

Curriculum vitae

Salvatore ORLANDO, PhD

Full Professor – Università Ca' Foscari Venezia

Head of Dept. of Environmental Science, Informatics and Statistics (DAIS)

Member of the Academic Senate of Ca' Foscari University

Vice President GRIN (National Ass. of Professors in Computer Science).

Associate Researcher at ISTI - CNR of Pisa, HPC Lab.

Date and Place of Birth: *June 14th, 1961 – Cosenza (Italy)*

Fiscal Code: *RLN SVT 61H14 D086A*

Home Address: *Via Raffaello Sanzio, 18/A – 56122 Pisa (Italy)*

Phone: +39 349 8625043

Email: orlando@unive.it

Homepage: <http://www.dais.unive.it/~orlando>

Biographical sketch.

Salvatore Orlando – Laurea/MSc (summa cum laude, 1985) and PhD (1991) in Computer Science, University of Pisa – was a post-doc fellow of the HP labs, and a post-doc fellow of the University of Pisa. In 1994 he joined as an assistant professor Ca' Foscari University of Venice, then associate professor since 2000, and finally full professor of Computer Science since 2017.

Salvatore Orlando is an adjunct researcher at ISTI-CNR of Pisa, and he was a visiting scientist at the School of Computer Science of Carnegie Mellon University (USA), and at University of Paderborn and Heinz Nixdorf Institut (Germany).

Since October 2020 Salvatore Orlando has been appointed as Head of Department of Environmental Science, Informatics and Statistics of Ca' Foscari University. He was the Coordinator of the BSc and MSc Programs in Computer Science of Ca' Foscari University since 2010 till 2019. In the period 2013-2014 he was also the delegate of the Rector for national and international university orientation.

His research interests include the design of scalable algorithms for data/web mining problems, efficient and effective solutions for information retrieval problems, distributed and P2P systems for information retrieval, parallel/distributed systems and programming environments. In the last years, he supervised the research work of many doctoral students of the PhD Program in Computer Science at Ca' Foscari University, 11 of which already completed their studies.

Salvatore Orlando (Google h-index = 33, Scopus h-index = 22) has published over 150 papers in international journals, book chapters, and conference/workshop proceedings. He has co-chaired conferences (EuroPVM/MPI, SASO, EDBT/ICDT), several conference tracks, and workshops. He has served on the program committees of international conferences, among which several premier conferences (ACM KDD, IEEE ICDM, ACM CIKM, SIAM SDM, ECML/PKDD) and many others, such as Europar, IPDPS, INFOSCALE, CCGRid, SASO, etc. Finally, he gave the tutorial "Mining query logs" at RuSSIR-2011 Summer School, at the European Conf. on Information Retrieval (ECIR) 2009, and at Int.l Conf. Web Intelligence (WI) 2009.

He has been local coordinator of and has participated in several National and European projects.

Education and Work experience

- 1980** High School Diploma - vote 60 of 60, Liceo Scientifico.
- 1985** Laurea degree in *Scienze dell'Informazione* (M.Sc. in Computer Science) - magna cum laude - University of Pisa.
- 1986** Consultant of a software house, Numerica Progetti s.r.l., for developing applications in the field of telecommunications, computer graphics, and multimedia data-bases.
- 1987** National military service.
- 1987-1990** Ph.D student of the Dept. of Computer Science of the University of Pisa. Received in 1991 the Ph.D. in Computer Science with a dissertation entitled "*Exploiting Synchronous Models for Distributed Memory Multiprocessor Architectures*" - Supervisor: Prof. M. Vanneschi - External referees: Prof. T. Hey (University of Southampton) e Dr. D. May. (Inmos Labs).
- 1991-1992** Fellow of Hewlett Packard - Pisa Science Center, for working on a joint project between Dept. of Computer Science of University of Pisa and Hewlett Packard. Title of the research: *Development of hardware/software environments for massively parallel architectures.*
- 1993** Post-doctoral fellow of the Dept. of Computer Science of the Pisa University.
- 1994-2000** Assistant Professor - Università Ca' Foscari Venezia.
- 2000-2017** Associate Professor - Università Ca' Foscari Venezia.
- 2017-today** Full Professor - Università Ca' Foscari Venezia.

Bibliometric Indices

	Google Scholar	Scopus
citations	4983	2110
h-index	37	23
i10-index	102	54

Awards

- 2016: "Best Student Paper Award" at DocEng 2016: ACM Symposium on Document Engineering (Rating GII-GRIN-SCIE: B) [46]
- 2015: "Best Paper Award" at SIGIR 15: ACM SIGIR Conference on Research and Development in Information Retrieval (Rating GII-GRIN-SCIE: A++) [55].
- 2015: "Notable Article in Computing" according to *ACM Computing Reviews* [55].
- 2013: "Notable Article in Computing" according to *ACM Computing Reviews* [17].
- 2011: "Best Paper Award Runner-up" at WSDM 2011: 4th ACM Conference on Web Search and Data Mining (Rating GII-GRIN-SCIE: A+) [69]

Teaching Activities

Academic Year	MSc, BSc, PhD Courses
1990-91	Series of lectures within the PhD course entitled “Parallel Architectures” by M. Vanneschi – Università di Pisa
1995-96, 96-97	Series of lectures for the courses: (a) Computer Architecture (BSc), (b) Lab of Computer Architecture (BSc) – Università Ca’ Foscari Venezia
1997-98, 98-99, 99-2000	Instructor of the following courses: (a) Distributed Systems (BSc), (b) Lab of Computer Architecture (BSc) – Università Ca’ Foscari Venezia
2000-01	Instructor of the following courses: (a) Computer Architecture (BSc), (b) Parallel Computing (PhD) – Università Ca’ Foscari Venezia
2001-02	Instructor of the following courses: (a) Computer Architecture A and B (BSc), (b) Parallel Computing with Lab (MSc) – Università Ca’ Foscari Venezia
2002-03, 03-04, 04-05, 05-06, 06-07, 07-08, 08-09	Instructor of the following courses: (a) Computer Architecture, mod I and II (BSc), (b) Parallel Computing with Lab (MSc), (c) Data Mining (MSc) – Università Ca’ Foscari Venezia
2009-10, 10-11, 11-12, 12-13	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) High Performance Computing (MSc), (c) Data and Web Mining (MSc) – Università Ca’ Foscari Venezia
2013-14	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) High Performance Computing (MSc) – Università Ca’ Foscari Venezia
2014-15	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) Data and Web Mining (MSc), (c) Information Retrieval (PhD) – Università Ca’ Foscari Venezia
2015-16, 16-17	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) High Performance Computing (MSc), (c) Information Retrieval and Web Search (MSc) – Università Ca’ Foscari Venezia
2017-18, 18-19, 19-20	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) Introduction to Coding (BA), (c) Information Retrieval and Web Search (MSc) – Università Ca’ Foscari Venezia
20-21	Instructor of the following courses: (a) Computer Architecture, mod II (BSc), (b) Computer Programming and Data Management (MA), (c) Information Retrieval and Web Search (MSc) – Università Ca’ Foscari Venezia
21-22	Instructor of the following courses: (a) Computer Programming and Data Management (MA), (b) Information Retrieval and Web Search (MSc) – Università Ca’ Foscari Venezia

Finally, Prof. Orlando was the supervisor of numerous BSc and MSc theses of the Computer Science programs at Università Ca’ Foscari Venezia, and of some other theses at Università di Pisa.

Scientific Supervision Activities

In the last years Prof. Orlando supervised many PhD students of the Doctoral Program in Computer Science at Ca’ Foscari. In particular, the following ten students got their PhD degree under his guidance: Paolo Palmerini (XVI ciclo, PhD, 2003), Claudio Silvestri (XVIII cycle, PhD, 2005), Francesco Lelli (XIX cycle, PhD 2007), Matteo Mordacchini (XIX cycle, PhD 2007), Claudio Lucchese (XX cycle, PhD 2008), Fernando Josè Braz (XXI cycle, PhD 2009), Gabriele Tolomei (XXIII cycle, PhD 2011), Andrea Triossi (XXIV cycle, PhD 2012), Daniele Broccolo (XXV cycle, PhD 2014), Francesco Lettich (XXVII cycle, PhD 2015), Mohsen Pourvali (XXIX cycle, 2016), Bruno Quintavalle (XXXI cycle, 2019), Francesco Busolin (XXXVII cycle).

In addition, Prof. Orlando supervised the activities of many post-doc fellows, funded by national and international projects and by research agreements with public and private entities.

University Services

- 1996-1998: member of the governing board of CICS (computer center) of Università Ca’ Foscari Venezia.

- 2000-today: member of the board of the doctoral program in Computer Science of Ca' Foscari.
- 2008-10: deputy head of Department of Computer Science at Ca' Foscari.
- 2000-today: member of the Teaching Committee of the BSc and MSc programs in Computer Science at Ca' Foscari.
- 2010-today: appointed as Head of the Teaching Committee of the BSc and MSc programs in Computer Science at Ca' Foscari.
- 2010-today: member of the governing board of CINI (National Inter-university Consortium for Informatics).
- 2011-today: delegate of the head of the DAIS department for teaching activities.
- 2013-14: delegate of the Rector for national and international university orientation.

Research projects

Coordination activity

- Leader of WP 5 - "Development and Application of Tools", within the project Ecomobility (2018-2019) - Interreg Italy-Croatia CBC Programme.
- Leader of Activity A2.4 - "Implementation of knowledge management systems", within the project Smart Cities: *ADAPT: Accessible Data for Accessible Proto-Types in Social Sector* - SCN_00447, funded by MIUR (2015-2017).
- Coordinator of the Ca' Foscari research unit of the national project: *TETRis - Servizi innovativi Open Source su TETRA* - PROGETTO PON01_00451" (2011-13)
- Coordinator of the Ca' Foscari research unit of the national project: *MOTUS - MObility and Tourism in Urban Scenarios* (2009-1012), funded by Ministero dello Sviluppo Economico - Program "Industria 2015 - Mobilità Sostenibile".
- Coordinator of the project: *Advanced Services for the info-mobility and tourism* (2009), funded by Fondo Sociale Europeo (FSE).
- Coordinator of the Ca' Foscari research unit of the national project PRIN 2004 *GeoPKDD: Geographic Privacy-aware Knowledge Discovery and Delivery* (2005-2006).
- Coordinator of a biannual research contract (2001-2003) between Tektronics S.p.A. and Università Ca' Foscari Venezia, focused on *Data Mining Analysis of Telephonic Traffic Data*.

Main participation to projects

- EU Coordination Action "MODAP (Mobility, Data Mining, and Privacy)" (2009-2012)
- EU ICT Policy Support programme: "ASSETS: Advanced Search Service and Enhanced Technological Solutions for the European Digital Library" (2010-12)
- IST EU FP6 Project: "SAPIR - Search In Audio Visual Content Using Peer-to-peer Information Retrieval" (2007-2009).
- European Research Network - IST FP6: "CoreGRID: Foundations, Software Infrastructures and Applications for large scale distributed, Grid and Peer-to-Peer Technologies" (2004-2008).
- European Integrated Project (IP) - IST FP6: "NextGRID: The Next Generation Grid" (2004-2007).
- National project MURST (2002-2003): "Architetture Software per infrastrutture di rete ad accesso eterogeneo (SAHARA)".
- National FIRB Project: "Grid.it: Piattaforme abilitanti per griglie computazionali a elevate prestazioni orientate a organizzazioni virtuali scalabili" (2003-2005).
- National strategic project: "Grid Computing: tecnologie abilitanti e applicazioni per eScience" (2002-2004).
- National strategic project: "ECD: Technologies and Services for Enhanced Contents Delivery" (2002-2004)
- National project ASI-PQE2000 "Sviluppo di applicazioni di osservazione della terra mediante sistemi e tool di calcolo ad alte prestazioni"

Scientific Activities

Chair / Co-Chair

- 2017: General Chair of EDBT/ICDT 2017.

- 2011: Program Global Chair of the topic "Parallel and Distributed Data Management" at Europar 2011.
- 2010: Co-Chair of the topic "Parallel and Distributed Databases, Data Mining, and Knowledge Discovery" of Europar 2004.
- 2010: Program Co-Chair of the topic "Parallel and Distributed Data Management" of Europar 2010.
- 2009-14: Awards Co-Chair of the High Performance Computing & Simulation Conference (HPCS)
- 2010: Program Co-Chair of SDIA 2010 - Scalable Data Intensive Applications - Special session of PDP 2010.
- 2008: Local Chair of the 2nd IEEE Int.l Conf. on Self-Adaptive and Self-Organizing Systems (SASO 2008).
- 2008: Workshop Chair of the 2nd Int. Conf. on Scalable Information Systems (INFOSCALE'08).
- 2008: Special Sessions & Workshops Co-Chair of the 2008 High Performance Computing and Simulation Conference (HPCS 2008).
- 2005: Program Co-Chair of the 8th High Performance and Distributed Mining Workshop (HPDM'05) - SIAM Conf. on Data Mining - 2005.
- 2003: Co-Chair of the Program Committee of EuroPVM/MPI '03 – Venezia – September 29 - Ottobre 2, 2003.

Member of Program Committees (main conferences)

- 2009,2016: ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD).
- 2008-13,2016: IEEE Int.l Conf. on Data Mining (ICDM).
- 2006-08: SIAM Conference on Data Mining (SDM).
- 2008,2012,2014-15: European Conf. on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD).
- 2012-13,2016: ACM Conference on Information and Knowledge Management (CIKM).
- 2015: IEEE International Parallel & Distributed Processing Symposium (IPDPS).
- 2004, 2010-11: European Conference on Parallel and Distributed Computing (Euro-Par).
- 2004,2005,2006,2007,2008: Euro PVM/MPI Conference (PVM/MPI).
- 2004-08: International Conference on Computational Science (ICCS).
- 2005-08,2011,2013-16: ACM Symposium on Applied Computing (SAC) - Track on DM
- 2007-16: Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD).
- etc.

Reviewer for the Journals:

- IEEE TKDE, IEEE TPDS, ACM TWEB, Springer KAIS, Elsevier Pattern Recognition Letters, etc.

Member of the editorial board:

- Journal of Data Mining & Digital Humanities (JDMDH), Episciences.org – Open Access, sponsored by CNRS, INRA and Inria.

Seminars and Tutorial

- 1997: Schloss Dagstuhl - Invited seminar at the event "Theory and Practice of Higher-Order Parallel Programming".
- 1999: Schloss Dagstuhl - Invited seminar at the event "High Level Parallel Programming: Applicability, Analysis and Performance".
- 2003: Invited seminar at the school on *Science and Applications of Advanced Computing Paradigms*, Università di Padova.
- 2009: Tutorial "Mining Query Logs" (with F. Silvestri) - European Conf. on Information Retrieval (ECIR 2009) [74].
- 2009: Tutorial "Query Log Analysis for Enhancing Web Search". The 2009 IEEE / WIC / ACM International Conferences on Web Intelligence (WI 09).
- 2011: Tutorial "Mining Query Logs to Improve Search Engines Operations" (with R. Perego & F. Silvestri). Web of Data: Joint RuSSIR/EDBT Summer School RuSSIR 2011.

- 2011: Invited seminar at Yandex (Moscow, Russia).
- 2013: Invited seminar at University of Glasgow (UK)
- 2014: Invited seminar at Linköping University (UK).
- 2015: Invited seminar at Università della Calabria (IT).
- 2017: Invited paper at International Conference on High Performance Computing & Simulation.

Publications

International Journal Papers

- [1] S. Calzavara, C. Lucchese, F. Marcuzzi, S. Orlando: Feature partitioning for robust tree ensembles and their certification in adversarial scenarios. *EURASIP J. Inf. Secur. - Springer 2021(1):12*
- [2] S. Calzavara, C. Lucchese, G. Tolomei, S.A. Abebe, S. Orlando: TREANT: Training Evasion-Aware Decision Trees. *Data Mining and Knowledge Discovery (DAMI)*, Springer, 34, pages 1390-1420 (2020).
- [3] E. Merico, M. Conte, F.M. Grasso, D. Cesari, A. Gambaro, A., E. Morabito, E. Gregoris, S. Orlando, A. Alebic-Juretic, V. Zubak, B. Mifka, D. Contini. Comparison of the impact of ships to size-segregated particle concentrations in two harbour cities of northern Adriatic Sea *ENVIRONMENTAL POLLUTION*, pp. 115175 (ISSN 0269-7491), 2020
- [4] F. Lettich, C. Lucchese, F. M. Nardini, S. Orlando, R. Perego, N. Tonello, and R. Venturini Parallel traversal of large ensembles of decision trees. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 10.1109/TPDS.2018.2860982, Volume: 30, Issue: 9, 2019.
- [5] M. Pourvali, S. Orlando. Enriching Documents by Linking Salient Entities and Lexical-Semantic Expansion in *JOURNAL OF INTELLIGENT SYSTEMS* (ISSN 0334-1860), Volume 29: Issue 1, 2020.
- [6] M. Pourvali, S. Orlando, H. Omidvarborna Topic Models and Fusion Methods: a Union to Improve Text Clustering and Cluster Labeling. *IJIMAI* 5(4): 28-34 (2019)
- [7] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, F. Silvestri, S. Trani. X-CLEAVER: Learning Ranking Ensembles by Growing and Pruning Trees. *ACM Transactions on Intelligent Systems and Technology (TIST)*, (ISSN:2157-6904), 9(6): 62:1-62:26 (2018).
- [8] S. Trani, C. Lucchese, R. Perego, D.E. Losada, D. Ceccarelli, S. Orlando. SEL: A unified algorithm for salient entity linking. *Computational Intelligence*, Wiley, (ISSN 1467-8640), 10.1111/coin.12147, Volume 34, Issue 1, pp. 2-29, February 2018.
- [9] F. Lettich, S. Orlando, C. Silvestri, C.S. Jensen. Manycore GPU processing of repeated range queries over streams of moving objects observations. *CONCURRENCY AND COMPUTATION*, Wiley, (ISSN 1532-0626), <http://dx.doi.org/10.1002/cpe.3881>, Volume 29, Issue 4 25 February 2017.
- [10] D. Dato, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonello, and R. Venturini. Fast ranking with additive ensembles of oblivious and non-oblivious regression trees. *ACM TOIS - Transactions on Information Systems*, ISSN: 1046-8188, 35(2), pag. 15:1-15:31 (2016)
- [11] G. Capannini, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonello. Quality versus efficiency in document scoring with learning-to-rank models. *INFORMATION PROCESSING & MANAGEMENT*, Elsevier, Article in press, Published online in 2016 (ISSN 0306-4573), <http://dx.doi.org/10.1016/j.ipm.2016.05.004>
- [12] F. Lettich, L.O. Alvares, V. Bogorny, S. Orlando, A. Raffaetà, C. Silvestri. Detecting avoidance behaviors between moving object trajectories. *DATA & KNOWLEDGE ENGINEERING*, Elsevier vol. 102, pp. 22-41, 2016 (ISSN 0169-023X)
- [13] Stefano Calzavara, Gabriele Tolomei, Andrea Casini, Michele Bugliesi, Salvatore Orlando. A Supervised Learning Approach to Protect Client Authentication on the Web. *ACM TWEB - Transactions on the Web*, 9(3):15, pp. 1-30, ISSN: 1559-1131 (2015)
- [14] C. Lucchese, S. Orlando, and R. Perego. A unifying framework for mining approximate top-k binary patterns. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 26(12): 2900-2913 (2014).
- [15] L. Leonardi, S. Orlando, A. Raffaetà, A. Roncato, C. Silvestri, G. Andrienko, and N. Andrienko. A general framework for trajectory data warehousing and visual OLAP. *Geoinformatica*, Springer, 18(2): 273-312 (2014).
- [16] F. Giummolè, S. Orlando, and G. Tolomei. A Study on Microblog and Search Engine User Behaviors: How Twitter Trending Topics Help Predict Google Hot Queries. *HUMAN, ASE Science Journal*, Vol. 2, No. 3, pp. 195-209, 2013, ISBN: 978-1-62561-001-0 (*Special issue - Best papers of the IEEE/ASE SocialCom2013 Conference*).
- [17] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, G. Tolomei. Discovering Tasks from Search Engine Query Logs. *ACM TOIS - Transactions on Information Systems*, 2013, Vol. 31, No. 3, pp. 1-43. , doi: 10.1145/2493175.2493179, ISSN: 1046-8188. *Selected by ACM Computing Reviews as a notable article in computing in 2013.*

- [18] F. Falchi, C. Lucchese, S. Orlando, R. Perego, F. Rabitti. Similarity Caching in Large-Scale Image Retrieval, *Information Processing and Management*, Elsevier, vol. 48, pp. 803-818, 2012, ISSN: 0306-4573.
- [19] A. Raffaetà, L. Leonardi, G. Marketos, G. Andrienko, N. Andrienko, E. Frentzos, N. Giatrakos, S. Orlando, N. Pelekis, A. Roncato, C. Silvestri. Visual Mobility Analysis using T-Warehouse, *International Journal of Data Warehousing and Mining*, IGI Global, ISSN: 1548-3924, Volume 7, Number 1, Pages 1–27, 2011.
- [20] C. Lucchese, C. Mastroianni, S. Orlando, D. Talia. Mining@home: Towards a Public Resource Computing Framework for Distributed Data Mining, *Concurrency & Computation: Practice and Experience*, Wiley, ISSN: 1532-0634, Volume 22, Issue 5, 2010, Pages 658-682.
- [21] F. Bonchi, F. Giannotti, C. Lucchese, S. Orlando, R. Perego, R. Trasarti. A Constraint-based Querying System for Exploratory Pattern Discovery, *Information Systems*, Elsevier, Volume 34, Issue 1, March 2009, pages: 3–27.
- [22] S. Orlando, R. Orsini, A. Raffaetà, A. Roncato, C. Silvestri. Trajectory Data Warehouses: Design and Implementation Issues, *Journal of Computing Science and Engineering (JCSE)*. KIISE Publ., Vol. 1, No. 2, pp. 240-261, Dicembre 2007 (*Special issue - Best papers of the 9th Int.l Conf. on Data Warehousing and Knowledge Discovery (DaWaK '07)*)
- [23] F. Lelli, E. Frizziero, M. Gulmini, G. Maron, S. Orlando, A. Petrucci, S. Squizzato. The many faces of the integration of instruments and the grid, *Int. Journal of Web and Grid Services (IJWGS)*, ISSN: 1741-1106, 3(3). pp. 239–266, 2007, Inderscience Publishers.
- [24] M. Marzolla, S. Orlando, M. Mordacchini. P2P Systems for Resource Discovery in a Dynamic Grid, *Parallel Computing Journal*, 33, pp. 339–358, 2007, Elsevier Science.
- [25] C. Silvestri, and S. Orlando. Approximate Mining of Frequent Patterns on Streams, *Int. Journal of Intelligent Data Analysis*, IOS Press, Vol. 11, No. 1, pp.49–73, 2007, ISSN 1088-467X.
- [26] F. Lelli, S. Orlando, G. Maron, S. Pinter. Bringing instruments into a Grid: An empiric approach. *WSEAS Transactions in Computers*, vol. 6, 2007, p. 153-159, ISSN: 1109-2750.
- [27] F. Silvestri, D. Puppini, D. Laforenza. Toward a Search Architecture for Software Components, *Concurrency & Computation: Practice & Experience*, ISSN: 1532-0634, Vol. 18, pp. 1317-1331, Wiley, 2006.
- [28] T. Fagni, R. Perego, F. Silvestri, and S. Orlando. Boosting the Performance of Web Search Engines: Caching and Prefetching Query Results by Exploiting Historical Usage Data. *ACM TOIS*, vol. 24, 2006, pp. 51-78.
- [29] C. Lucchese, S. Orlando, and R. Perego. Fast and Memory Efficient Mining of Frequent Closed Itemsets, *IEEE Transactions on Knowledge and Data Engineering*, January 2006 (Vol. 18, No. 1), pp. 21-36.
- [30] S. Orlando, P. Palmerini, and R. Perego. Mixed Data and Task Parallelism with HPF and PVM. *CLUSTER COMPUTING: The Journal of Networks, Software and Applications*, Baltzer Science Publishers, Vol.3, n.3, 2000, pp. 201–213.
- [31] S. Orlando and R. Perego. *COLT_{HPF}*: a Run-Time Support for the High-Level Coordination of HPF Tasks. *Concurrency: Practice and Experience*, ISSN: 1532-0634, Wiley, Vol. 11(8), pp. 407–434, 1999.
- [32] S. Orlando and R. Perego. SUPPLE: an Efficient Run-Time Support for Non-Uniform Parallel Loops. *Journal of Systems Architecture*, Special Issue on New Trends in Programming and Execution Models for Parallel Architectures, Heterogeneously Distributed Systems and Mobile Computing, Elsevier Science publisher, Vol. 45, pp. 1323–1343, 1999.
- [33] S. Orlando and R. Perego. A Comparison of Implementation Strategies for Non-Uniform Data Parallel Computations. *Journal of Parallel and Distributed Computing*, Academic Press, vol. 52, n. 2, August 1998, pp. 132-149.
- [34] S. Orlando and R. Perego. Exploiting Partial Replication in Unbalanced Parallel Loops Scheduling on Multicomputers. *Microprocessing and Microprogramming - The EUROMICRO J.*, Elsevier Science publisher, 41, 1996.
- [35] B. Bacci, M. Danelutto, S. Orlando, S. Pelagatti, and M. Vanneschi. *P³L*: a Structured High-level Parallel Language and its Structured Support. *Concurrency: Practice and Experience*, Wiley, Vol. 7(3), pp. 225–255, 1995.
- [36] M. Danelutto, R. Di Meglio, S. Orlando, S. Pelagatti, and M. Vanneschi. A Methodology for the Development and the Support of Massively Parallel Programs. *Future Generation Computer Systems J.*, 8:205–220, 1992. Also published as selected paper in *Programming Languages for Parallel Processing*, IEEE Computer Society Press, D. B. Skillicorn and D. Talia Ed., 319–334, 1994.

International Conference Papers

- [37] S. Assefa Abebe, C. Lucchese and S. Orlando. EiFFFeL: Enforcing Fairness in Forests by Flipping Leaves. ACM SAC 2022, to appear.
- [38] F. Busolin, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, S. Trani: Learning Early Exit Strategies for Additive Ranking Ensembles. SIGIR 2021: 2217-2221
- [39] C. Lucchese, F.M. Nardini, R. Perego, S. Orlando, S. Trani. Query-level Early Exit for Additive Learning-to-Rank Ensembles. ACM SIGIR 2020, pp. 2033-2036.

- [40] C. Lucchese, F.M. Nardini, R. Perego, S. Orlando, S. Trani. Selective Gradient Boosting for Effective Learning to Rank. *SIGIR 2018*: 155-164.
- [41] M. Coletto, C. Lucchese, S. Orlando. Do Violent People Smile: Social Media Analysis of their Profile Pictures. *WWW (Companion Volume) 2018*: 1465-1468
- [42] F. Lettich, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonello, R. Venturini. Multicore/Manycore Parallel Traversal of Large Forests of Regression Trees. *HPCS 2017*: 915
- [43] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, S. Trani. X-DART: Blending Dropout and Pruning for Efficient Learning to Rank. *SIGIR 2017*: 1077-1080
- [44] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonello, R. Venturini. QuickScorer: Efficient Traversal of Large Ensembles of Decision Trees. *ECML/PKDD 2017 (Skopje)*, LNAI 10536, Springer.
- [45] Claudio Lucchese, Salvatore Orlando, Raffaele Perego. Evaluating Top-K Approximate Patterns via Text Clustering. *DaWaK 2016*: 114-127, LNCS 9829, Springer 2016.
- [46] Salvatore Trani, Diego Ceccarelli, Claudio Lucchese, Salvatore Orlando, Raffaele Perego. SEL: A Unified Algorithm for Entity Linking and Saliency Detection. *ACM Symposium on Document Engineering, DocEng 2016*: 85-94 (**Best Student Paper**).
- [47] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Gabriele Tolomei. Learning to Rank User Queries to Detect Search Tasks. *ACM on International Conference on the Theory of Information Retrieval, ICTIR 2016*: 157-166.
- [48] C.L. Rocha, I.R. Brilhante, F. Lettich, J.A.F. de Macêdo, A. Raffaetà, R.M. de Castro Andrade, S. Orlando. TPRED: a Spatio-Temporal Location Predictor Framework. *20th International Database Engineering & Applications Symposium, IDEAS 2016*, ACM, pp. 34-42.
- [49] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonello, R. Venturini. Exploiting CPU SIMD Extensions to Speed-up Document Scoring with Tree Ensembles. *The 39th ACM SIGIR 2016*: 833-836 (short).
- [50] M. Coletto, C. Lucchese, S. Orlando, R. Perego. Polarized User and Topic Tracking in Twitter. *The 39th ACM SIGIR 2016*: 945-948 (short).
- [51] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, F. Silvestri, S. Trani: Post-Learning Optimization of Tree Ensembles for Efficient Ranking. *The 39th ACM SIGIR 2016*: 949-952 (short).
- [52] M. Pourvali, S. Orlando, M. Gharagozloo. Improving Clustering Quality by Automatic Text Summarization. *11th Asia Information Retrieval Societies Conference, AIRS 2015*, LNCS 9460, Springer, pp. 292-303.
- [53] F. Lettich, S. Orlando, C. Silvestri. Processing streams of spatial k-NN queries and position updates on manycore GPUs. *ACM SIGSPATIAL/GIS 26:1-26:10 2015*.
- [54] Claudio Lucchese, Salvatore Orlando, Raffaele Perego. Supervised Evaluation of Top-k Itemset Mining Algorithms. *17th Int.l Conf. on Data Warehousing and Knowledge Discovery (DaWaK)*, 2015, LNCS 9263, Springer 2015, Springer, pp 82-94.
- [55] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Nicola Tonello, Rossano Venturini. QuickScorer: a Fast Algorithm to Rank Documents with Additive Ensembles of Regression Trees. *The 38th ACM SIGIR 2015*, pages 73-82, ACM, 2015 (**Best Paper**).
- [56] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Nicola Tonello. Speeding up Document Ranking with Rank-based Features. *The 38th Annual ACM SIGIR Conference (SIGIR 2015)*, Short paper, pp. 895-898.
- [57] Salvatore Trani, Diego Ceccarelli, Claudio Lucchese, Salvatore Orlando, Raffaele Perego. Manual Annotation of Semi-Structured Documents for Entity-Linking. *ACM CIKM 2014*: 2075-2077 (Demo paper).
- [58] Diego Ceccarelli, Claudio Lucchese, Salvatore Orlando, Raffaele Perego, Salvatore Trani. Dexter 2.0 - an Open Source Tool for Semantically Enriching Data. *International Semantic Web Conference (Posters & Demos) 2014*: 417-420.
- [59] Stefano Calzavara, Gabriele Tolomei, Michele Bugliesi, Salvatore Orlando. Quite a mess in my cookie jar!: leveraging machine learning to protect web authentication. *Proc. of WWW Conference 2014*: 189-200.
- [60] Claudio Silvestri, Francesco Lettich, Salvatore Orlando, and Christian S. Jensen. GPU-based Computing of Repeated Range Queries over Moving Objects. *22nd Euromicro Int.l Conf. on Parallel, Distributed and Network-Based Processing (PDP 2014)*, pages 640-647.
- [61] D. Ceccarelli, C. Lucchese, S. Orlando, R. Perego, and S. Trani. Learning relatedness measures for entity linking. *Proc of. ACM CIKM 2013*, pages 139-148.
- [62] D. Broccolo, C. Macdonald, S. Orlando, I. Ounis, R. Perego, F. Silvestri, and N. Tonello. Load-Sensitive Selective Pruning for Distributed Search. *Proc of. ACM CIKM 2013*, pages 379-388.
- [63] D. Broccolo, C. Macdonald, S. Orlando, I. Ounis, R. Perego, F. Silvestri, and N. Tonello. Query Processing in Highly-Loaded Search Engines". *Proc. of SPIRE 2013*, LNCS, Springer, ISBN 978-3-319-02431-8, pages 49-55.

- [64] F. Giummolè, S. Orlando, and G. Tolomei. Trending Topics on Twitter Improve the Prediction of Google Hot Queries. *Proc. of 2013 ASE/IEEE Socialcom 2013*, pp. 39-44, 2013.
- [65] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, and G. Tolomei. Modeling and Predicting the Task-by-Task Behavior of Search Engine Users. *Proc. of 10th Int.l Conference on Open Research Areas in Information Retrieval, OAIR 2013, RIAO series*. pages 77–84.
- [66] C. Silvestri, S. Orlando. Exploiting GPUs in Frequent Itemset Mining. *20th Euromicro Int.l Conf. on Parallel, Distributed and Network-Based Processing (PDP 2012)*. Garching, Germany, 15-17 Feb, 2012, p. 416-425, IEEE Computer Society, ISBN: 9780769546339.
- [67] A. Triossi, S. Orlando, A. Raffaetà, T. Fruehwirth. Compiling CHR to parallel hardware. *Proc. of the 14th Symposium on Principles and Practice of Declarative Programming*. p. 173-184, ACM, ISBN: 9781450315227, 2012.
- [68] D. Ceccarelli, C. Lucchese, S. Orlando, R. Perego, F. Silvestri. Caching query-biased snippets for efficient retrieval. *Proc. of 14th Int.l Conf. on Extending Database Technologies (EDBT 2011)*, Upsala, Sweden, March 22-24, 2011.
- [69] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, G. Tolomei. Detecting Task-based Query Sessions using Collaborative Knowledge. In *Proc. of 4th ACM WSDM*, pages 277–286, 2011.
- [70] C. Lucchese, S. Orlando, R. Perego. Mining Top-K Patterns from Binary Datasets in presence of Noise. In *Proc. of the 10th SIAM Int.l Conf. on Data Mining (SDM'10)*, April 2010.
- [71] C. Lucchese, S. Orlando, R. Perego. Generative Pattern Model for Mining Binary Datasets. In *Proc. of the 2010 ACM Symposium on Applied Computing (SAC '10)*, Track on Data Mining.
- [72] L. Leonardi, G. Marketos, E. Frentzos, N. Giatrakos, S. Orlando, N. Pelekis, A. Raffaetà, A. Roncato, C. Silvestri, Y. Theodoridis. T-Warehouse: Visual OLAP Analysis on Trajectory Data. *Proc. of ACM ICDE 2010, Demo paper*.
- [73] M. Mordacchini, P. Dazzi, G. Tolomei, R. Baraglia, S. Orlando, F. Silvestri. Challenges in Designing an Interest-based Distributed Aggregation of Users in P2P Systems. *Int.l Conf. on Ultra Modern Telecommunications 2009*.
- [74] S. Orlando, F. Silvestri. Mining Query Logs. *ECIR 2009, LNCS, Volume 5478/2009*, pages 814-817, Springer.
- [75] F. Falchi, C. Lucchese, S. Orlando, R. Perego, F. Rabitti. Caching Content-based Queries for Robust and Efficient Image Retrieval. *Proc. of 12th EDBT*, pages 780-790, 2009, ACM.
- [76] L. Leonardi, S. Orlando, A. Raffaetà, A. Roncato, C. Silvestri. Frequent Spatio-Temporal Patterns in Trajectory Data Warehouses. In *Proc. of the 2009 ACM Symposium on Applied Computing (SAC '09)*, Track on Data Mining, pages 1433-1440, 2009.
- [77] F. Silvestri, R. Baraglia, C. Lucchese, S. Orlando, R. Perego. (Query) History Teaches Everything, Including the Future". Paper invited to *LA-Web 2008*, pp. 12-22, IEEE CS press, October 28th-30th, , Brazil, 2008.
- [78] G. DaCosta, S. Orlando, M. D. Dikaiakos. Multi-set DHT for interval queries on dynamic data. *1st Int.l Conf. on Data Management in Grid and P2P Systems (Globe'2008)*, LNCS, Springer.
- [79] C. Gennaro, M. Mordacchini, S. Orlando, F. Rabitti. Processing Complex Similarity Queries in Peer-to-Peer Networks. *Proc. of The 23rd Annual ACM Symposium on Applied Computing (SAC '08)*, 2008, Track on "Engineering Large-Scale Distributed Systems"
- [80] F. Lelli, G. Maron, S. Orlando. Towards Response Time Estimation in Web Services. *IEEE International Conference on Web Services (ICWS 2007)*, pp. 1138-1139, IEEE, ISBN: 9780769529240, 2007.
- [81] F. Lelli, G. Maron, S. Orlando. Client Side Estimation of a Remote Service Execution. *15th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2007)*, October 24-26, 2007, Istanbul, Turkey.
- [82] G. Da Costa, M. D. Dikaiakos, S. Orlando. Nine months in the life of EGEE: a look from the South. *15th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2007)*, October 24-26, 2007, Istanbul, Turkey. pages 281-287.
- [83] C. Lucchese, S. Orlando, R. Perego. Parallel Mining of Frequent Closed Patterns: Harnessing Modern Computer Architectures. *Proc. of the IEEE Int.l Conf. on Data Mining 2007 (ICDM 07)*. pp 242–251, 2007.
- [84] S. Orlando, R. Orsini, A. Raffaetà, A. Roncato, C. Silvestri. Spatio-Temporal Aggregations in Trajectory Data Warehouses. *9th Int.l Conf. on Data Warehousing and Knowledge Discovery (DaWaK '07)*, LNCS 4654, Springer, 2007, pp. 66–77.
- [85] C. Lucchese, S. Orlando, R. Perego, F. Silvestri. Mining Query Logs to Optimize Index Partitioning in Parallel Web Search Engines. *Proc. of 2nd Int. Conf. on Scalable Information Systems (INFOSCALE'07)*, 2007.
- [86] R. Baraglia, C. Lucchese, S. Orlando, M. Serrano, F. Silvestri. A Privacy Preserving Web Recommender System. *Proc. of The 21th Annual ACM Symposium on Applied Computing (SAC '06)*, 2006, pp. 559-563.
- [87] F. Bonchi, F. Giannotti, C. Lucchese, S. Orlando, R. Perego, R. Trasarti. CONQUEST: a Constraint-based Querying System for Exploratory Pattern Discovery. *Proc. of ACM ICDE 2006, Demo paper*, pp. 159.

- [88] M. Marzolla, M. Mordacchini, S. Orlando. Tree Vector Indexes: Efficient Range Queries for Dynamic Content on Peer-to-Peer Networks. Proc. of the 14th Euromicro Conference on Parallel, Distributed and Network based Processing (PDP 2006), IEEE Press, pp. 457–464, 2006.
- [89] C. Lucchese, S. Orlando, R. Perego. Mining frequent closed itemsets out-of-core. Proc. of the SIAM Conference on Data Mining (SDM '06), pp. 417–427, 2006.
- [90] G. Maron, F. Lelli, S. Orlando. Improving the Performance of XML Based Technologies by Caching and Reusing Information. IEEE Int.l Conference on Web Services (ICWS 2006), pp. 689–700, 2006.
- [91] C. Silvestri, S. Orlando. Distributed Approximate Mining of Frequent Patterns. Proc. of the 2005 ACM Symposium on Applied Computing, SAC 2005, special track on Data Mining.
- [92] D. Puppini, F. Silvestri, D. Laforenza, S. Orlando. A Search Engine for Software Components. Proc. of the 2004 IEEE/WIC/ACM Int. Conf. on Web Intelligence (WI 2004). Beijing, China. September 20-24, 2004. pp. 495-498.
- [93] F. Silvestri, S. Orlando, R. Perego. Assigning Identifiers to Documents to Enhance the Clustering Property of Fulltext Indexes. Proc. of the 27th SIGIR Conference, pp. 305–312, 2004.
- [94] F. Silvestri, S. Orlando, R. Perego. WINGS: a Parallel Indexer for Web Contents. Proc. of International Conference on Computational Science 2004 (ICCS'04) - Krakov, Poland. LNCS, Springer, to appear.
- [95] P. Palmerini, S. Orlando, R. Perego. Statistical Properties of Transactional Databases. Proceedings of the 2004 ACM Symposium on Applied Computing, SAC 2004, special track on Data Mining, pp. 515–519, March 14-17, 2004, Nicosia, Cyprus.
- [96] S. Orlando, C. Silvestri, R. Perego. A new algorithm for gap constrained sequence mining. Proceedings of the 2004 ACM Symposium on Applied Computing, SAC 2004, special track on Data Mining, pp. 540–547, March 14-17, 2004, Nicosia, Cyprus. ISBN:1-58113-812-1.
- [97] F. Silvestri, S. Orlando, R. Perego. Assigning Document Identifiers to Enhance Compressibility of Web Indexes. Proceedings of the 2004 ACM Symposium on Applied Computing, SAC 2004, special track on Data Mining, pp. 600-605, March 14-17, 2004, Nicosia, Cyprus.
- [98] T. Fagni, S. Orlando, P. Palmerini, R. Perego, F. Silvestri. A Hybrid Strategy for Caching Web Search Engine Results. Poster paper Twelfth International World Wide Web Conference (WWW2003), May 20-24, 2003, Budapest, Hungary.
- [99] R. Baraglia, M. Danelutto, D. Laforenza, S. Orlando, P. Palmerini, P. Pesciullesi, R. Perego, M. Vanneschi. AssistConf: a Grid configuration tool for the Assist parallel programming environment. Proc. of the 11th Euromicro Conference on Parallel Distributed and Network based Processing (PDP2003) Genoa, Italy, February 5-7, 2003. IEEE CS Press, pp. 193-200.
- [100] S. Orlando, P. Palmerini, R. Perego, F. Silvestri. Adaptive and Resource-Aware Mining of Frequent Sets. Proc. of IEEE ICDM, 2002. IEEE CS Press, pp. 338–345.
- [101] S. Orlando, P. Palmerini, R. Perego, F. Silvestri. Scheduling High Performance Data Mining Tasks on a Data Grid Environment. Proc. of Euro-Par, 27-30 August 2002. LNCS 2400 - Springer, pp. 375-384.
- [102] S. Orlando, P. Palmerini, R. Perego, F. Silvestri. An Efficient Parallel and Distributed Algorithm for Counting Frequent Sets. Proc. of Int. Conf. VECPAR 2002, June 2002, Porto, Portugal. Appeared as Selected Paper in the LNCS 2565, Springer, pp. 197-204.
- [103] S. Orlando, R. Perego, F. Silvestri. Design of a Parallel and Distributed Web Search Engine. 2001 Parallel Computing Conference (ParCo 2001), 4-7 September 2001. Imperial College Press, pp. 197–204.
- [104] S. Orlando, P. Palmerini, R. Perego. Enhancing the Apriori Algorithm for Frequent Set Counting. Proceedings of DaWak 2001, September 5-7, 2001, Munich, Germany, Lecture Notes In Computer Science 2114, Springer, 2001, pp. 71–82.
- [105] S. Ciarpaglini, L. Folchi, S. Orlando, S. Pelagatti, and R. Perego. Integrating Task and Data Parallelism with taskHPF. In *Proc. of PDPTA'2000 - 2000 International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, Nevada, USA, June 26-29, 2000, pp. 2485–2491.
- [106] S. Orlando, P. Palmerini, and R. Perego. Coordinating HPF programs to mix task and data parallelism. In *Proc. of the 2000 ACM Symposium on Applied Computing (SAC'2000) - Special Track on Coordination Models, Languages and Application*, Villa Olmo, Como, Italy, March 19-21, 2000, ACM Press, pp. 240–247.
- [107] S. Orlando and R. Perego. Scheduling Data-Parallel Computations on Heterogeneous and Time-Shared Environments. In *Proc. of EUROPAR'98*, Southampton, UK, September 1998, LNCS 1470 - Springer Verlag - Pag. 356-366.
- [108] S. Orlando and R. Perego. An MPI-Based Run-Time Support to Coordinate HPF Tasks. In *Proc. of the EuroPVM/MPI98 Conference*, Liverpool, UK, September 1998, LNCS 1497 - Springer Verlag - Pag 289–296.
- [109] S. Orlando and R. Perego. A Support for Non-Uniform Parallel Loops and its Application to a Flame Simulation Code. In *Proc. of Irregular '97 - Fourth Int. Symposium on Solving Irregularly Structured Problems in Parallel*, LNCS 1253, Springer Verlag, pages 186–197, June 1997. Paderborn (Germany).

- [110] S. Orlando and R. Perego. A template for non-uniform parallel loops based on dynamic scheduling and prefetching techniques. In *Proc. of the 1996 ACM Int. Conf. on Supercomputing*, May 1996. Philadelphia-PA.
- [111] S. Orlando and R. Perego. A Load Balancing Strategy for Iterated Parallel Loop Scheduling. In *Proc. of the HPCN Europe 96 – Int. Conf. and Exhibition*, pages 978–979. LNCS 1067 - Springer Verlag, April 1996. Bruxelles.
- [112] S. Orlando and R. Perego. An Efficient Template for the Highly Parallel Implementation of Non-Uniform Loops. In *Proc of the 2nd EUROPEAN SCHOOL OF COMPUTER SCIENCE – Parallel Progr. Environ. for High Perf. Computing*, pages 31–34, April 1996. Alpe d’Huez - France.
- [113] B. Bacci, M. Danelutto, S. Pelagatti, S. Orlando, and M. Vanneschi. Summarizing an experiment in parallel programming language design. In *Proc. of the HPCN Europe 95 – Int. Conf. and Exhibition*, pages 7–13, LNCS 919 - Springer Verlag, May 1995.
- [114] S. Orlando and R. Perego. Scheduling Non-Uniform Parallel Loops on Highly Parallel Multicomputers. In *Proc. of the IEEE Int. Conf. on High Performance Computing – HiPC95*, pages 547–553. MacGraw Hill, December 1995. New Delhi, India.
- [115] B. Bacci, M. Danelutto, S. Pelagatti, S. Orlando, and M. Vanneschi. Unbalanced Computations onto a Transputer Grid. In *Proceedings of The 1994 Transputer Research and Application Conference*, pages 268–282, October 1994. Athens, Georgia, USA.
- [116] S. Orlando and M. Vanneschi. Compiling Geometric Paradigms through Local Communications. In *Proc. of IFIP Int. Conf. on Applications in Parallel and Distributed Computing*, pages 175–184, Caracas, Venezuela, April 1994.
- [117] B. Bacci, M. Danelutto, S. Orlando, S. Pelagatti, and M. Vanneschi. Efficient Compilation of Structured Parallel Programs for Distributed Memory MIMD Machines. In *Proc. of Int. Conf. ParCo93 - Parallel Computing 93*, pages 565–568, 1994.
- [118] S. Orlando and M. Vanneschi. A Synchronous Model for Static Optimization of Multicomputer Programs. In *Proc. of the IFIP Conf. on Progr. Env. for Parallel Computing*, pages 193–214, Edinburgh, Scotland, April 1992. IFIP Transactions Series, North-Holland.
- [119] F. Baiardi and S. Orlando. Strategies for a Massively Parallel Implementation of Simulated Annealing. In *Proc. of Int. Conf. PARLE ’89*, pages 273–287, Eindhoven, The Netherlands, June 1989. LNCS 366 Spinger-Verlag.

Chapters in International Research Books

- [120] F.M. Nardini, S. Orlando, R. Perego, A. Raffaetà, C. Renso, C. Silvestri. Analysing Trajectories of Mobile Users: From Data Warehouses to Recommender Systems. In: Flesca, Greco, Masciari, Saccà (eds). "A Comprehensive Guide Through the Italian Database Research Over the Last 25 Years". Studies in Big Data, vol 31. Springer, 2018.
- [121] R. Baraglia, C. Lucchese, S. Orlando, R. Perego, F. Silvestri. Preserving Privacy in Web Recommender Systems. In *Privacy Aware Knowledge Discovery: Novel Applications and New techniques*, F. Bonchi and E. Ferrari eds., pages 369–391, Taylor and Francis publisher, 2010.
- [122] S. Orlando, A. Raffaetà, A. Roncato, C. Silvestri. Warehousing and Mining Streams of Mobile Object Observations In *"Intelligent Techniques for Warehousing and Mining Sensor Network Data"*, Alfredo Cuzzocrea eds., IGI Global publication, ISBN: 978-1-60566-328-9, 2010.
- [123] D. Barbalace, C. Lucchese, C. Mastroianni, S. Orlando, D. Talia. Mining@home: Public Resource Computing for Distributed Data Mining. In *From Grids to Service and Pervasive Computing*, 2008, pp. 217-227, SPRINGER, ISBN: 9780387094540
- [124] F. Bonchi, F. Giannotti, C. Lucchese, S. Orlando, R. Perego, R. Trasarti. On Interactive Pattern Mining from Relational Databases. *Knowledge Discovery in Inductive Databases*, Revised Selected and Invited Papers. LNCS 4747, pp. 42–62, 2007, Springer.
- [125] D. Laforenza, C. Lucchese, S. Orlando, R. Perego, D. Puppini, F. Silvestri. On the Value of Query Logs for Modern Information Retrieval. In A. Soro, G. Paddeu, G. Armano. *Distributed Agent-based Retrieval Tools*. Polimetrica International Scientific Publisher, 2006.
- [126] C. Lucchese, S. Orlando, R. Perego, C. Silvestri. Mining Frequent Closed Itemsets from Distributed repositories. In D. Talia, A. Bilas, M. Dikaiakos. *Knowledge and Data Management in Grids*. vol. 3 of CoreGrid Series, Spinger, 2007.
- [127] R. Baraglia, M. Danelutto, T. Fagni, D. Laforenza, S. Orlando, A. Paccosi, N. Tonello, M. Vanneschi, C. Zoccolo. HPC Application Execution on Grids. Getov, Laforenza, Reinefeld eds., *Future Generation Grids*, CoreGRID Series, Vol 3, pp. 263-282, Springer, 2006. ISBN 0-387-27935-0.
- [128] D. Puppini, F. Silvestri, D. Laforenza, S. Orlando. Toward GRIDLE: A Way to Build Grid Applications Searching Through an Ecosystem of Components. In Jose C. Cunha and Omer F. Rana eds., *Grid Computing: Software Environments and Tools*, Springer Verlag, 2006, ISBN: 1-85233-998-52004.

- [129] B. Bacci, B. Cantalupo, M. Danelutto, S. Orlando, D. Pasetto, S. Pelagatti, M. Vanneschi. An environment for structured parallel programming. In Grandinetti, Kowalick eds., *Advances in High Performance Computing*, pages 219–234, Kluwer Academic Publishers, 1997.
- [130] S. Orlando. Compile-time Parallelizing Techniques. In G. Balbo and M. Vanneschi, editors, *General Purpose Parallel Computers: Architectures, Programming Environments, and Tools*, chapter 8, pages 209–250. ETS Pisa, 1995.

International Workshop Papers

- [131] C. Silvestri, F. Cagnin, F. Lettich, S. Orlando, M. Wachowicz. Mining condensed spatial co-location patterns. ACM SIGSPATIAL International Workshop on Mobile Geographic Information Systems (MobiGIS) 2015: 84-87.
- [132] D. Ceccarelli, C. Lucchese, S. Orlando, R. Perego, and S. Trani. Dexter: an Open Source Framework for Entity Linking. Proc. of the CIKM 2013 Workshop on Exploiting Semantic Annotation (ESAIR), 2013, pages 17-20.
- [133] G. Tolomei, S. Orlando, D. Ceccarelli, C. Lucchese. Twitter Anticipates Bursts of Requests for Wikipedia Articles. Proc. of the CIKM 2013 Workshop on Data-driven User Behavioral Modelling and Mining from Social Media (DUBMOD@CIKM), 2013, pages 5-8.
- [134] S. Orlando, F. Pizzolon, and G. Tolomei. SEED: A Framework for Extracting Social Events from Press News. Proc. of the 2nd Int.l Workshop on Web of Linked Entities (WOLE 2013), WWW2013 Workshop. WWW (Companion Volume), pages 1285-1294.
- [135] G. Tolomei, S. Orlando, F. Silvestri. Towards a task-based search and recommender systems. ICDE Workshops, 2010, pages 333–336
- [136] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, G. Tolomei. Detecting Task-Based Query Sessions Using Collaborative Knowledge. Web Intelligence/IAT Workshops 2010, pages 128-131.
- [137] A. Triossi, S. Orlando, A. Raffaetà, F. Raiser, T. Frühwirth. Constraint-based hardware synthesis. Proc. of 24th Workshop on (Constraint) Logic Programming, 2010.
- [138] F. Falchi, C. Lucchese, S. Orlando, R. Perego, F. Rabitt. A metric cache for similarity search. Proc. of the 2008 ACM CIKM workshop on Large-Scale distributed systems for information retrieval (LSDS-IR) 2008, pages 43-50, ACM.
- [139] C. Gennaro, M. Mordacchini, S. Orlando, F. Rabitti. MRout: A Peer-to-Peer Routing Index for Similarity Search in Metric Spaces. Proc. of the 5th Int. Workshop on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P 2007) in conjunction with VLDB 2007. 2007.
- [140] F. Braz, S. Orlando, R. Orsini, A. Raffaetà, A. Roncato, C. Silvestri. Approximate Aggregations in Trajectory Data Warehouses. Proc. ICDE Workshop on Spatio-Temporal Data Mining 2007 (STD07), IEEE Press.
- [141] M. Marzolla, M. Mordacchini, S. Orlando. A P2P Resource Discovery System Based on a Forest of Trees. 17th Int. Workshop on Database and Expert Systems Applications (DEXA'06) - GLOBE06, pp. 261-265, IEEE press.
- [142] E. Frizziero, M. Gulmini, F. Lelli, G. Maron, A. Oh, S. Orlando, A. Petrucci, S. Squizzato, S. Traldi. Instrument Element: A New Grid component that Enables the Control of Remote Instrumentation. 6th IEEE Int. Symposium on Cluster Computing and the Grid Workshops (CCGRIDW'06) - Workshop on Grid Testbeds, 2006.
- [143] C. Lucchese, S. Orlando, R. Perego. Distributed Mining of Frequent Closed Itemsets: Some Preliminary Results. Proc. of the 8th Int. Workshop on High Performance and Distributed Mining (HPDM '05), in conjunction with SIAM Data Mining Conf. (SDM'05). April, 2005.
- [144] M. Marzolla, M. Mordacchini, S. Orlando. Resource Discovery in a Dynamic Grid Environment. 16th Int. Workshop on Database and Expert Systems Applications (DEXA'05) - GLOBE05 - pp. 356-360, IEEE press.
- [145] C. Silvestri, S. Orlando. Approximate Mining of Frequent Patterns on Streams. 2nd Int. ECML-PKDD Workshop on Knowledge Discovery from Data Streams, 2005.
- [146] C. Lucchese, S. Orlando, R. Perego. DCI_{Closed}: A Fast and Memory Efficient Algorithm to Mine Frequent Closed Itemsets. Proceedings of the Workshop on Frequent Itemset Mining Implementations (FIMI'04), in conjunction with ICDM'04, November 1-4, 2004. Brighton, UK.
- [147] C. Silvestri, S. Orlando. Distributed association mining: an approximate method. Proceedings of the 7th Int. Workshop on High Performance Distributed Mining (HPDM'04), in conjunction with 4th International SIAM Conference on Data Mining, Orlando, Florida.
- [148] C. Lucchese, S. Orlando, P. Palmerini, R. Perego, F. Silvestri. kDCI: a Multi-Strategy Algorithm for Mining Frequent Sets. Proceedings of the Workshop on Frequent Itemset Mining Implementations (FIMI'03), in conjunction with ICDM'03. 19 November 2003, Melbourne, Florida, USA.
- [149] S. Orlando, R. Perego, C. Silvestri. CCSM: an Efficient Algorithm for Constrained Sequence Mining. Proc. of the 6th Int.l Workshop on High Performance Data Mining: Pervasive and Data Stream Mining, in conjunction with Third International SIAM Conference on Data Mining (SDM'03), May 1-3, 2003, San Francisco, CA.

- [150] S. Orlando, P. Palmerini, R. Perego, F. Silvestri. A Scalable Multi-Strategy Algorithm for Counting Frequent Sets. Proc. of the 5th Workshop on High Performance Data Mining, April 2002, Washington, USA.
- [151] C. Gennaro, S. Orlando, and R. Perego. Integrating HPF in a Skeleton Based Parallel Language. Proceedings of the 9th Euromicro Workshop on Parallel and Distributed Processing, PDP2001, Mantova, Italy February 7th-9th, 2001, IEEE CS Press, pp. 195–202.
- [152] R. Baraglia, D. Laforenza, S. Orlando, P. Palmerini, and R. Perego. Implementation Issues in the Design of I/O Intensive Data Mining Applications on Clusters of Workstations. In Proc. of IPDPS-2000 - Workshop on High Perf. Data Mining, Cancun, Mexico, 2000, LNCS 1800 - Springer Verlag, pp. 350-357.
- [153] S. Orlando and R. Perego. A Coordination Layer for Exploiting Task Parallelism with HPF. In Proc. of the Fourth ACM Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR98), Carnegie Mellon University, Pittsburgh, PA, USA, May 1998, LNCS 1511 - Springer Verlag - Pag. 386–393.
- [154] F. Baiardi and S. Orlando. Synchronous Execution Models for Massively Parallel Systems. In Proc. of Int. Workshop on Supercomputing Tools for Science and Engineering, pages 177–185, Pisa, Italy, December 1989. Franco Angeli.

Edited Books, Journals, Book Chapters

- [155] Volker Markl, Salvatore Orlando, Bernhard Mitschang, Periklis Andritsos, Kai-Uwe Sattler, Sebastian Breß. Proc. of the 20th International Conference on Extending Database Technology, EDBT 2017, Venice, Italy, March 21-24, 2017. OpenProceedings.org 2017, ISBN 978-3-89318-073-8.
- [156] G.M. Di Nunzio, F.M. Nardini, S. Orlando (Eds.). Proceedings of the 7th Italian Information Retrieval Workshop (IIR 2016), Venezia, 2016. CEUR Workshop Proceedings 1653, CEUR-WS.org 2016.
- [157] S. Orlando, G. Antoniu, A. Ghoting, M. S. Pérez-Hernández (Eds.). Introduction to Topic 5: Parallel and Distributed Data Management. Euro-Par 2011: pp. 351-352, LNCS 6852/2011, Springer.
- [158] R. Sakellariou, S. Orlando, J.-L. Larriba-Pey, S. Parthasarathy, D. Zeinalipour-Yazti (Eds.). *Parallel and Distributed Data Management*. Euro-Par 2010, pp. 316, LNCS 6271/2010, Springer.
- [159] S. Orlando, D. Laforenza (Eds.). Selected Papers from the EUROPVM/MPI 2003 Conference, Venice, Italy, 29 September-2 October 2003. *Int. Journal of High Performance Computing Applications*, Spring 2005, Volume 19, No. 1, Sage Publications.
- [160] D. B. Skillicorn, A. Hameurlain, P. Watson, S. Orlando (Eds.). *Topic 5: Parallel and Distributed Databases, Data Mining and Knowledge Discovery*. Euro-Par 2004, pp. 346, LNCS 3149, Springer.
- [161] J. Dongarra, D. Laforenza, S. Orlando (Eds.). *Recent Advances in Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) - 10th European PVM/MPI Users' Group Meeting*, Venice, Italy, September 29 - October 2, 2003, Proceedings. Lecture Notes in Computer Science, 2840, Springer.

National Conference/Workshop Papers

- [162] F. Lettich, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonellotto, R. Venturini. GPU-based Parallelization of QuickScorer to Speed-up Document Ranking with Tree Ensembles. IIR 2016, CEUR Workshop Proceedings 1653, CEUR-WS.org 2016.
- [163] C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, F. Silvestri, S. Trani. Post-Learning Optimization of Tree Ensembles. IIR 2016, CEUR Workshop Proceedings 1653, CEUR-WS.org 2016.
[c108] viewexportask othersask others Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Nicola Tonellotto, Rossano Venturini: Speeding-up Document Scoring with Tree Ensembles using CPU SIMD Extensions. IIR 2016
- [164] D. Dato, C. Lucchese, F.M. Nardini, S. Orlando, R. Perego, N. Tonellotto, R. Venturini. Ranking Documents Efficiently with QuickScorer. SEBD 2016, pp. 166-173, Matematicamente.it 2016, ISBN 9788896354889.
- [165] C. Silvestri, F. Lettich, S. Orlando, C.S. Jensen. A wait-free output data structure for GPU-based streaming query processing. 23rd Italian Symposium on Advanced Database Systems (SEBD 2015): 232-239.
- [166] G. Capannini, D. Dato, C. Lucchese, M. Mori, F.M. Nardini, S. Orlando, R. Perego, N. Tonellotto. QuickRank: a C++ Suite of Learning to Rank Algorithms. Proc. of the 6th Italian Information Retrieval Workshop (IIR2015). CEUR Workshop Proceedings 1404, CEUR-WS.org 2015.
- [167] M. Coletto, C. Lucchese, S. Orlando, R. Perego. Electoral Predictions with Twitter: a Machine-Learning approach Proc. of the 6th Italian Information Retrieval Workshop (IIR2015). CEUR Workshop Proceedings 1404, CEUR-WS.org 2015.

- [168] L. Leonardi, S. Orlando, A. Raffaetà, A. Roncato, C. Silvestri. Visual OLAP analysis on network constrained trajectories. Proc. of the 18th Italian Symposium on Advanced Database Systems (SEBD 2010), Rimini, Italy, June 20-23.
- [169] F. Falchi, C. Lucchese, S. Orlando, R. Perego, Fausto Rabitti. Caching Algorithms for Similarity Search. Proc. of the 17th Italian Symposium on Advanced Database Systems (SEBD 2009).
- [170] C. Gennaro, M. Mordacchini, S. Orlando, F. Rabitti. A Scalable Distributed Data Structure for Multi-Feature Similarity Search. Proc. of the 16th Italian Symposium on Advanced Database Systems (SEBD 2008), Mondello (PA), Italy, June 22-25.
- [171] S. Orlando, A. Raffaetà, A. Roncato, C. Silvestri. Trajectory Data Warehouses: Storing and Aggregating Frequent ST Patterns. Proc. of the 16th Italian Symposium on Advanced Database Systems (SEBD 2008), Mondello (PA), Italy, June 22-25.
- [172] F. Braz, S. Orlando, R. Orsini, A. Raffaetà, A. Roncato, C. Silvestri. Designing a Spatio-Temporal Trajectory Data Warehouse. In Proc. of XXII Brazilian Symposium on Databases, 2007.
- [173] S. Orlando, R. Orsini, A. Raffaetà, A. Roncato, C. Silvestri. Trajectory Data Warehouses: Design Issues and Use Cases. Proc. of the 15th Italian Symposium on Advanced Database Systems (SEBD 2007), pp. 208–219, June 17-20, 2007, Torre Canne - Fasano (BR)
- [174] S. Orlando, R. Perego, F. Silvestri. Assigning Document Identifiers to Enhance Compressibility of Fulltext Indices. Proc. of the 12th Italian Symposium on Advanced Database Systems (SEBD 2004), pp. 222–229.
- [175] P. Criscione, R. Scopigno, and S. Orlando. Metodologia di programmazione parallela P^3L : un'applicazione a problemi di geometria computazionale. In *Proceedings of the Annual AICA Conference*, pages 221–229, September 1995. Chia, Italy.
- [176] B. Bacci, M. Danelutto, S. Pelagatti, S. Orlando, and M. Vanneschi. Methodologies and tools for massively parallel programming: the P^3L approach. In *Proceedings of the Annual AICA Conference*, pages 471–484, September 1994. Palermo, Italy.
- [177] R. Bisiani and S. Orlando. Supporting Parallel Programming with Distributed Artificial Intelligence Techniques. In *Proceedings of the Annual AICA Conference*, pages 441–448, September 1994. Palermo, Italy.
- [178] S. Antonelli, M. Danelutto, R. Di Meglio, S. Orlando, and S. Pelagatti. Una metodologia di programmazione per macchine a parallelismo massiccio - applicazione al problema del ray tracing. In *Atti del Workshop su Elaborazione Parallela: Ricerca ed Applicazioni*, pages 139–156, Roma, Italy, 1992.
- [179] S. Orlando and R. Perego. Un Modello per il Controllo di Transazioni Concorrenti in Ambiente Distribuito. In *Atti del Congresso Annuale AICA*, pages 71–84, Trieste, Italy, October 1989.

Other Works

- [180] R. Baraglia, D. Laforenza, S. Orlando, P. Palmerini, R. Perego. Design of Efficient Input/Output Intensive Data Mining Applications. ERCIM News, no. 44, 2000, a publication of the European Community in Information Technology.
- [181] R. Perego, S. Orlando. Task Parallelism in an HPF Framework. ERCIM News, no. 34, 1998, a publication of the European Community in Information Technology.
- [182] S. Orlando. Combining Template-based Task Parallelism and HPF Data Parallelism. Schloss Dagstuhl – Invited Seminar on “High Level Parallel Programming: Applicability, Analysis and Performance”, April 1999.
- [183] S. Orlando. Load Balancing techniques for Non-uniform Data-Parallel Computations. Schloss Dagstuhl – Invited Seminar on “Theory and Practice of Higher-Order Parallel Programming”, February 1997.
- [184] S. Orlando. La sicurezza nelle transazioni elettroniche. In L. Pilotti ed., *La comunicazione in rete per le PMI*, chapter 6, pages 183–206. Il Sole 24 Ore Libri, December, 1996.
- [185] S. Orlando, R. Perego. A Template for the Implementation of Non-Uniform Loops on Massively Parallel Machines. In *Science and Supercomputing at CINECA, 1995 Report*, G. Erbacci and M. Voli (Eds.), 1996, pages 192–199.