Daniela Tulone

Engineering, R&D Department
via S.Martino della, 56
00185 Rome, Italy

Phone: +39 06 492591
Fax: +39 06 4453278
E-mail: daniela.tulone@eng.it

RESEARCH INTERESTS: fault-tolerance in wired and wireless sensor networks with particular focus on enhancing the system performance and scalability via *trade-offs* between data accuracy and efficiency, via *statistical* techniques and via *quorum systems.* Consistency of shared data, coordination protocols, and secure distributed systems.
Recent work in sensor networks includes data collection and query answering, fault detection and diagnosis, time synchronization, data consistency in highly mobile networks such as vehicular networks, resource management such as energy conservation, sensing and interpreting images.

EDUCATION

- **Ph.D., Computer Science.** January 2003–December 2005.
  University of Pisa, Italy, and Massachusetts Institute of Technology, USA.
  Thesis: *Mechanisms for energy conservation in wireless sensor networks.*
- **M.S., Computer Science.** September 1996–January 1998.
  Courant Institute of Mathematical Sciences, New York University. GPA 3.9/4.0.
- **Laurea, Mathematics.** November 1989–March 1994.
  University of Catania, Italy. Graduated with the highest honors (110/110 *cum laude*).
  Thesis: *ETNA, Extensible Theorem Prover in NAtural deduction, in Set Theory.*

ADDITIONAL INTENSIVE POST-GRADUATE SCHOOLS

- **Advanced School on Cryptography.** Barcelona, Spain. February 2004.
  Department of Mathematics, Universitat Politecnica de Catalunya.
- **Advanced School on Mobile Computing.** Pisa, Italy. September 2003.
  Scuola Normale Superiore, Pisa.
- **International Post-Graduate School.** Cortona, Italy. July 1995.
  Scuola Matematica Interuniversitaria. Title: *Computational Complexity.*
  Entrance limited to 20 European Ph.D. students and young researchers.
- **Advanced School on Industrial Modeling.** March 1994–November 1994.
  University of Catania (Italy), University of Kaiserslautern (Germany), and IBM.
  Entrance limited to 15 post-graduate students, determined by a competitive exam. Topics covered: numerical techniques, splines, and optimization models. Final exam score 60/60.

WORK EXPERIENCE

- **Engineering Ingegneria Informatica, R&D Division.**
  Senior Researcher, Rome, Italy. November 2007–present.
  Responsible for the formation of a new research lab on wireless sensor networks. Currently involved in the definition of different European research projects and applicative projects. Application domains of interests are environmental monitoring, e-health, and automotive.
- **UCLA, Center for Embedded Networking Sensor Center (CENS).**
  Visiting Researcher, Los Angeles, USA. October 2006–August 2007.
- **MIT, Computer Science and Artificial Intelligence Laboratory.**
  Post-Doc Researcher, Cambridge, USA. January 2006–July 2006.
- **C.N.R. (Italian National Research Council), Dependable Computing Research Lab.**
  Research Associate, Pisa, Italy. November 2002–September 2003.
- **Bell–Laboratories, Secure System Research Department.**
  Staff Member, Murray Hill, NJ, USA. April 2000–October 2002.

- **AT&T Labs, WorldNet Internet Department.**
  Software Engineer, Holmel, NJ, USA. March 1998–January 2000.
  Lead developer for 3-tier help-desk and web-based systems.
- **Courant Institute of Mathematical Sciences.**
  C.N.R. Exchange Researcher, New York, USA. September 1996–September 1997.
- **Department of Mathematics,** University of Catania.
  C.N.R. Associate Researcher, Catania, Italy. February 1995–June 1996.

HONORS AND AWARDS

- **Best paper award.** IEEE International Conference of Communication (ICC), *Network Security and Information Assurance Symposium.*
  D. Tulone. *A secure and scalable digital time-stamping service*, June 2006.
- **Ph.D. grant from University of Pisa.** November 2002.
  Awarded (previous written and oral competition) one of the 6 Ph.D. grants from the University of Pisa for the academic years 2003–2005.
- Selected for the **MIT-Italy program** for the years 2004–2005.
- **C.N.R. (Italian National Research Council) research grant.** October 2002.
  Awarded a research grant to perform research at the ISTI CNR, Pisa.
- **C.N.R. research fellowship.** October 1995.
  Awarded **one of eight** fellowships from C.N.R. to perform research at the Courant Institute of Mathematical Sciences. This fellowship was determined through **national competition**.
- **C.N.R. research fellowship.** December 1994.
  Awarded **one of ten** fellowships from C.N.R. to conduct research on Logic and Computer Algebra at the Department of Mathematics, University of Catania. This fellowship was determined through **national competition**.
- **Consorzio Archimede grant.** March 1994.
  Awarded (previous written and oral competition) one of the 15 grants from Consorzio Archimede to work on Mathematical models.

PROJECTS

Monitoring in Sensor Networks.

- **Answering queries and detecting anomalies via time series forecasting.**
  with Samuel Madden (MIT)

  A crucial problem in long-lived sensor applications consists in designing on-line data collection and querying mechanisms that are able to analyze large volume of real-time sensor data and detect anomalies while conserving energy. We proposed an adaptable energy-efficient framework, called SAF, for approximately answering queries at the sink and detecting anomalies and data similarities, based on time series forecasting [4, 5, 6, 8, 22]. SAF relies on a class of lightweight and adaptable time series models built at sensors, and on a suite of novel algorithms for monitoring and adapting the prediction models, and for detecting *data anomalies*, *node similarities* and *data correlation* at no additional communication cost. These properties make SAF a flexible tool for monitoring physical phenomena and studying their evolution. SAF can be applied for instance to monitor the quality of water, air pollution, and to study climate variations, volcanic and seismic phenomena. It has applications also to agriculture and to e-health applications (e.g., to remotely monitor patients at risk or to assist doctors and nurses in case of emergency).
- **Fault detection and diagnosis, and robust data aggregation.**
  with Mani Srivastava (UCLA)

Embedded networked sensors are a powerful instrument with the potential of advancing knowledge by monitoring the physical world at unprecedented scales and resolutions. However, for an instrument to be useful, the information it provides must be of high integrity. This requirement is hard to be guaranteed in case of low-cost sensor networks since sensors are error-prone and sensitive to small environmental variations. As a result, the design of an *on-line detection and diagnosis* system is of paramount importance to provide guarantees on the data produced at sensors and to repair faults when possible. We proposed a general on-line framework, called *Inspect*, for detecting and diagnosing sensor faults. *Inspect* is able to distinguish among sensor malfunctioning, mis-calibration, significant variations in the phenomenon under analysis, and transient variations that are irrelevant to the application [11]. It can work in synergy with SAF, thus providing an energy-efficient real-time monitoring system which provides user-tunable confidence bounds.

– **A fire detection and rescue system.**

There is a growing concern for the increasing trend in wood fires for their tremendous costs, vegetation destruction, and in some case fatalities. One of the difficulties in performing efficient rescue operations is related to the latency in triggering fire alarms and to the lack of information related to the geographical regions affected by the fire and their possible correlations. We are currently applying our framework *SAF* and *Inspect* to design a fire detection and rescue system that relies on (1) a static sensor network dispersed in the wood to monitor environmental attributes that are relevant in detecting a fire and evaluating its impact, (2) body sensors worn by firefighters, and (3) a centralized basestation. SAF and Inspect play a crucial role in detecting and localizing a fire while keeping false alarms low, and in detecting geographical-temporal correlations crucial when designing efficient rescue strategies in the presence of a malicious coordinated attack. The rescue operations coordinated at the basestation are adaptive and rely on lightweight prediction models built at static sensors and at mobile body sensors.

– **A robust in-network data aggregation in sensor networks.**

Another key problem arising in very large sensor networks consists in designing robust in-network data aggregation mechanisms that are able to perform in-situ pre-processing and analyze local sensor data thus reducing the amount of communication and enhancing the network bandwidth usage. The techniques developed in SAF has been applied to design a lightweight data aggregation that can be carried out at sensors by detecting similarities and fusing sensor data produced at similar nodes [5].

Building blocks in sensor applications.

– **High-precision and energy-efficient time estimation.**

Time synchronization plays a crucial role in real-time monitoring applications, in node coordination and in several network tasks. Moreover, some sensor applications such as *seismic applications* and *structural health* monitoring systems require a clock accuracy of a few microseconds, which poses additional challenges for those applications. We studied the time synchronization problem from a novel perspective, which is complementary to the well-studied clock synchronization problem and consists of reducing the error growth between consecutive synchronization using clock information [1, 9, 13]. We proposed a suite of deterministic and probabilistic protocols that are both of theoretical and practical interest. In fact, they lead to a refinement of the *optimality bound* for external clock synchronization, and they can be applied to: (1) conserve energy in long-term applications, (2) improve the clock accuracy by *a few constant factors*, and (3) enhance the robustness of the clock in case of network partitions and malicious attacks.

– **Conserving energy and enhancing QoS via Quorum Systems.**
with Erik Demaine (MIT)

Very large sensor applications involving several hundreds of sensor nodes, such as environmental monitoring applications, bring up scalability and performance issues. We studied techniques that enhance the performance of large sensor networks and that can be applied in designing building blocks for routing, data management, and adaptive reconfiguration. In particular, we analyzed for the first time quorum systems (a technique designed to enhance the scalability and performance of distributed systems) in the specific context of sensor networks and energy conservation, and showed the unsuitability of previous quorum systems and metrics. This observation motivated us to revise quorums systems and their metrics by taking into account sensor limitations and properties, and to propose a novel class of quorum systems [10, 22].

Vehicular networks.

- **Data consistency in highly mobile networks.**
  Ensuring strong data guarantees in highly mobile networks is a challenging problem, which is crucial in several network tasks requiring node coordination (e.g., low duty cycle, object tracking, routing assistance) [2, 7]. More precisely, we studied under which mobility conditions it is possible to ensure strong data guarantees when node paths and speed are unknown such as in vehicular networks, and we provided a *minimum set* of mobility constraints that are satisfiable also in case of low density networks. We also proposed a novel class of quorum systems that is *provably optimal* in terms of communication costs and that does not violate the consistency property in case of high mobility. Simulation results based on the restricted random waypoint on a city section confirmed our theoretical results and indicate their applicability to vehicular networks.
- **Urban traffic monitoring system.**
  We are currently designing a traffic monitoring system based on static and mobile sensor nodes, which provides vehicles with alternative routes in case of traffic jam or emergency situations. Information obtained by vehicular sensor networks (e.g., public transportation) and combined by fixed sensor nodes provided with cameras, give an estimate of the local traffic flow and a short-term prediction. Such an estimate is transmitted to a centralized basestation where it is processed and used to maintain a global forecasting traffic model, which is used to design alternative routes to be transmitted to vehicles.

Secure distributed computing.

- **The Fleet system.**
  with Mike Reiter (CMU) and Dahlia Malkhi (Microsoft Research)
  Fleet is a middleware system implementing a secure distributed data repository for persistent objects [16]. It is primarily targeted for supporting critical applications such as governmental and financial applications, e-banking. Fleet is designed to be highly available, scalable to large numbers of servers and clients. It uses Byzantine quorum systems to improve the performance of the system and comprises a suite of novel intrusion-tolerant protocols based on quorum systems. We built a prototype e-voting application on top of the Fleet system and deployed Fleet to the American Defense Advanced Research Projects Agency (DARPA) in September 2001.
- **A secure and scalable digital time-stamping system.**
  Secure digital timestamps play a crucial role in many applications that rely on the correctness of time sensitive information. Previous time-stamping systems are based on *linking schemes*, which provide a relative temporal order by linking requests together. However, these systems do not scale well, and have coarse granularity and high latency. As a result, they are not suitable for applications requiring fine-grained and short-lived timestamps (e.g., stock trading, e-auctions, aggregation of real-time sensitive information, and temporal access control). We proposed a novel *provably correct* scheme based on real-time timestamps and Byzantine quorum systems, which overcomes those drawbacks and leads to performance enhancement [11, 13]. Its fine granularity, improved scalability and efficiency make our scheme particularly suitable to real-time applications such as quoting

applications. We are currently working on applying these results to design a e-stock market system.

- **Efficiency and scalability in intrusion-tolerant data consistency protocols.**
  Coordination protocols that are resilient to malicious attacks are notorious for having high computational and communication cost. We studied ways to improve the efficiency and scalability of intrusion-tolerant (Byzantine-tolerant) coordination protocols in case of very large data. We addressed this problem using a randomized approach and one-way collision-resistant hash functions [12, 23], and using an optimistic information dispersal protocol that dynamically adapts as failures are detected [20]. We also explored weaker data consistency models, such as causal consistency in shared memory [15] and in a dynamic client/server model [14].

- **Randomized automatic theorem proving.**
  with Chee Yap (NYU)

  We studied the problem of efficiently proving or disproving conjectures in Set Theory [25] and in Geometry [3, 17]. We proposed a novel hybrid symbolic-numerical approach that proves the validity of a geometric conjecture *by examples*, by generating instances of its construction. Our work led to the *first polynomial-time* probabilistic method for proving conjectures in Geometry, and to the extension of the Schwartz's randomized zero-test for multivariate polynomials to radical expressions. Our system is part of the Core package for the exact numerical computation.

PUBLICATIONS

Journals

1. D. Tulone. *On the feasibility of global time estimation under isolation conditions in wireless sensor networks.* Algorithmica, 49(4), pp. 386-411. Accepted June 2005.
2. D. Tulone. *Ensuring data consistency in highly mobile networks via quorum systems.* In Ad Hoc Networks, vol. 5, issue 3, pp. 1251-1271.
3. D. Tulone, C. K. Yap, C. Li. *Randomized zero testing of radical expressions and Elementary Geometry theorem proving.* In Automated Deduction in Geometry, LNCS 2061, pp. 58–82, Springer 01.

Under revision:

4. D. Tulone, S. Madden. *Time series forecasting for efficiently answering queries and detecting similarities in sensor networks.* Journal submission, July 2006.
5. D. Tulone. *A hierarchical model-based framework for answering queries in very large sensor networks.* Journal submission, September 2007.

Conferences

Wireless sensor networks:

6. D. Tulone, S. Madden. *An energy-efficient querying framework in sensor networks for detecting node similarities.* In Proc. of the 9th Intl. ACM Symp. on Modeling, Analysis and Simulation of Wireless and Mobile Systems, pp. 191–300, Oct 06.
7. D. Tulone. *Is it possible to ensure strong data guarantees in highly mobile networks?* In Proc. of the 5th Annual Mediterranean Workshop of Ad hoc Networks (MedHoc), Jun 2006.
8. D. Tulone, S. Madden. *PAQ: Time series forecasting for approximate query answering in sensor networks.* In Proc. of the 3rd European Conf. in Wireless Sensor Networks, pp. 21–37, Feb 06.
9. D. Tulone. *A resource–efficient time estimation for wireless sensor networks.* In Proc. of the 4th ACM Workshop of Principles of Mobile Computing, pp. 52–59, Oct 04.
10. D. Tulone, E. Demaine. *Redesigning quorum systems for energy conservation in sensor networks.* In Proc. 2nd Intl. Conf. on Wireless Algorithms, Systems, and Applications, pp. 147–157, Aug 2007.
11. D. Tulone, M. Srivastava. *Inspect: a General Framework for On-Line Detection and Diagnosis of Sensor Faults.* In Proc. of ITA Conf., Sept 07.

Secure distributed computing:

12. D. Tulone. *A secure and scalable digital time–stamping service.* In Proc. of the IEEE Intl. Conf. on Communications (ICC): Network Security and Information Assurance Symp., Jun 06. **Best paper**.
13. D. Tulone. *Enhancing efficiency of Intrusion–tolerant coordination protocols via hash functions.* In Proc. of the 10th ACM Conf. Euro–Par 2004, pp. 587-595, Sept 04. An extended version can be found in Tech. Report, ISTI CNR, Pisa. Apr 03.
14. D. Tulone. *How efficiently and accurately can a process get the reference time?* Intl. Symp. on Distributed Computing, Brief announcement, pp. 25-32, Oct 03.
15. D. Tulone. *Ensuring causal consistency in Internet–based services with arbitrary failures.* In 6th IEEE Intl. Symp. on Autonomous Decentralized Systems, fast abstract, pp. 17-18, April 03.
16. R. Baldoni, C. Spaziani, S. Tucci–Piergiovanni, D. Tulone. *An implementation of causal memories using the writing semantic.* In Proc. of 6th Intl Conf. On Principles of Distributed Sys, pp. 41-50, Dec 02.
17. D. Malkhi, M. K. Reiter, D. Tulone, E. Ziskind. *Persistent objects in the Fleet system.* In Proc. of the 2nd IEEE Darpa Information Survivability Conference and Exposition, Vol. 2, pp. 1126-1137, Jun 01.

Randomized algorithms:

18. D. Tulone, C. Yap, C. Li. *Randomized zero testing of radical expressions and Elementary Geometry theorem proving.* In Proc. of Intl. Workshop on Automated Deduction in Geometry, Sept 00.

Other manuscripts:

19. D. Tulone. *A secure and scalable digital time-stamping service.* In preparation for journal submission.
20. D. Tulone. *An optimistic adaptable information dispersal and recovery scheme.* Manuscript, MIT CSAIL, May 06. In preparation for publication.
21. D. Tulone. *On the design of an intrusion-tolerant distributed function.* Manuscript, MIT CSAIL, Mar 06. In preparation for publication.
22. D. Tulone. *Mechanisms for energy conservation in wireless sensor networks.* Ph.D. thesis, Department of Computer Science, University of Pisa, Dec 05.
23. D. Tulone. *Efficiency and intrusion-tolerance: reconcilable aspects of the same coin?* Ph.D. thesis proposal, Dept. of Computer Science, Univ. of Pisa. Mar 04.
24. D. Tulone, C. Yap. *Cookbook bounds for multi–term recursions.* Manuscript, New York Univ., Feb 98.
25. D. Tulone. *ETNA: Extensible Theorem Prover in NAtural deduction, in Set Theory.* M.S. thesis (tesi di laurea), Dept. of Mathematics, Univ. of Catania, Mar 94.

TALKS & PRESENTATIONS

Invited talks:

- *A secure and scalable digital time-stamping system for e-commerce and real-time applications.* ETH, Zurich. May 2007.
- *SAF: a Similarity-based Adaptable query Framework for sensor networks.* Computer Science Department, La Sapienza, Roma, April 2007.
- *Time series forecasting for efficiently answering queries and detecting similarities in sensor networks.* Invited talk at IBM Research Lab (Zurich), and at Politecnico di Milano (Italy). Jul 2006.
- *SAF: An energy-efficient framework for approximate query answering and detecting similarities in sensor networks.* Yahoo Research and UPF, Barcelona. Sept 2006.
- *Trade-offs between data accuracy and efficiency in wireless sensor networks.* MIT. Feb 2005.
- *Persistent Objects in the Fleet System.* Invited talk at Universita' *La Sapienza*, at University of Salerno, and at University of Bologna. Jun 2002.
- *The Principles behind the Fleet System.* Bell-Labs, Murray Hill. May 2002.

Other talks:

- *Mechanisms for energy conservation in wireless sensor networks.* University of Pisa. Jun 2006.
- *A secure and scalable digital times-tamping service.* ICC, Istanbul. Jun 2006.
- *Is it possible to guarantee data consistency in highly mobile networks?* MedHoc, Lipari. Jun 2006.

- *PAQ: Time series forecasting for approximate query answering in sensor networks.* EWSN, Zurich. Feb 2006.
- *Resource efficient time estimation in wireless sensor networks.* DIALM-POMC, Philadelphia. Oct 2004.
- *Enhancing efficiency of intrusion–tolerant coordination protocols via one–way hash functions.* EURO-PAR, Pisa. Aug 2004.
- *Efficiency and intrusion–tolerance: reconcilable aspects of the same coin?* University of Pisa. Mar 2004.
- *How efficiently and accurately can a process get the reference time?* DISC, Sorrento, Italy. Oct 2003.
- *Enhancing efficiency via hash functions.* ISTI CNR, Pisa. May 2003.
- *Ensuring causal consistency in Internet–based services with arbitrary failures.* ISADS, Pisa. Apr 2003.
- *Randomized Zero Testing of Radical Expressions and Elementary Geometry Theorem Proving.* ADG, Zurich. Sept 2000.
- *Reliability of a mechanical real-time scheduler.* Department of Applied Mathematics, University of Kaiserslautern. Sept 1994.

SOFTWARE

- Author of the *Prover system*, a probabilistic automatic theorem prover, which is part of the *Core library*, a library designed for the exact geometric computation. It is downloadable at http://cs.nyu.edu/exact/core/prover. Deployed in March 1999.
- Co-author of the Fleet system, a middleware system that implements an intrusion–tolerant data repository for persistent shared objects. Deployed to Darpa in September 2001.
- SAF framework for approximately answering queries and detecting data similarities [4]. May 2006.

COMPUTER SKILLS

- *Programming Languages:* Java, C, C++, Lisp, Ada.
- *Operating Systems:* Unix and Windows platforms.
- *Distributed System:* RMI, CORBA, EJB, JINI.
- *Algebraic Systems:* Mathematica, Maple.