Daniela Tulone
UCLA, 3563 Boelter Hall
Los Angeles, CA 90095-1596
E-mail: tulone@ee.ucla.edu

December 28, 2006

Professor Ali H. Sayed
Electrical Engineering Department
University of California
Los Angeles, CA 90095-1594

Dear Professor Sayed,

I am glad to have the opportunity to apply for a Visiting Assistant Professor position at your Department. Enclosed is a copy of my recent curriculum vitae, which includes a list of my publications and references, and my teaching and research statement.

I am currently working as a 6-month visiting Post-Doc at the EE Department and at CENS with Professor Mani Srivastava. Previously, I worked at MIT as a Post-Doc researcher. I hold a Ph.D. in Computer Science from University of Pisa and MIT, a M.S. in Computer Science from the Courant Institute of Mathematical Sciences, and a B.S. and M.S. in Mathematics from University of Catania. Working at the EE Department for 2 years as a Visiting Assistant Professor would give me the opportunity to foster my research collaborations initiated with people at UCLA during the last two months, and my teaching experience, so important for my future academic career.

I appreciate being considered for your faculty position. Please let me know if I can provide you with any additional information to support my application.

Sincerely,

Daniela Tulone

Daniela Tulone

Center for Embedded Networked Sensing
UCLA, 3563 Boelter Hall
Los Angeles, CA 90095-1596

E-mail: tulone@ee.ucla.edu
http://theory.csail.mit.edu/~tulone

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 *quorum systems.* Consistency of shared data, coordination protocols, and secure distributed systems.
Recent work in wireless sensor networks includes approximate query answering via time series forecasting, time synchronization, data consistency in highly mobile networks, and data integrity.

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.
- **B.S., and M.S., (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, optimization models. Final exam score 60/60.

WORK EXPERIENCE

- **UCLA, EE Department and CENS.**
  Visiting Post-Doc Researcher, Los Angeles, USA. October 2006–April 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.**
  Member of Technical Staff, 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. Researcher, Catania, Italy. February 1995–June 1996.

HONORS AND AWARDS

- **Best paper award.** ICC, Network Security and Information Assurance Symposium. Istanbul, June 2006. D. Tulone. *A secure and scalable digital time-stamping service.*
- **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** (about 250 applicants).
- **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 one of the 15 grants from Consorzio Archimede to work on mathematical models.

PROJECTS

Sensor Networks

- **Approximate query answers via time series forecasting.**
  with Samuel Madden (MIT)
  We proposed an adaptable energy-efficient query framework, called SAF, for approximately answer queries at the sink and detecting anomalies and data similarities, based on time series forecasting [1, 5, 7, 17, 22]. SAF relies on a class of lightweight and adaptable time series models, and on a suite of novel algorithms for monitoring and adapting the prediction model, and for detecting anomalies and data similarities at no additional communication cost. This is obtained using a novel definition of data similarity based not on raw data but on data models, and using a novel clustering algorithm that is provably optimal in the number of clusters. SAF consumes very little energy and is highly flexible, offering user-defined bounds, and tunable data rates and strategies. Simulation results based on real data traces confirmed the above mentioned features.
- **Time synchronization in sensor networks.**
  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 [2, 8, 12]. We propose a suite of deterministic and probabilistic time estimation algorithms that reduce the error growth by at least a factor of 2 by exploiting either the sign of the clock deviation, or the clock history (time series). These algorithms 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 conserve energy, or to improve the accuracy and robustness of the clock in case of node isolation, or network partitions, or malicious failures.

- **Quorum systems as tools for energy conservation.**
  with Erik Demaine (MIT)
  Quorum systems are known techniques designed to enhance the performance and scalability of distributed systems. They tolerate failures and reduce the access cost per operation and the system load balance, which are all desirable features in very large sensor networks. We analyzed for the first time quorum 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 (e.g., load balance, access cost, and quorum capacity) by taking into account sensor limitations and properties. We applied these metrics to analyze our novel *geographic-based quorum* system, and proposed an efficient quorum-based data diffusion protocol [9, 22].
- **Data consistency in highly mobile networks.**

  We studied the problem of ensuring strong data guarantees in highly mobile networks, which is crucial in network tasks requiring strong coordination (e.g., low duty cycle, object tracking, routing assistance) [3, 6]. More precisely, we studied under which mobility conditions it is possible to ensure strong data guarantees when node paths and speed are unknown, and provided a *minimum set* of mobility constraints that are satisfiable also in case of low density networks. In addition, we showed that under our mobility model previous quorum systems fail to guarantee data consistency, and proposed a novel class of quorum systems, called MDQ, and a quorum system that is provably optimal in terms of communication cost. Simulation results based on the random waypoint and the restricted random waypoint on a city section confirmed our theoretical results.
- **Data integrity.**
  with Mani Srivastava (UCLA)
  The focus of this project is on ensuring data integrity in sensor networks via statistical techniques. More precisely, our goal is to study the causes of faults and their temporal and geographical scope, develop diagnosis algorithms, and validate them in real-life. More precisely, we have developed a fault detection and diagnosis general framework [19], which comprises a suite of models and novel algorithms for classifying faults, detecting their temporal and geographical scope, and when possible for repairing them. Indeed, the ultimate goal of this project is to guarantee the integrity of data produced at sensors or fused by intermediate nodes.

- **Secure distributed computing**
  - **The Fleet system.**
    with Mike Reiter (CMU) and Dahlia Malkhi (Microsoft Research)
    Fleet is a middleware system implementing a distributed data repository for persistent objects [15]. It is primarily targeted for supporting highly critical applications: in particular, the objects it stores maintain their correct semantics despite malicious attacks. 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, comprises a suite of novel intrusion-tolerant protocols based on quorum systems, and is dynamically extensible with new object types. We built a prototype voting application on top of the Fleet system. Fleet was deployed to Darpa in September 2001.
  - **A secure and scalable digital times-tamping 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 a performance enhancement [10, 12]. Its fine granularity, improved scalability and efficiency make our scheme particularly suitable to real-time applications, and applicable to mobile e-commerce.

- **Efficiency and scalability in Byzantine-tolerant data consistency protocols.**
  Intrusion-tolerant coordination protocols are notorious to be very costly from the computational and communication viewpoint. We studied ways to improve the efficiency and scalability of 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 [11, 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 a shared memory [14] and in a dynamic client/server model [13].

- **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 [4, 16]. 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 lead us 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, S. Madden. *Time series forecasting for efficiently answering queries and detecting similarities in sensor networks.* Submitted to Transactions on Sensor Networks.

2. D. Tulone. *On the feasibility of global time estimation under isolation conditions in wireless sensor networks.* To appear in Algorithmica.

3. D. Tulone. *Ensuring data consistency in highly mobile networks via quorum systems.* To appear in Ad Hoc Networks.

4. 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.

Conferences

<u>Wireless sensor networks:</u>

5. D. Tulone, S. Madden. *An energy-efficient querying framework in sensor networks for detecting node similarities.* To appear in Proc. of the 9th Intl. ACM Symp. on Modeling, Analysis and Simulation of Wireless and Mobile Systems, Oct 06.
6. 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.
7. 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.
8. 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.
9. D. Tulone, E. Demaine. *Redesigning quorum systems for energy conservation in sensor networks.* Submitted to conference.

<u>Secure distributed computing:</u>

10. 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**.
11. 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.

12. 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.
13. 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.
14. 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.
15. 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:

16. 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:

17. D. Tulone, S. Madden. *A hierarchical model-based query framework for large sensor networks.* Invited MSWIM paper for Computer Networks journal submission. In preparation.
18. D. Tulone. *A secure and scalable digital time-stamping service.* In preparation for journal submission.
19. D. Tulone, M. Srivastava. *Detecting and diagnosing faults in sensor networks via time series forecasting.* Manuscipt in preparation for conference 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:

- *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 querying answering and detecting similarities in sensor networks based on time series forecasting.* 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.

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.
- *The Principles behind the Fleet System.* Bell-Labs, Murray Hill. May 2002.
- *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 [1]. 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, CoCoA.

INTERNAL REFERENCES

Professor Mani Srivastava
EE Department, UCLA
E-mail: mbs@ucla.edu
Phone: 310-267-2098

Professor Mark Hansen
CENS Statistics, UCLA
E-mail: cocteau@stat.ucla.edu
Phone: 310-206-8375

Professor Deborah Estrin
Director, Center for Embedded Networked Sensing
EE and CS Department, UCLA
Email: destrin@cs.ucla.edu
Phone: 310-206-3923

EXTERNAL REFERENCES

Assistant Professor Stefano Basagni
Northeastern University, Boston, MA
Dept. of Electrical and Computer Engineering
E-mail: basagni@ece.neu.edu
Phone: 617-373-3061
(Ph.D. thesis's reviewer)

Assistant Professor Samuel Madden
Massachusetts Institute of Technology
CSAIL Laboratory
E-mail: madden@csail.mit.edu
Phone: (617)258-6643
(Mentor during my Post-Doc at MIT)

Professor Andrea Maggiolo
University of Pisa, Italy
Computer Science Department
E-mail: maggiolo@di.unipi.it
Phone: +39 050 2212 75
(Chair of the Ph.D. program)

Professor Alfredo Ferro
University of Catania, Italy
Dept. of Mathematics and Computer Science
E-mail: ferro@dmi.unict.it
Phone: +39 095 733 7032
(Advisor in Catania and at NYU)

Professor Pascal Felber
Universit de Neuchtel, Switzerland
Institut d'informatique
E-mail: pascal.felber@unine.ch
Phone: +41 32 718 27 09
(Collaborated with him at Bell-Labs and after)

Research and Teaching Statement

Daniela Tulone
EE Department and CENS, UCLA
tulone@ee.ucla.edu

1 Research statement

My primary research goal is to study mechanisms that improve the performance and the robustness of wireless sensor applications and of distributed client/server systems in wired networks. My research work has been driven by a strong interest in both theory and practice, I have been always fascinated by the impact that mathematical theories can have on technology and on real applications. As a result, my research work tends to build on mathematical foundations and apply them to practical problems, often leading to the realization of a prototype. My most recent work has focussed on wireless sensor networks, which brings up new challenges from both the theoretical and system viewpoint, and offers an ideal ground for interdisciplinary collaborations, which I find very interesting. I am currently working as a 6-month visiting Post-Doc with Professor Mani Srivastava at the EE Department and at CENS, which I find very enriching and vibrant research environments. Interacting with people, getting exposed to their projects, and learning about the challenges encountered during their deployment is very interesting and leaves large space to apply mathematical and statistical tools, as well as distributed computing techniques. I believe that the interaction among different research areas can lead to substantial improvements. Working at the EE Department for 2 years as a Visiting Assistant Professor would give me the opportunity to foster my research collaborations with people at UCLA and my teaching experience, so important for my future academic career.

1.1 Background

The strong interest for both theory and practice started during my Master in Mathematics drove me along my diverse professional path. Diving deeper into numerical and optimization techniques and being exposed to the European academic and industrial environment during a post-graduate school in Industrial Modeling right after my Master, was a clear step towards blending theory and practice. However, as my dominant interest was for Computer Science, I started working as a C.N.R. research associate on Automatic Theorem Proving in Geometry at University of Catania. During that period I took the decision of diving deeper into Computer Science and moving to the United States to study Computer Science. Having won an Italian national C.N.R. competition for young researchers for performing research abroad, I moved to the Courant Institute of Mathematical Science where I worked with Professor Chee Yap on algorithms, and at the same time I earned my Master in Computer Science. That year was important for disclosing a wide spectrum of interesting problems and applications, and for injecting in me the curiosity for industry. In that respect working at AT&T as a software engineer tremendously helped me to improve my technical skills and enter in touch with the real world, and most importantly it showed me that research and academia was most attractive to me. Indeed, during that period I continued my collaboration at NYU with Professor Chee Yap. That work led us to the first polynomial-time probabilistic method for proving geometric conjectures, and the system I designed has been part of the Core library since 1999. Since then I always have been in research, working as a staff member in the Secure System Research Department at Bell-Labs, Murray Hill, on the design of a secure and scalable distributed data repository and of data consistency protocols, and later as Ph.D. student to better pursue an academic career. My Ph.D. thesis work was done at MIT in collaboration with people from the theory and system group. It showed that statistical techniques, such as time series forecasting, and distributed and fault-tolerant techniques can be applied to sensor networks to conserve energy.

1.2 Approach

My approach draws from mathematical and statistical models, combinatorial techniques, such as quorum systems and distributed computing, to design protocols and techniques that can lead to practical solutions. The diverse experiences cumulated in academia, research labs and shortly in industry, have often helped me to look at problems from a different perspective and led me to enhanced solutions. For instance, the study of the clock error growth between consecutive clock adjustments using hardware clock information led me to a refinement of the optimality bound for external clock synchronization, and to a practical way to improve the clock accuracy and reduce energy consumption in sensor networks. The work done on clock synchronization and on distributed computing helped me to transform the hard problem of designing a secure fine-grained and scalable digital time-stamping into a simpler problem, and led me to an enhanced scheme. This work was awarded as best paper at ICC 2006. Other successful examples of interactions among different research areas are given by my work on data consistency in highly mobile ad hoc networks, which I viewed as a fault-tolerant problem with continuous failures and recoveries, and my work on proving geometric conjectures that lead to the first practical polynomial-time approach for combining numerical and symbolic techniques. These are encouraging examples showing the potentiality of collaborations among different research areas. I find particularly interesting and enriching working with people with different backgrounds, as it often brings up new ideas and different viewpoints.

1.3 Research accomplishments in sensor networks

My Ph.D. thesis work, which was done at MIT, proposed techniques for energy conservation in sensor networks. In particular, I showed that time series forecasting and distributed computing techniques, such as quorum systems, can be adapted to meet the limitations and characteristics of sensor networks and to reduce energy consumption. More precisely, I applied time series models to the problem of efficiently collecting data produced at sensors, and of accurately and robustly estimating the time. In addition, I analyzed quorum system techniques in the specific context of energy conservation in sensor networks, thus proposing revised quorum systems and metrics, and studied the problem of efficiently ensuring data guarantees also in the presence of high mobile nodes. I summarize below some of these results.

Approximate query answers. With Professor Sam Madden I proposed an adaptable energy-efficient query framework, called SAF, for approximately answer queries at the sink and detecting anomalies and data similarities, based on time series forecasting [1]. SAF relies on a class of lightweight and adaptable time series models tailor-fit to be computed at sensors, and on a suite of novel algorithms for monitoring and adapting the prediction model, and for detecting anomalies and data similarities at no additional communication cost. This is obtained using a novel definition of data similarity, based not on raw data but on data models, and using a novel clustering algorithm that is provably optimal in the number of clusters. SAF consumes very little energy and is highly flexible. It offers user-defined bounds and the possibility of dynamically tuning the data rate at sensors, and of choosing energy-efficient strategies. Simulations results based on real data traces confirmed the above mentioned features. The framework proposed in SAF is generic and has a wide spectrum of potential applications ranging from object tracking to data integrity and data calibration.

Time synchronization. I 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 [2]. More precisely, I proposed a suite of deterministic and probabilistic time estimation algorithms that reduce the error growth by at least a factor of 2 by exploiting either the sign of the clock deviation, or the clock history (time series). These algorithms 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 conserve energy in long-running environmental applications, or to improve the accuracy and robustness of the time estimate in case of network partitions or malicious failures.

Data consistency in highly mobile ad hoc networks. I studied the problem of ensuring strong data guarantees in highly mobile networks, which is crucial in network tasks requiring strong coordination (e.g., low duty cycle, object tracking, routing assistance). More precisely, I studied under which mobility conditions it is possible to ensure strong data guarantees when node paths and speed are unknown, and I derived a *minimum set* of mobility constraints that are satisfiable also in case of low density networks. In addition, I showed that under our mobility model previous quorum systems fail to guarantee data consistency, and proposed a novel class of quorum systems, called MDQ, and a quorum system that is provably optimal in terms of communication cost. Simulation results based on the random waypoint and on the restricted random waypoint on a city section confirmed these theoretical results.

Quorum systems as tools for energy conservation. Quorum systems are known techniques designed to enhance the performance and scalability of distributed systems. They tolerate failures and reduce the access cost per operation and the system load balance, which are all desirable features in very large sensor networks. With Professor Erik Demaine I analyzed for the first time quorum systems in the specific context of sensor networks and regarded them as tool for energy conservation. Showing the unsuitability of previous quorum systems and metrics designed for wired networks, motivated us to revise quorums systems and their metrics (e.g., load balance, access cost, and quorum capacity) by taking into account sensor limitations and properties. We applied these metrics to analyze a novel geographic-based quorum system, and proposed an efficient quorum-based data diffusion protocol.

1.4 Research accomplishments in secure distributed computing

While I was at Bell-Labs in the Secure System Research Department, I worked on the design of distributed coordination protocols that are resilient to malicious failures and highly scalable, and I continued working on these problems while I was at University of Pisa. I briefly describe below some of the projects I worked on.

The Fleet system. At Bell-Labs I worked with Professor Mike Reiter and Professor Dahlia Malkhi on the Fleet project. Fleet is a middleware system implementing a distributed data repository for persistent objects. It is primarily targeted for supporting highly critical applications: in particular, the objects it stores maintain correct semantics despite the arbitrary failure of a limited number of servers and, for some object types, of clients allowed to invoke methods on those objects. 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, comprises a suite of novel intrusion-tolerant protocols based on quorum systems, and is dynamically extensible with new object types. We built a prototype voting application on top of the Fleet system. Fleet was deployed to Darpa in September 2001.

A secure and scalable digital time-stamping system. This project originated while attending a talk on cryptography, and it was developed during my first year of Ph.D. at University of Pisa. Secure digital time-stamps 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 time-stamps (e.g., stock trading, e-auctions, aggregation of real-time sensitive information, and temporal access control). I proposed a novel *provably correct* scheme based on real-time time-stamps and Byzantine quorum systems, which overcomes those drawbacks and leads to a performance enhancement [5]. Its fine granularity, improved scalability and efficiency make our scheme particularly suitable to real-time applications. With Massimiliano Cuzzo, who is a senior engineer in a financial company, we are currently designing a prototype system for a stock trading application based on this time-stamping service.

Efficiency and scalability in Byzantine-tolerant data consistency protocols. Intrusion-tolerant coordination protocols are notorious to be very costly from the computational and communication viewpoint. While I was at University of Pisa, I studied ways to improve the efficiency and scalability of Byzantine-tolerant coordination protocols in case of very large data, using randomization and one-way collision-resistant hash functions, and a novel optimistic information dispersal protocol that dynamically adapts as failures are detected. During my visit at Universita' La Sapienza, with Professor Roberto Baldoni I also started exploring causal consistency models in shared memory and later in dynamic client/server models.

1.5 Research accomplishments in automatic theorem proving.

My Master thesis in Mathematics was on methods to efficiently prove and disprove conjectures in Set theory under the advise of Professor Alfredo Ferro. After my Master I continued working with him on automatic theorem provers and their interaction with algebraic systems, and while at NYU I focused on studying efficient methods to prove geometric conjectures. With Professor Chee Yap I proposed a novel hybrid symbolic-numeric approach that proves the validity of a geometric conjecture *by examples*, by generating instances of its construction. Our work led us 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.

1.6 On-going and future work

In the last two months with Professor Mani Srivastava I started working on the data integrity problem in sensor networks. Our goal is to detect and classify sensor faults and repair them when possible, using time series forecasting techniques. Indeed, the SAF framework can be adapted to efficiently detect anomalies and distinguish between sensor faults and an abnormal behavior of the physical phenomenon. We have already developed a suite of algorithms for diagnosing faults and detecting their temporal and geographical scope. We also plan to validate this framework using data traces and include a mechanism for data calibration and for ensuring data integrity, and ultimately design a tool for sensor networks.

In the last two months I started collaborating with other faculty at UCLA, such as with Professor Deborah Estrin, Professor Mark Hansen, and their students. For instance, we are applying statistical models to the Cyclops project to efficiently detect object motion at sensors and important patterns in the behavior of birds. Another on-going collaboration regards the problem of accurately and robustly providing a time estimate in a wide area deployment in Mexico for the study of earthquakes by applying some of the techniques I developed for clock synchronization. Other faculty members with whom I might establish collaborations in the future are Professor Gregory Pottie, Professor Izhak Rubin, and Professor Lieven Vandenberghe.
In the near future I also plan to continue working on mobility, more precisely to investigate the use of statistical models to predict the node motion and tackle to the problem of network partitions. Other problems that I would like to investigate in the long-run are data security in sensor networks, and the definition of weaker consistency models for ad hoc networks that are able to provide some sort of data guarantee also in case of sporadic network partitions and that can be applied to vehicular networks.

2 Teaching statement

Although I really enjoyed working in research labs, I find an academic job, which combines research and teaching, most attractive for me. In several occasions, during my Masters and my Ph.D., I had the chance to mentor young students, giving them orientation, and helping them in technical issues. In all those cases I found interacting with students a rewarding way to transmit my experience and my enthusiasm for scientific research, which also motivated me to improve my technical and communication skills. I always try my best to stimulate their curiosity, based on their background and interests, and help them going beyond of that particular problem, showing the relevance of those techniques for applications and for building a solid foundation for their career. It was rewarding to see that some of these students remained attracted by research and decided to go for a Ph.D., and that others with industrial interests decided to investing more time and energy in improving their technical skills as a preparation for their future job. With regard to subject matter I am prepared to teach the following courses:

EE 103	Applied Numerical Computing.
EE 236A	Linear Programming.
EE 136	Introduction to Engineering Optimization techniques.
EE 131A	Probability.
EEM 116	Logic Design of Digital Systems.
MATH33A	Linear Algebra and Application.
MATH31B	Integration and Infinite Series.
MATH32A	Calculus of Several Variables.
EEM116C	Computer Systems Architecture.
EE205A	Matrix Analysis for Scientists and Engineers
EE 209S	Special topics in Embedded Computing Systems.
EE 202B	Distributed Embedded Systems.
EE210A	Adaptive Filtering.
EE 241B	Kalman Filters.

I would also be glad to teach a graduate course aimed at showing the application of statistical and distributed computing techniques to sensor networks, by analyzing some fundamental problems such as time synchronization, data calibration, query answering, routing, and data consistency.

References

[1] D. Tulone, S. Madden. *Time series forecasting for efficiently answering queries and detecting similarities in sensor networks.* Submitted to Transactions on Sensor Networks.

[2] D. Tulone. *On the feasibility of global time estimation under isolation conditions in wireless sensor networks.* To appear in Algorithmica.

[3] D. Tulone. *Ensuring data consistency in highly mobile networks via quorum systems.* To appear in Ad Hoc Networks.

[4] 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.

[5] 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**.