH ACTIVITIES

2021	Online webinar at OASIS TOSCA Implementation Stories series.
2021	Online webinar at H-CLOUD technical community event.

2020	Online webinar at Software Practice Advancement group, British Computer Society.
2020	Online webinar on serverless computing, Eficode webinar series.
2020	H-CLOUD project communications task force.
2019	Project presentation at SummerSOC school, Crete, Greece.
2017	Presentation at DevOpsDays event, Warsaw, Poland (500+ attendees).
2016	Presentation at OW2 event, CloudExpo, Excel London.
2015	Presentation at Computing Measurement Group, Mumbai, India (100+ attendees).
2015	Featured interview on Imperial College homepage about the DICE research project.
2015	Presentation at OW2 event, CloudExpo, Excel London.
2015	Presentation, RELATE ITN Doctoral School, Würzburg, Germany.
2014	Presentation, Future Internet Assembly, Athens.
2014	Presentation at OW2 event, CloudExpo, Excel London.
2013	Imperial Fringe Festival, Cloud Computing stand.
2013	Featured interview on Imperial College website to advertise JRF programme.
2013	MODAClouds presentation, ICT 2013, Vilnius, Lithuania.
2011-15	ACM SIGMETRICS LinkedIn group coordinator.

SOFTWARE

2024	Sourceforge's "Rising Star" recognition for the LINE scientific software (line-solver.sf.net).
2022	Sourceforge's "Community Leader" recognition for Java Modelling Tools (jmt.sf.net).

PATENTS

2012	*Simulation Techniques for Predicting In-Memory Database Systems Performance.* US Patent, US9111022 B2, assignee SAP, co-inventors: S. Kraft, A. Jula.
2011	*Characterizing Web Workloads For Quality of Service Prediction.* US Patent, US8560618 B2, assignee SAP, co-inventor: S. Pacheco-Sanchez.
2011	*Predicting performance of a consolidated virtualized computing environment.* US Patent, US 9164785 B2, assignee SAP, co-inventors: S. Kraft, D. Krishnamurthy.
2011	*Estimating service resource consumption based on response time.* European Patent, EP20100013643, assignee SAP, co-inventors: S. Kraft, S. Pacheco-Sanchez, S. Dawson.

INVITED TALKS AND SEMINARS

2024	Keynote, ACM/SPEC International Conference on Performance Engineering (ICPE).
2024	Keynote, International Conference on Big Data and Artificial Intelligence (BDA).
2023	Keynote, IEEE International Symposium on Software Reliability Engineering (ISSRE).
2023	Invited talk, TeaPACS 2023, Int'l Workshop on Teaching Performance Analysis of Comp. Sys.
2022	Keynote, AUSPDC 2022 symposium, CORE Australasian Computer Science Week.
2022	Tata Consulting Services, Mumbai, India.
2021	Invited talk, SODALITE Final workshop, Cloud Expo Europe Frankfurt.
2021	Invited talk, Huawei European Autonomous Network Forum, Huawei Ireland.
2020	Invited talk, Huawei Autonomous Network Conference, Huawei Ireland.
2019	Virginia Tech, US.
2019	University of Warwick, UK.
2017	University of Delft, The Netherlands.
2017	Umeå University, Sweden.
2017	Keynote, 1st Vienna Software Seminar on DevOps/Continuous Delivery, Vienna, Austria.
2016	Invited talk, 2nd Workshop on Performance Analysis of Big data Systems (PABS), Delft, NL.
2016	University of Pavia, Italy.
2016	University of Leeds, UK, School of Computing Weekly Colloquium.
2016	Budapest University of Technology and Economics, Hungary.
2015	University of Vienna, Austria.
2015	Gran Sasso Science Institute, Italy.

2015	University of York, UK.
2015	Tata Consulting Services, Mumbai, India.
2015	Keynote, TACTiCS Symp. on Performance Engineering, Mumbai, India.
2014	INRIA Grenoble, France.
2013	University of Edinburgh, UK.
2012	BCAM, Spain.
2011	Politecnico di Milano, Italy.
2011	Keynote, 5th Int. Workshop on Practical Applic. of Stoch. Modelling (PASM), Karlsruhe, DE.
2010	University of Florence, Italy.
2010	University Tor Vergata, Rome, Italy.
2010	BCAM - Basque Center for Applied Mathematics, Bilbao, Spain.
2010	IFIP Working Group 7.3 Workshop, Namur, Belgium.
2010	University of Munich (LMU), Germany.
2010	University of Venice (Cá Foscari), Italy.
2009	Northeastern University, Boston.
2009	INRIA-Grenoble Rhone Alpes, France.
2008	Mathworks, SIMULINK group, Boston, MA.
2008	SAP Research, Belfast, UK.
2007	Dagstuhl - Numerical Methods for Structured Markov Chains.
2007	College of William & Mary, Williamsburg, VA.
2006	IFIP Working Group 7.3 Workshop, St. Malo, France.

Full List of Publications

Refereed Journal Publications

J65.	TSC	X. Wu, H. Yu, G. Casale, G. Gao. Towards Cost-Optimal Policies for DAGs to Utilize IaaS Clouds with Online Learning, in *ACM Trans. on Services Computing*, full paper, to appear in 2025.
J64.	TOMPECS	Y. Gao, G. Casale, R. Singhal. Performance Modeling of Distributed Data Processing in Microservice Applications, in *ACM Trans. on Modeling and Performance Evaluation of Computer Systems*, 9(4):1–30, 2024.
J63.	JNSM	A. Gias, Y. Gao, M. Sheldon, J. A. Perusquía, O. O'Brien, G. Casale. SampleHST-X: A Point and Collective Anomaly-Aware Trace Sampling Pipeline with Approximate Half Space Trees, in *Journal of Network and Systems Management*, 33(3):44, 2024.
J62.	TSE	Z. Niu, G. Casale. Neural Density Estimation of Response Times in Layered Software Systems, *IEEE Trans. Softw. Eng.*, 50(3):636-650, Mar 2024.
J61.	TOMPECS	D. Olliaro, G. Casale, A. Marin, S. Rossi. A product-form network for systems with job stealing policies, *ACM Trans. on Modeling and Performance Evaluation of Computer Systems*, 9(2): 6:1-6:26, 2024.
J60.	TOMACS	G. Casale, Y. Gao, Z. Niu, L. Zhu. LN: A Flexible Algorithmic Framework for Layered Queueing Network Analysis, *ACM Trans. on Modeling and Comp. Simulation*, 34(3):1–26, Jul 2024.
J59.	TMC	S. Tuli, G. Casale, N. Jennings. PreGAN+: Semi-Supervised Fault Prediction and Preemptive Migration in Dynamic Mobile Edge Environments, *IEEE Trans. on Mobile Computing*, 23:6881-6895, Jun 2024.
J58.	SPE	L. Zhu, D. A. Tamburri, G. Casale. RADF: Architecture Decomposition for Function as a Service, *Software: Practice and Experience*, 54(4):566-594, Apr 2024.
J57.	TC	G. Casale, M. Roveri. Scheduling Inputs in Early Exit Neural Networks, *IEEE Trans. on Computers*, 73(2):451-465, Feb 2024.
J56.	TC	S. Tuli, G. Casale, N. Jennings. SciNet: Co-Design of Resource Management in Cloud Computing Environments, *IEEE Trans. on Computers*, 72:3590-3602, Dec 2023.
J55.	TNSM	S. Tuli, G. Casale, N. Jennings. CILP: Co-simulation based Imitation Learner for Dynamic Resource Provisioning in Cloud Computing Environments, *IEEE Trans. on Network and Service Management*, 20(4):4448-4460, Dec 2023.

J54. JNCA S. Tuli, F. Mirhakimi, S. Pallewatta, S. Zawad, G. Casale, B. Javadi, F. Yan, R. Buyya, N. R. Jennings. AI Augmented Edge and Fog Computing: Trends and Challenges, *Elsevier J. of Network and Computer Applications*, 216, article 103648, Jul 2023.

J53. TOMPECS X. Wu, F. De Pellegrini, G. Casale. Delay and Price Differentiation in Cloud Computing: A Service Model, Supporting Architectures, and Performance, *ACM Trans. on Modeling and Performance Evaluation of Computer Systems*, 8(3), Article 6, pp. 1-40, 2023.

J52. TOMACS R. Wang, G. Casale, A. Filieri. Estimating Multiclass Service Demand Distributions Using Markovian Arrival Processes, *ACM Trans. on Modeling and Comp. Simulation*, 33(1-2):11-26, Feb 2023.

J51. TNSM S. Tuli, G. Casale, N. Jennings. DRAGON: Decentralized Fault Tolerance in Edge Federations, *IEEE Trans. on Network and Service Management*, 20(1):276-291, Mar 2023.

J50. SR S. Tuli, G. Casale, N. Jennings. SimTune: Bridging the Simulator Reality Gap for Resource Management in Edge-Cloud Computing, *Nature Scientific Reports*, 12, article 19158, Nov 2022.

J49. TMC S. Tuli, G. Casale, N. Jennings. SplitPlace: AI Augmented Splitting and Placement of Large-Scale Neural Networks in Mobile Edge Environments, *IEEE Trans. on Mobile Computing*, 22:5539-5554, Sep 2023.

J48. TPDS J. Soikkeli, G. Casale, L. Munoz Gonzalez, E. Lupu. Redundancy Planning for Cost Efficient Resilience to Cyber Attacks, *IEEE Trans. on Dependable and Secure Computing*, 20(2):1154–1168, Mar-Apr 2023.

J47. PVLDB S. Tuli, G. Casale, N. Jennings. TranAD: Deep Transformer Networks for Anomaly Detection in Multivariate Time Series Data, *PVLDB*, 15(6):1201-1214.

J46. TPDS S. Tuli, G. Casale, N. Jennings. GOSH: Task Scheduling using Deep Surrogate Models in Fog Computing Environments, *IEEE Trans. on Parallel and Distributed Systems*, 33:2821-2833, Nov 2022.

J45. TSC S. Tuli, S.S. Gill, P. Garraghan, R. Buyya, G. Casale, N. Jennings. START: Straggler Prediction and Mitigation for Cloud Computing Environments using Encoder LSTM Networks, *IEEE Trans. on Services Computing*, 16:615-627, Jan-Feb 2023.

J44. TPDS S. Tuli, G. Casale, N. Jennings. MCDS: AI Augmented Workflow Scheduling in Mobile Edge Cloud Computing Systems, *IEEE Trans. on Parallel and Distributed Systems*, 33(11):2794–2807, Nov 2022.

J43. TPDS S. Tuli, S. Poojara, S. Srirama, G. Casale, N. Jennings. COSCO: Container Orchestration using Co-Simulation and Gradient Based Optimization for Fog Computing Environments, *IEEE Trans. on Parallel and Distributed Systems*, 33(1):101–116, Jan 2022.

J42. PEVA G. Casale, P.G. Harrison, W.H. Ong. Facilitating Load-Dependent Queueing Analysis Through Factorization, in *Perform. Eval.*, Elsevier, vol. 152, Nov 2021.

J41. JSS S. Tuli, S.S. Gill, M. Xu, P. Garraghan, R. Bahsoon, S. Dustdar, R. Sakellariou, O. Rana, R. Buyya, G. Casale, N. Jennings. HUNTER: AI based Holistic Resource Management for Sustainable Cloud Computing, *J. of Systems and Software*, Elsevier, article 111124, Oct 2021.

J40. AAP I. Perez, G. Casale. Variational inequalities and mean-field approximations for partially observed systems of queueing networks, *Advances in Applied Prob.*, 53(3), 687-715, Sep 2021.

J39. TON G. Casale, N. Gast. Performance analysis of list-based caches with non-uniform access, *IEEE/ACM Trans. on Networking*, 29(2):651–664, Apr 2021.

J38. ACCESS A. Alnafessah, A. Ul Gias, R. Wang, L. Zhu, G. Casale, A. Filieri. Quality-Aware DevOps research: where do we stand?, *IEEE Access*, Mar 2021.

J37. ACCESS A. Alnafessah, G. Casale. TRACK-Plus: Optimizing Artificial Neural Networks for Hybrid Anomaly Detection in Data Streaming Systems, *IEEE Access*, Jul 2020.

J36. PEVA L. Zhu, G. Casale, I. Perez. Fluid Approximation of Closed Queueing Networks with Discriminatory Processor Sharing, *Perform. Eval.*, Elsevier, Vol. 139, Jun 2020.

J35. TOMPECS X. Wu, F. De Pellegrini, G. Gao, G. Casale. A Framework for Allocating Server Time to Spot and On-demand Services in Cloud Computing, *ACM Trans. on Modeling and Perform. Eval. Comp. Sys.*, 4(4), 1–31, Dec 2019.

J34. CLUS A. Alnafessah, G. Casale. Artificial Neural Networks Based Techniques for Anomaly Detection in Apache Spark, *Cluster Computing*, Springer, Oct 2019.

J33. SICS G. Casale, M. Artač, W.-J. van den Heuvel, *et al*. RADON: Rational Decomposition and Orchestration for Serverless Computing, *Software-Intensive Cyber-Physical Systems*, Springer, Aug 2019.

J32. JSS S. S. Gill, P. Garraghan, V. Stankovski, G. Casale, *et al*. Holistic Resource Management for Sustainable and Reliable Cloud Computing: An Innovative Solution to Global Challenge, *Journal of Systems and Software*, Elsevier, May 2019.

J31. CSUR R. Buyya, S. N. Srirama, G. Casale, R. Calheiros, Y. Simmhan, B. Varghese, *et al*. A Manifesto for Future Generation Cloud Computing: Research Directions for the Next Decade, *ACM Computing Surveys*, 51(5):105, 2019.

J30. TOE D. Tamburri, G. Casale. Cognitive Distance and Research Output in Computing Education: A Case-Study, *IEEE Trans. on Education*, 1–9, Oct 2018.

J29. TOMPECS W. Wang, G. Casale, A. Kattepur, and M. Nambiar. QMLE: a Methodology for Statistical Inference of Service Demands from Queueing Data, *ACM Trans. on Modeling and Perform. Eval. Comp. Sys.*, 17:1-28, 2018.

J28. POMACS G. Casale. Accelerating Performance Inference over Closed Systems by Asymptotic Methods, *PACM on Meas. Anal. Comput. Syst.* journal, 1(1):1–23, Jun 2017.

J27. TREL J.F. Peréz, G. Casale. LINE : A Scalable Tool for Evaluating Software Applications in Unreliable Environments, *IEEE Trans. on Reliability*, 2017.

J26. TOMPECS K. Molka, G. Casale. Contention-Aware Workload Placement for In-Memory Databases in Cloud Environments, *ACM Trans. on Modeling and Perform. Eval. Comp. Sys.*, 2(1):1–29, Oct 2016.

J25. EJOR G. Casale, A. Sansottera, P. Cremonesi. Compact Markov-Modulated Models for Multiclass Trace Fitting, *European J. of Oper. Research*, 255(3):822–833, Nov 2016.

J24. TOMACS W. Wang, G. Casale, C. Sutton. A Bayesian Approach to Parameter Inference in Queueing Networks, *ACM Trans. on Modeling and Comp. Simulation*, 27(1):1–26, Nov 2016.

J23. CLUS D. J. Dubois, G. Casale. OptiSpot: Minimizing Application Deployment Cost using Spot Cloud Resources, *Cluster Computing*, 19(2):893–909, Springer, Mar 2016.

J22. TOMACS G. Casale, V. de Nitto-Personé, E. Smirni. QRF: An Optimization-Based Framework for Evaluating Complex Stochastic Networks, *ACM Trans. on Modeling and Computer Simulation*, 26(3):15, Jan 2016.

J21. PEVA S. Spinner, G. Casale, F. Brosig, S. Kounev. Evaluating Approaches to Resource Demand Estimation, *Perform. Eval.*, Elsevier 92:51-71, Oct 2015.

J20. PEVA G. Casale, J. F. Pérez, W. Wang. QD-AMVA: Evaluating Systems with Queue-Dependent Service Requirements, *Perform. Eval.*, Elsevier, 91:80-98, Sep 2015.

J19. TSE J.F. Pérez, G. Casale, S. Pacheco-Sanchez. Estimating Computational Requirements in Multi-Threaded Applications, *IEEE Trans. on Softw. Eng.*, 41(3):264–278, Mar 2015.

J18. PEVA G. Casale, M. Tribastone, P.G. Harrison. Blending Randomness in Closed Queueing Network Models, *Perform. Eval.*, Elsevier, 82:15-38, Dec 2014.

J17. JISA D. Ardagna, G. Casale, M. Ciavotta, J. F. Pérez, W. Wang. Quality-of-Service in Cloud Computing: Modelling Techniques and Their Applications, *J. of Internet Services and Applications*, 5(1), Springer, Sep 2014.

J16. PEVA J. Anselmi, G. Casale. Heavy-Traffic Revenue Maximization in Parallel Multiclass Queues, *Perform. Eval.*, 70(1):806–821, Elsevier, Oct 2013.

J15. SOSYM S. Kraft, G. Casale, D. Krishnamurthy, D. Greer, P. Kilpatrick. Performance Models of Storage Contention in Cloud Environments, *J. of Software and Systems Modeling*, Springer, Oct 2013.

J14. TSE G. Casale, N. Mi, L. Cherkasova, E. Smirni. Dealing With Burstiness in Multi-Tier Applications: Models and Their Parameterization, *IEEE Trans. on Softw. Eng.*, 38(5): 1040-1053, Sep/Oct 2012.

J13. TSE G. Casale, A. Kalbasi, D. Krishnamurthy, J. Rolia. BURN: Enabling Workload Burstiness in Customized Service Benchmarks, *IEEE Trans. on Softw. Eng.*, 38(4):778–793, Jul/Aug 2012.

J12. TNSM N. Mi, G. Casale, E. Smirni. ASIdE: Using Autocorrelation-Based Size Estimation for Scheduling Bursty Workloads, *IEEE Trans. on Network and Service Management*, 9(2):198–212, Jun 2012.

J11. PEVA G. Casale. Exact Analysis of Performance Models by the Methods of Moments, *Perform. Eval.*, Elsevier, 68(6):873-896, Jun 2011.

J10. PEVA G. Casale. A Generalized Method of Moments for Closed Queueing Networks, *Perform. Eval.*, Elsevier, 68(2):180-0, Feb 2011.

J9. PEVA G. Casale. Approximating Passage Time Distributions in Queueing Models by Bayesian Expansion, *Perform. Eval.*, Elsevier, 67(11):1076–1091, Nov 2010.

J8. TC G. Casale, N. Mi, E. Smirni. Model-Driven System Capacity Planning Under Workload Burstiness, *IEEE Trans. on Computers*, 59(1):66-80, Jan 2010.

J7. PEVA G. Casale, E. Z. Zhang, E. Smirni. KPC-Toolbox: Best Recipes for Automatic Trace Fitting Using Markovian Arrival Processes, *Perform. Eval.*, Elsevier, 67(9):873–896, Sep 2010.

J6. JISA N. Mi, G. Casale, L. Cherkasova, E. Smirni. Sizing Multi-Tier Systems with Temporal Dependence: Benchmarks and Analytic Models, *Springer J. of Internet Services and Applications*, 1(2):117–134, Nov 2010.

J5. PEVA G. Casale, E. Z. Zhang, E. Smirni. Trace Data Characterization and Fitting for Markov Modeling, *Perform. Eval.*, Elsevier, 67(2):61-79, Feb 2010.

J4. TSE G. Casale. CoMoM: Efficient Class-Oriented Evaluation of Multiclass Performance Models, *IEEE Trans. on Softw. Eng.*, 35(2):162-177, Mar/Apr 2009.

J3. TC G. Casale, R. R. Muntz, G. Serazzi. Geometric Bounds: a Non-Iterative Analysis Technique for Closed Queueing Networks, *IEEE Trans. on Computers*, 57(6):780–794, Jun 2008.

J2. QUESTA G. Casale. A Note on Stable Flow-Equivalent Aggregation in Closed Networks, *Springer Queueing Systems*, 60(3):193–2, Dec 2008.

J1. COR E. Amaldi, M. Bruglieri, G. Casale. A Two-Phase Relaxation-Based Heuristic for the Maximum Feasible Subsystem Problem, *Computers & Operations Research*, Elsevier, 35(5):1465–1482, May 2008.

REFEREED CONFERENCE PUBLICATIONS

Conference papers appearing in journal proceedings are marked with an asterisk.

C72. IPDPS B. Sun, R. Pinciroli, G. Casale, E. Smirni. DeepBAT: Performance and Cost Optimization of Serverless Inference Using Transformers, in *Proc. of IEEE IPDPS*, June 2025, full paper.

C71. INFOCOM Y. Chen, Z. Niu, M. Roveri, G. Casale. CEED: Collaborative Early Exit Neural Network Inference at the Edge, in *Proc. of INFOCOM*, May 2025, full paper. [acc. rate: 18.7%].

C70*. ICPE Journal First track paper at *ICPE*, published in journal special issue [J64], May 2025.

C69. ESOCC Y. Gao, R. Sala, D. Ardagna, G. Casale. Deep Surrogate Models of Serverless Batch Processing Services, in *Proc. of ESOCC*, Feb 2025, full paper, 15 pages.

C68. MASCOTS Y. Zhou, M. Sheldon, G. Casale. Approximating Closed Queueing Networks in Semi-Markov Random Environments, in *Proc. of MASCOTS*, Oct 2024, full paper, 8 pages.

C67. WSC W. Plumb, A. Bottle, G. Casale. Bayesian Optimization for Clinical Pathway Decomposition from Aggregate Data, in *Proc. of INFORMS Winter Simulation Conf.*, Dec 2024, 12 pages.

C66. DSN Z. Niu, M. Roveri, G. Casale. ChainNet: A Customized Graph Neural Network Model for Loss-aware Edge AI Service Deployment, in *Proc. of IEEE/IFIP DSN*, Jun 2024 [acc rate: 20.6%].

C65. WSC W. Plumb, A. Bottle, G. Casale, A. Liddle. Clinical Pathway Clustering Using Surrogate Likelihoods and Reliability Validation, in *Proc. of INFORMS Winter Simulation Conf.*, 12 pages, Dec 2023.

C64. CNSM Y. Chen, M. Roveri, S. Tuli, G. Casale. Coupling QoS Co-Simulation with Online Adaptive Arrival Forecasting, in *Proc. of IFIP/IEEE CNSM*, Nov 2023.

C63. NOMS A. Gias, Y. Gao, M. Sheldon, J. A. Perusquía, O. O'Brien, G. Casale. SampleHST: Efficient On-the-Fly Selection of Distributed Traces, in *Proc. of IEEE/IFIP NOMS*, 10 pages, May 2023.

C62. INFOCOM S. Tuli, G. Casale, L. Cherkasova, N. Jennings. DeepFT: Fault-Tolerant Edge Computing using a Self-Supervised Deep Surrogate Model, in *Proc. of IEEE INFOCOM*, 10 pages, May 2023 [acc. rate: 19.2%].

C61. ICSOC R. Wang, G. Casale, A. Filieri. Enhancing Performance Modeling of Serverless Functions via Static Analysis, in *Proc. of ICSOC*, 15 pages, Dec 2022. **Best Paper Award.**

C60. QEST G. Casale, Y. Gao, Z. Niu, L. Zhu. LN: a Meta-Solver for Layered Queueing Network Analysis, in *Proc. of QEST*, 22 pages, Sep 2022.

C59. IWQOS Y. Gao, G. Casale. JCSP: Joint Caching and Service Placement for Edge Computing Systems, in *Proc. of IEEE/ACM IWQoS*, 10 pages, Jun 2022. [acc. rate 24.7%]

C58. DSN S. Tuli, G. Casale, N. Jennings. CAROL: Confidence-Aware Resilience Model for Edge Federations, in *Proc. of IFIP/IEEE DSN*, 13 pages, Jun 2022. [acc. rate 18.7%]. **Best paper award**.

C57*. VLDB Regular paper at *VLDB 2022*, proceedings in PVLDB journal [J47], Sep 2022.

C56. INFOCOM S. Tuli, G. Casale, N. Jennings. PreGAN: Preemptive Migration Prediction Network for Proactive Fault-Tolerant Edge Computing, in *Proc. of IEEE INFOCOM*, 10 pages, May 2022 [acc. rate: 19.9%], **Best paper award**.

C55. MASCOTS Z. Niu, G. Casale. A Mixture Density Network Approach to Predicting Response Times in Layered Systems, in *Proc. of IEEE MASCOTS*, 8 pages, Nov 2021.

C54. MASCOTS Y. Chen, G. Casale. Deep Learning Models for Automated Identification of Scheduling Policies, in *Proc. of IEEE MASCOTS*, 8 pages, Nov 2021.

C53*. PERFORMANCE Regular paper at *IFIP PERFORMANCE 2021*, proceedings in PEVA [J41], Oct 2021.

C52. QEST R. Wang, G. Casale, A. Filieri. Service Demand Distribution Estimation for Microservices Using Markovian Arrival Processes, in *Proc. of QEST*, 18 pages, Aug 2021.

C51. CCGRID G. Russo Russo, V. Cardellini, G. Casale, F. Lo Presti. MEAD: Model-Based Vertical Auto-Scaling for Data Stream Processing, in *Proc. of IEEE/ACM CCGRID*, 10 pages, May 2021.

C50. MASCOTS A. Gias, G. Casale. COCOA: Cold Start Aware Capacity Planning for Function-as-a-Service Platforms, in *Proc. of IEEE MASCOTS*, 8 pages, Dec 2020.

C49. WSC G. Casale. Integrated performance evaluation of extended queueing network models with LINE, in *Proc. of Winter Simulation Conference*, 12 pages, Dec 2020.

C48. ICDCS A. Gias, G. Casale, M. Woodside. ATOM: Model-Driven Autoscaling for Microservices, in *Proc. of IEEE ICDCS*, 11 pages, Jul 2019.

C47. INFOCOM G. Casale. Analyzing replacement policies in list-based caches with non-uniform access costs, in *Proc. of IEEE INFOCOM*, 432–440, Apr 2018. Best-in-Session Presentation Award. [acc. rate: 19.2%].

C46. NOMS S. Dipietro, R. Buyya, G. Casale. PAX: Partition-Aware Autoscaling for the Cassandra NoSQL Database, in *Proc. of IEEE/IFIP NOMS*, Apr 2018.

C45*. SIGMETRICS Regular paper at *ACM SIGMETRICS 2017*, proceedings published in PACM journal [J28]. **Best Paper Award**. [acc. rate: 13%].

C44. CLOUD J. Wen, L. Ren, F. Yan, D. Dubois, G. Casale, E. Smirni. How to Supercharge the Amazon T2: Observations and Suggestions, in *Proc. of IEEE CLOUD*, Jun 2017.

C43. IM K. Molka, G. Casale. Energy-Efficient Resource Allocation and Provisioning for In-Memory Database Clusters, in *Proc. of IFIP/IEEE IM*, 2017. **Best Student Paper Award**.

C42. VALUETOOLS S. Dipietro, G. Casale, G. Serazzi. A Queueing Network Model for Performance Prediction of Apache Cassandra, in *Proc. of VALUETOOLS*, Nov 2016.

C41. MASCOTS P. Jamshidi, G. Casale. An Uncertainty-Aware Approach to Optimal Configuration of Stream Processing Systems, in *Proc. of IEEE MASCOTS*, Sep 2016.

C40. QRS R. Osman, J. F. Peréz, G. Casale. Quantifying the Impact of Replication on the Quality-of-Service in Cloud Databases, in *Proc. of IEEE QRS*, Aug 2016.

C39. CLOUD D. J. Dubois, C. Trubiani, G. Casale. Model-driven Application Refactoring to Minimize Deployment Costs in Preemptible Cloud Resources, in *Proc. of IEEE CLOUD*, Jun 2016. [acc. rate: 15%].

C38. ICPE W. Wang, G. Casale, A. Kattepur, M. K. Nambiar. Maximum Likelihood Estimation of Closed Queueing Network Demands from Queue Length Data, in *Proc. of ACM/SPEC ICPE*, Mar 2016. **Best Paper Award**.

C37. CNSM K. Molka, G. Casale. Experiments or Simulation? A Characterization of Evaluation Methods for In-Memory Databases, in *Proc. of IFIP/IEEE CNSM*, Nov 2015.

C36*. PERFORMANCE Regular paper at *IFIP PERFORMANCE 2015*, published in journal special issue [J20], Oct 2015.

C35. ICCAC D. J. Dubois, G. Casale. Autonomic Provisioning and Application Mapping on Spot Cloud Resources, in *Proc. of IEEE ICCAC*, Sep 2015. **Best Paper Award**.

C34. CLOUD J. Wen, L. Lu, G. Casale, E. Smirni. Less can be More: micro-Managing VMs in Amazon EC2, in *Proc. of IEEE CLOUD*, Jun 2015. **Best Paper Award**. [acc. rate: 17%].

C33. NOMS K. Molka, G. Casale, T. Molka, L. Moore. Memory-Aware Sizing for In-Memory Databases, in *Proc. of IEEE/IFIP NOMS*, May 2014. [acc. rate: 19%].

C32*. PERFORMANCE Regular paper at *IFIP PERFORMANCE 2013*, proceedings in PEVA [J16].

C31. MASCOTS W. Wang, G. Casale. Bayesian Service Demand Estimation with Gibbs Sampling, in *Proc. of IEEE MASCOTS*, Aug 2013.

C30. MASCOTS J. F. Peréz, S. Pacheco-Sanchez, G. Casale. An offline demand estimation method for multi-threaded applications, in *Proc. of IEEE MASCOTS*, Aug 2013. **Finalist for Best Paper Award**.

C29. DSN A. Sansottera, G. Casale, P. Cremonesi. Analysis of Second-Order Marked Markovian Arrival Processes, in *Proc. of IEEE/IFIP DSN*, Jun 2013 [acc. rate: 22%].

C28. IM S. Musabbi, D. Krishnamurthy, G. Casale. RPO: Runtime Web Server Optimization Under Simultaneous Multithreading, in *Proc. of IFIP/IEEE IM*, May 2013.

C27. CLOUD S. Kraft, G. Casale, A. Jula, P. Kilpatrick, D. Greer. WIQ: Work-Intensive Query Scheduling for In-Memory Database Systems, in *Proc. of IEEE CLOUD*, Jun 2012 [acc. rate:17%].

C26. ICPE G. Casale, P. Harrison. A Class of Tractable Models for Run-Time Performance Evaluation, in *Proc. of ACM/SPEC ICPE*, 63-74, Apr 2012. **Best Paper Award**.

C25. QEST G. Casale, M. Tribastone. Fluid Analysis of Queueing in Two-Stage Random Environments, in *Proc. of QEST*, Aachen, Germany, 21-30, Sep 2011.

C24. CLOUD S. Pacheco-Sanchez, G. Casale, B. Scotney, S. McClean, G. Parr, S. Dawson. Markovian Workload Characterization for QoS Prediction in the Cloud, in *Proc. of IEEE CLOUD*, 147-154, Washington D.C., Jun 2011 [acc. rate: 18%].

C23. DSN V. De Nitto, G. Casale, E. Smirni. Approximate Analysis of Blocking Queueing Networks with Temporal Dependence, in *Proc. of IEEE/IFIP DSN*, 574 - 585, Hong Kong, China, IEEE Press, in Jun 2011.

C22. VALUETOOLS M.Makaronidis, G. Casale. Efficient Parallelization of the Method of Moments for Queueing Networks Using Multi-Modular Algebra , in *Proc. of VALUETOOLS*, ACM, May 2011.

C21. ICPE S.Kraft, G. Casale, D.Krishnamurthy, D.Greer, P.Kilpatrick. I/O Performance Prediction in Consolidated Virtualized Environments, in *Proc. of ACM/SPEC ICPE*, 295-306, ACM, Mar 2011.

C20*. PERFORMANCE Regular paper at *IFIP PERFORMANCE 2010*, proceedings in PEVA [J9].

C19. SIGMETRICS G. Casale, N. Mi, E. Smirni. CWS: A Model-Driven Scheduling Policy for Correlated Workloads, in *Proc. of ACM SIGMETRICS*, New York, NY, ACM, 251-262, Jun 2010 [acc. rate: 16%].

C18. ICPE H.Li, G. Casale, T.Ellahi., SLA Driven Planning and Optimization of Enterprise Applications, in *Proc. of ACM/SPEC ICPE*, San Jose, CA, 117–128, Jan 2010.

C17. MIDDLEWARE G. Casale, A.Kalbasi, D.Khrishnamurthy, J. Rolia. Automatic Stress Testing of Multi-Tier Systems by Dynamic Bottleneck Switch Generation, in *Proc. of ACM/IFIP/USENIX MIDDLEWARE*, Urbana-Champaign, Illinois, Springer LNCS 5896, 393–413, Dec 2009 [acc. rate:19%].

C16. ICAC N.Mi G. Casale, L.Cherkasova, E. Smirni. Injecting Realistic Burstiness Into a Traditional Client-Server Benchmark, in *Proc. of IEEE ICAC*, Barcelona, Spain, 149–158, IEEE Press, Jun 2009, [acc. rate:16%].

C15. DSN G. Casale, E. Smirni. MAP-AMVA: Approximate Mean Value Analysis of Bursty Systems, in *Proc. of IEEE/IFIP DSN*, Estoril, Portugal, 409–418, IEEE Press, Jun 2009.

C14. VALUETOOLS S.Kraft, S.Pacheco-Sanchez, G. Casale, S.Dawson. Estimating Service Resource Consumption From Response Time Measurements, in *Proc. of VALUETOOLS*, Pisa, Italy, ACM, Oct 2009.

C13. MASCOTS N. Mi, G. Casale, Q. Zhang, A. Riska, E. Smirni. Autocorrelation-Driven Load Control in Distributed Systems, in *Proc. of IEEE MASCOTS*, London, UK, 269–278, IEEE Press, Sep 2009 [acc. rate: 20%].

C12. QEST G. Casale. The Multi-branched Method of Moments for Queueing Networks, in *Proc. of QEST*, Budapest, Hungary, 227–236, IEEE Press, Sep 2009.

C11. SIGMETRICS G. Casale, N. Mi, E. Smirni. Bound Analysis of Closed Queueing Networks with Workload Burstiness, in *Proc. of ACM SIGMETRICS*, 13-24, Annapolis, MD, ACM, Jun 2008 [acc. rate: 17%].

C10. DSN — N. Mi, G. Casale, E. Smirni. Scheduling for Performance and Availability in Systems with Temporal Dependent Workloads, *in Proc. of IEEE/IFIP DSN*, 336–345, Anchorage, AK, IEEE Press, Jun 2008.

C9. MIDDLEWARE — N. Mi, G. Casale, L. Cherkasova, E. Smirni. Burstiness in Multi-Tier Applications: Symptoms, Causes, and New Models, in *Proc. of ACM/IFIP/USENIX MIDDLEWARE*, 265–286, Leuven, Belgium, Springer LNCS 5346, Dec 2008. **Best Paper Award**, [acc. rate: 18%].

C8. QEST — G. Casale, E. Z. Zhang, E. Smirni. KPC-Toolbox: Simple Yet Effective Trace Fitting Using Markovian Arrival Processes, in *Proc. of QEST*, St.Malo, France, 83–92, IEEE Press, Sep 2008. **Best Student Paper Award**.

C7. ANSS — M. Bertoli, G. Casale, G. Serazzi. The JMT Simulator for Performance Evaluation of Non Product-Form Queueing Networks, in *Proc. of the 38th Annual Simulation Symposium*, 3 – 10, IEEE Press, 2007.

C6. MASCOTS — J. Anselmi, G. Casale, P. Cremonesi. Approximate Solution of Multiclass Queueing Networks with Region Constraints, in *Proc. of IEEE MASCOTS*, 1-5, Istanbul, Turkey, IEEE Press, 2007.

C5. SIGMETRICS — G. Casale. An Efficient Algorithm for the Exact Analysis of Multiclass Queueing Networks with Large Population Sizes, in *Proc. of joint ACM SIGMETRICS/IFIP PERFORMANCE 2006*, St. Malo, France, 169–180, 2006, ACM. [acc. rate: 14%].

C4. MASCOTS — G. Casale, R.R. Muntz, G. Serazzi. A New Class of Non-Iterative Bounds for Closed Queueing Networks, in *Proc. IEEE MASCOTS*, Monterey, US, 69 – 76, Sep 2006, IEEE Press.

C3. QEST — M. Bertoli, G. Casale, G. Serazzi. Java Modelling Tools: an Open Source Suite for Queueing Network Modelling and Workload Analysis, in *Proc. of QEST*, Riverside, US, Sep 2006, 119-120, IEEE Press, 2006.

C2. QEST — G. Casale. On Single Class Load-Dependent Normalizing Constant Equations, in *Proc. of QEST*, Riverside, US, Sep 2006, 333 – 342, IEEE Press, 2006.

C1. MASCOTS — G. Casale, G. Serazzi. Bottlenecks Identification in Multiclass Queueing Networks using Convex Polytopes, in *Proc. of MASCOTS*, 223-230, 2004, IEEE Press. **Best Student Paper Award**.

BOOK CHAPTERS

B5. IEEE/WILEY — A. Alnafessah, G. Russo Russo, V. Cardellini, G. Casale, F. Lo Presti. AI-driven performance management in data-intensive applications, in *Communication Networks and Service Management in the Era of Artificial Intelligence and Machine Learning*, IEEE/Wiley, 2021.

B4. SPRINGER — G. Casale, P. G. Harrison. AutoCAT: Automated Product-Form Solution of Stochastic Models, in *Matrix-Analytic Methods in Stochastic Models*, Springer, 27:57–85, 2013.

B3. LNCS — G. Casale, M. Gribaudo, G. Serazzi. *Tools for Performance Evaluation of Computer Systems: Historical Evolution and Perspectives*, in *Proc. of IFIP PERFORM*, Springer LNCS 6821, 24–37, Vienna, Austria, Oct 2010.

B2. ICP — G. Casale, G. Serazzi. Stabilisation Techniques for Load-Dependent Algorithms. Book chapter in *J.A.Barria Ed., Communication Networks and Computer Systems, Imperial College Press*, 2006.

B1. OCG — G. Casale, G. Serazzi. Estimating Bottlenecks of Very Large Models. Book chapter in *G.Kotsis Ed., Performance Evaluation: Stories and Perspective, Vienna, Austria*, Dec 2003, Austrian Computing Group (OCG) Press.

OTHER PUBLICATIONS

The following list includes letters, magazines, position papers, and invited papers.

O12. TEAPACS — G. Casale. Performance evaluation teaching in the age of cloud computing, in ACM Performance Evaluation Review, *Proc. of the 2nd International Workshop on Teaching Performance Analysis of Computer Systems*, 2023. (Invited).

O11. CLOUD S. Tuli, G. Casale, N. Jennings. MetaNet: Automated Dynamic Selection of Scheduling Policies in Cloud Environments, in *Proc. of IEEE CLOUD*, 11 pages, Jul 2022. (Invited).

O10. ITL S. Tuli, S. Gill, G. Casale, N. Jennings. iThermoFog: IoT-Fog based Automatic Thermal Profile Creation for Cloud Data Centers using Artificial Intelligence Techniques, *Internet Technology Letters*, Wiley, 2020.

O9. CF S. Di Pietro, G. Casale. SD: A Divergence-Based Estimation Method for Service Demands in Cloud Systems, in *Proc. of FiCloud*, 197-204, 2019. (Invited)

O8. IEEECC G. Kecskemeti, G. Casale, D. N. Jha, J. Lyon, R. Ranjan. Modelling and Simulation Challenges in Internet of Things. *IEEE Cloud Computing*, 4(1):62-69, 2017.

O7. CF G. Casale, C. Chesta, P. Deussen, E. Di Nitto, P. Gouvas, S. Koussouris, V. Stankovski, A. Symeonidis, V. Vlassiou, A. Zafeiropoulos and Z. Zhao. Current and Future Challenges of Software Engineering for Services and Applications, in *Proc. of CloudForward*, Madrid, 2016.

O6. CROSSCLOUD D. Petcu, E. Di Nitto, D. Ardagna, A. Solberg, G. Casale. Towards Multi-Clouds Engineering, in *Proc. of CrossCloud workshop*, Toronto, 2014.

O5. PASTA S. Pacheco-Sanchez, G. Casale, B. Scotney, S. McClean, G. Parr. A Case Study of Demand Estimation for a Multi-Threaded ERP Application, in *Proc. of PASTA workshop*, Sep 2011.

O4. PASTA G. Casale, M. Tribastone. Process-Algebraic Modeling of Priority Queueing Networks, in *Proc. of PASTA workshop*, Sep 2010.

O3. ICPE J. Rolia, D. Khrishnamurthy, G. Casale, S.Dawson. APE: Automated Performance Evaluation of Customized Services, in *Proc. of ACM/SPEC ICPE*, San Jose, California, 3–14, Jan 2010. (Invited keynote)

O2. NGS G. Casale, N. Mi, E. Smirni. Versatile Models of Systems Using MAP Queueing Networks, in *Proc. of NSF-NGS Workshop* (at IPDPS 2008), IEEE Press, 2008.

O1. NGS E. Smirni, Q. Zhang, N. Mi, A. Riska, G. Casale. New Results on the Performance Effects of Autocorrelated Flows in Systems, in *Proc. of NSF-NGS Workshop* (at IPDPS 2007), IEEE Press, 2007.

WORK-IN-PROGRESS PAPERS

W24. WOSP-C R. Dobres, Z. Niu, G. Casale. Approximating Fork-Join Systems via Mixed Model Transformations, short paper in *Proc. of WOSP-C workshop*, May 2024.

W23. ICPE Z. Li, G. Casale. Matrix Network Analyzer: a New Decomposition Algorithm for Phase-type Queueing Networks, short paper in *Proc. of ICPE Emerging Research Track*, May 2024.

W22. AAMAS M. Sheldon, D. Paccagnan, G. Casale. Cournot Queueing Games with Applications to Mobility Systems, short paper in *Proc. of AAMAS*, May 2024.

W21. MAMA J. Bor, G. Casale, W. Knottenbelt, E. Smirni, A. Stathopoulous. Fitting with matrix exponential mixtures generated by discrete probabilistic scaling, *Special Issue on MAMA Workshop*, ACM Perf. Eval. Rev., 1(2): 15-17, 2023.

W20. SIGMETRICS S. Tuli, G. Casale, N. Jennings. Learning to Dynamically Select Cost Optimal Schedulers in Cloud Computing Environments, in ACM Performance Evaluation Review, Special Issue on ACM SIGMETRICS / Performance 2022 Posters, Sep 2022.

W19. ML4SYS S. Tuli, G. Casale, N. Jennings. Generative Optimization Networks for Memory Efficient Data Generation, *Machine Learning for Systems* workshop at NeurIPS 2021, Short paper, Dec 2021, [acc. rate: $\mathbf{9.2}\%$, 347 submissions].

W18. CLOUD L. Zhu, G. Giotis, V. Tountopoulos, G. Casale. RDOF: Deployment Optimization for Function as a Service, in *Proc. of CLOUD*, 7 pages, Sep 2021 (short paper).

W17. TOSME M. Sheldon, G. Casale. TauSSA: Simulating Markovian Queueing Networks with Tau Leaping, *TOSME* workshop at Performance 2021, Short paper, Nov 2021.

W16. VALUETOOLS A. Alnafessah, G. Casale. TRACK: Optimizing Artificial Neural Networks for Anomaly Detection in Spark Streaming Systems, *Proc of. Valuetools*, Short paper, May 2020, ACM.

W15. MAMA G. Casale, P. G. Harrison, O. W. Hong. Novel Solutions for Closed Queueing Networks with Load-Dependent Stations, *Special Issue on MAMA Workshop*, ACM Perf. Eval. Rev., Sep 2019, ACM.

W14. MAM G. Casale, G. Horvath, J.F. Pérez. A Matrix-Analytic Approximation for Closed Queueing Networks with General FCFS Nodes, the 9th International Conference on Matrix-Analytic Methods in Stochastic Models (MAM9), Budapest, Jun 2016.

W13. MASCOTS K. Molka, G. Casale. Efficient Memory Occupancy Models for In-Memory Databases, *Proc of. IEEE MASCOTS*, Sep 2016, ACM, Short paper.

W12. MAMA W. Wang, G. Casale. Maximum Likelihood Estimation of Closed Queueing Network Demands from Queue Length Data, *Special Issue on MAMA Workshop*, ACM Perf. Eval. Rev., Sep 2015, ACM.

W11. MISE G. Casale, *et al.* DICE: Quality-driven development of data-intensive cloud applications, in *Proc. of Modeling in Software Engineering (MISE) Workshop*, May 2015, ACM.

W10. MICAS W. Wang, G. Casale. Evaluating Weighted Round Robin Load Balancing for Cloud Web Services, in *Proc. of Management of Resources and Services in Cloud and Sky Computing (MICAS) Workshop*, Sep 2014, IEEE Press.

W9. MICAS J. Moschetta, G. Casale. OFBench: an Enterprise Application Benchmark for Cloud Resource Management Studies, in *Proc. of Management of Resources and Services in Cloud and Sky Computing (MICAS) Workshop*, Sep 2012, IEEE Press.

W8. MISE D. Ardagna, E. Di Nitto, G. Casale, *et al.* MODAClouds: A model-driven approach for the design and execution of applications on multiple Clouds, in *Proc. of Modeling in Software Engineering (MISE) Workshop*, May 2012, ACM.

W7. DCPERF G. Casale, S. Kraft, D. Krishnamurthy. A Model of Storage I/O Performance Interference in Virtualized Systems, in *Proc. of 1st Data Center Performance Workshop*, Jun 2011.

W6. NMSC G. Casale, P. G. Harrison, M. Vigliotti. Product-Form Approximation of Tandem Queues via Matrix Geometric Methods, in *Proc. of Numerical Solution of Markov Chains (NSMC) Workshop*, Sep 2010.

W5. HOTMETRICS G. Casale, A.Kalbasi, D.Khrishnamurthy, J. Rolia. Automatically Generating Bursty Benchmarks for Multi-Tier Systems, *Special Issue on the 2nd ACM HOTMETRICS workshop*, ACM Perf. Eval. Rev. 37(3):32–37, Dec 2009, ACM.

W4. HOTMETRICS A.Riska, N.Mi, G. Casale, E. Smirni. Feasibility Regions: Exploiting Trade-offs between Power and Performance in Disk Drives, *Special Issue on the 2nd ACM HOTMETRICS workshop*, ACM Perf. Eval. Rev. 37(3):43–48, Dec 2009, ACM.

W3. HOTMETRICS G. Casale, N. Mi, L. Cherkasova, E. Smirni. How to Parameterize Models with Bursty Workloads, *Special Issue on the 1st ACM HOTMETRICS workshop*, ACM Perf. Eval. Rev. 36(2):38-44, Jun 2008.

W2. MAMA G. Casale. CoMoM: Class-Oriented Evaluation of Multiclass Models, *Special Issue on MAMA Workshop*, ACM Perf. Eval. Rev. 36(2):38-44, ACM, Sep 2008, ACM.

W1. MAMA G. Casale, E.Z. Zhang, E. Smirni. Characterization of Moments and Autocorrelation in MAPs, *Special Issue on MAMA Workshop* in ACM Perf. Eval. Rev., 35(2):27–29, Sep 2007, ACM.

DEMOS

D8. ICPE G. Casale. Automated multi-paradigm analysis of extended and layered queueing models with LINE, in *Proc. of ICPE*, 2019. **Best Demo Award**.

D7. QEST C. Li, T. Altamimi, M. H. Zargari, G. Casale, and D. Petriu. Tulsa: A Tool for Transforming UML to Layered Queueing Networks for Performance Analysis of Data Intensive Applications, in *Proc. of QEST*, 2017.

D6. ICPE G. Casale, M. Cazzoli, J. Shuai, V.S. Lopes, G. Serazzi, L. Zhu. Generalized Synchronizations and Capacity Constraints for Java Modelling Tools, in *Proc. of ACM/SPEC ICPE*, 2017.

D5. MASCOTS D. J. Dubois, G. Casale. Performance Prediction for Burstable Cloud Resources, in *Proc. of VALUETOOLS*, Nov 2016.

D4. QUDOS W. Wang, J. Pérez, G. Casale. Filling the Gap: a Tool to Automate Parameter Estimation for Software Performance Models, in *Proc. of QUDOS*, 2015, ACM.

D3. ICPE S. Spinner, G. Casale, X. Zhu, S. Kounev. LibReDE: A Library for Resource Demand Estimation, in *Proc. of ICPE*, 2014, ACM.

D2.	MAM	G. Casale, E.Smirni. KPC-Toolbox: Fitting Markovian Arrival Processes and Phase-Type Distributions with MATLAB, ACM Perf. Eval. Rev.,39(4):47, Mar 2012.
D1.	SIGMETRICS	E.Z.Zhang, G. Casale, E.Smirni. KPC-Toolbox: Best Recipes Toward Automatization of Workload Fitting, ACM Perf. Eval. Rev. 36(2):134-136, Sep 2008.

TUTORIALS

T8.	ICPE	S. Tuli, G. Casale. Optimizing the Performance of Fog Computing Environments Using AI and Co-Simulation, at *ACM/SPEC ICPE 2022*.
T7.	ICPE	A. van Hoorn, A. Gias, L. Zhu, G. Casale, T. Düllmann, M. Wurster. Performance Engineering for Microservices and Serverless Applications: the RADON approach, at *ACM/SPEC ICPE 2020* (invited to be repeated at ICPE 2021).
T6.	PERFORMANCE	G. Casale, G. Serazzi, L. Zhu. Performance Evaluation with Java Modelling Tools: a Hands-On Introduction, at *IFIP Performance 2017*.
T5.	ICPE	G. Casale, S. Spinner, W. Wang. Automated Parameterization of Performance Models from Measurements, at *ACM/SPEC ICPE 2016*.
T4.	UCC	D.Ardagna, N. Ferry, G. Casale. Model-Driven Management of Multi-Cloud Applications, at *IEEE/ACM UCC 2014*.
T3.	ASE	D.Ardagna, N. Ferry, G. Casale, M.Almeida, J.F.Pérez. MDD-CLOUD - Model Driven Design of Cloud Applications with "a priori" Quality of Service Guarantees, at *IEEE/ACM ASE 2014*.
T2.	SIGMETRICS	G. Casale. Building Accurate Workload Models using Markovian Arrival Processes, at *ACM SIGMETRICS 2011*, Jun 2011.
T1.	ICPE	G. Casale, G. Serazzi. Quantitative System Evaluation with Java Modelling Tools, at *ACM/SPEC ICPE 2011*, Mar 2011.

TEACHING

CURRENT TEACHING

2021–	Probability and Statistics, Imperial College London.
2021–	Scheduling and Resource Allocation, Imperial College London.

PAST TEACHING

2015–23	Simulation and Modelling, Imperial College London.
2014–20	Operations Research, Imperial College London.
2018–20	Performance Engineering, Imperial College London.
2016–17	Performance Engineering (Half module), Imperial College London.
2014–15	Performance Analysis (Half module), Imperial College London.
2011–14	Performance Analysis, Imperial College London.
2005–06	Capacity planning (Dimensionamento dei Sistemi Informatici), Politecnico di Milano.
2003	Enterprise Systems (Impianti Informatici), Politecnico di Milano, Italy.

PROJECT SUPERVISION

A UG final project is equivalent to a BSc thesis. A MSc final project is equivalent to a MSc thesis. Group projects are 3-month projects with a development focus. ISO and Topics projects are literature surveys. MRes projects are 4-month research projects that award the same number of credits of a course.

2010–20	23 UG final projects, 1 MRes project, 37 MSc final projects, 11 UG group projects, 3 MSc group project, 3 ISO projects, 1 Topics project.
2003–06	15 final projects for UG and MSc students at Politecnico di Milano, Italy.

Supervision and mentoring

Research Supervision

Postdoctoral researchers (7): Eleftherios Anastasiadis (2021), Jose Antonio Perusquia Cortes (2021), Chen Li (2016-2018), Pooyan Jamshidi (2015-16), Rasha Osman (2015), Daniel J. Dubois (2014-16), Juan F. Pérez (2013-15).

Pre-doc researchers (11): Adam Eljaafari (2023), James Stadler (2023), Harshit Mawandia (2022), Yicheng Gao (2021), Matthew Sheldon (2021), Sajal Mittal (2021), John Yao (2019), Lulai Zhu (2019-21), Vitor S. Lopez (2016), Shuai Jiang (2015), Tatiana Ustinova (2015-17).

PhD students (17): Zhuoyuan Li (2024-), Wenxiang Luo (2023-), Yaqi Zhu (2023-), Yichong Chen (2022-), Julianna Bor (2022-; co-supervision), Matthew Sheldon (2021-), Yicheng Gao (2020-), Zifeng Niu (2020-), William Plumb (2020-), Shreshth Tuli (2020-23; co-supervision; **SPEC dissertation award winner**), Runan Wang (2019-; co-supervision), Ahmad Alnafessah (2017-22), Alim Gias (2017-22), Lulai Zhu (2017-22), Salvatore Dipietro (2015-19), Karsten Molka (2013-17), Weikun Wang (2013-16).

Industrial PhD mentorships (2): Stephan Kraft (2009-11), Sergio Pacheco-Sanchez (2009-11).

University Service

PhD Vivas and Juries

Internal examiner (9): Imperial College London (2024, 2023, 2022, 2020, 2018, 2×2017, 2014, 2013, 2012).

External examiner (17): University of Melbourne (Australia, 2025), Ecole Polytechnique (France, 2023), University of Galway (Ireland, 2023), University of Antwerp (Belgium, 2023), University of Lund (Sweden, 2022), University of Manchester (UK, 2022), University of Rome Tor Vergata (Italy, 2021), University of Florence (Italy, 2020), University of Melbourne (Australia, 2020), University of Calgary (Canada, 2017), Umeå University (Sweden, 2017), Delft University (The Netherlands, 2017), Carleton University (Canada, 2016), Newcastle University (UK, 2015), University of Turin (Italy, 2013), Politecnico di Milano (Italy, 2012), DTU Copenhagen (Denmark, 2010).

Interim evaluations (3): University of Pisa (2022), University of Florence (2021), Umeå University (2016).

Administration

All the activities below have been carried out at Imperial College London.

Ongoing roles and activities:

2018-	PhD admissions tutor, Department of Computing.
2020-	Academic staff representative (non-professorial), Department of Computing.
2023-	Board member, Science and Solutions for a Changing Planet DTP, Grantham Institute.
2018-	Panel Member for Computing, Faculty of Enginering President's PhD scholarships panel.

Past roles and activities:

2023	JP Morgan PhD fellowships shortlisting panel, Faculty of Engineering.
2023	Computing workload working group, Department of Computing.
2022	Work location framework committee, Department of Computing.
2017-18	PhD admissions tutor, HiPEDS Centre for Doctoral Training.
2016-18	PhD scholarship committee, HiPEDS Centre for Doctoral Training.
2015-	PhD scholarship committee, Department of Computing.
2017	Panel co-chair, Google European PhD fellowships departmental shortlisting panel.
2016-18	Deputy PhD admissions tutor, Department of Computing.

2015	Departmental EPSRC doctoral prize committee.
2013-16	Internal lunch seminars coordinator.
2012	Panel member, PhD Careers Event.
2012	Judge for Google PhD Poster Competition.