

**COMPUTER SCIENCE**  
**Graduate Program Progress Report**

**Name:** Bruno Ribeiro

**Background:**

**Education:**

- **M.S.**, Computer Engineering. COPPE Research Institute, Federal University of Rio de Janeiro, March 2001 — July 2003.
- **B.S.**, Computer Science. Computer Science Department, Federal University of Rio de Janeiro, 1996 — 2001.

**Honors and Achievements**

- Sprint Scholarship grant, 2006 — 2007
- CAPES PhD. Scholarship, 2004 — 2008
- CNPq M.S. Scholarship, 2001 — 2003
- CNPq Undergraduate Research Assistantship, 1997 — 1999

**Advisor/Group:**

Professor Don Towsley, Computer Networks Laboratory (NETWORKS).

- Research Assitantships:
  - UMass – Department of Computer Science (under supervision of Don Towsley)  
Dates: 2004 — present  
*Network Measurement.* A foundational study on the value of information obtained from some network measurement techniques.  
*Wireless mobility.* A study on the stochastic geometrical nature of ad-hoc wireless networks.
  - UFRJ – Department of Computer Science (under supervision of Edmundo de Souza e Silva)  
Dates: 2001—2004  
*Path diversity.* Packet path diversity for continuous media applications. Also compared path diversity with packet block interleaving. (Master thesis)  
*Network dimensioning.* A study hired by the CPqD research institute on the impact of QoS Service Level Agreements for the dimensioning of IP networks.
  - UFRJ – Department of Pure and Applied Mathematics (under supervision of S. C. Coutinho)  
Dates: 1997—1999 (Undergraduate Research Assistant)  
*Commutative Algebraic Fields.* Studied Gröebner basis and its applications.

## Research:

### Peer-Reviewed Publications (in English)

- Bruno Ribeiro, Don Towsley, Tao Ye, and Jean Bolot. *Fisher Information of Sampled Packets: an Application to Flow Size Estimation*. To appear at ACM SIGCOMM Internet Measurement Conference (IMC'06), Rio de Janeiro, Brazil, October 2006.
- Weifeng Chen, Yong Huang, Bruno Ribeiro, Kyoungwon Suh, Honggang Zhang, Edmundo de Souza e Silva, Jim Kurose, and Don Towsley. *Exploiting the IPID field to infer network path and end-system characteristics*. Passive and Active Measurement Workshop (PAM'05), Boston, March 2005.
- Ribeiro, B. F. M., Rodrigues, P. H. A., Marcondes, C. A. C. *A SIP/H.323 Signaling Gateway Implementation for IP Telephony*. IEEE Latin American Network Operations and Management Symposium (LANOMS'01), 2001, Belo Horizonte, Minas Gerais.
- Coutinho, S. C., Ribeiro, B. F. M. *On Holomorphic Foliations without Algebraic Solutions*. Experimental Mathematics, v.10, n.4, p.529 - 536, 2001.

### Patents

- *Packet sampling and the use of TCP protocol information for flow size estimation*. Sprint wants patent the idea and is willing to cover all patent costs as long as UMass IPO agrees with some of their conditions. The next step is to talk to the UMass IPO.
- *Fisher information and stream processing for flow size estimation*. Work in progress but I cannot talk about it (Sprint asked me not to). Sprint wants patent it under the same conditions mentioned above. Apologies for the vague title.

### Thesis

- M.Sc Thesis: *On the Efficiency of Path Diversity for Continuous Media Applications*. Federal University of Rio de Janeiro, July 2003.
  - Committee: Edmundo de Souza e Silva (UFRJ), Valmir Barbosa (UFRJ), Virgilio Almeida (UFMG), Don Towsley (UMass).

### Critical analysis (Ad-hoc reviews)

1. IEEE/ACM Transactions on Networking.
2. IEEE INFOCOM'03.
3. IEEE GLOBECOM'02.
4. IEEE ICC'02, ICC'03.

### Non-reviewed Publications

- Bruno Ribeiro, Edmundo de Souza e Silva and Don Towsley. *On the Efficiency of Path Diversity for Continuous Media Applications*. UMass CMPSCI Technical Report 05-19.

- Bruno Ribeiro. *Joint Control and Routing Optimizations on Overlay Networks with an OSP Budget Constraint*. UMass CMPSCI Technical Report 06-34. August, 2006.

### **Publications in Progress**

- Weifeng Chen, Bruno Ribeiro, Gerome Miklau, and Don Towsley, *Vulnerabilities of Prefix-preserving Anonymization of IP Traces*.
- Bruno Ribeiro, Tao Ye, and Don Towsley, *Fisher information and stream processing for flow size estimation*.

### **Presentations**

1. Presented *Fisher Information of Sampled Packets: an Application to Flow Size Estimation* (Ribeiro et. al.) at IMC'06, October 2006.
2. Presented *Fisher Information for Network Estimators* at IFIP W.G.7.3 Workshop (invited speaker), October 2006.
3. Seminar presentations: 7+.
4. Networks group presentations: 2.
5. Sprint Advanced Technology Laboratories presentations: 3.

### **Internships:**

- Sprint Advanced Technology Laboratories, Burlingame, CA, Summer 2006
- Sprint Advanced Technology Laboratories, Burlingame, CA, Summer 2005

### **Service:**

- Tea Totaller, Spring 2004 semester.
- Volunteer during CS Saturday poster presentation 2006.
- Volunteer during Candidate's Week 2005.
- Maintainer, Computer Networks Group publications page, Fall 2006 – present.

### **Teaching:**

#### **Prior Experience**

- Co-Instructor, Undergraduate Computer Networks course, UFRJ, 1st semester of 2003
- Teaching Assistant, Undergraduate Computer Networks course, UFRJ, 2nd semester of 2002
- Teaching Assistant, Undergraduate Computer Networks course, UFRJ, 1st semester of 2002

- Teaching Assistant, Undergraduate Computer Networks course, UFRJ, 2nd semester of 2001

**Core Curriculum:**

Area	Course	Semester	Grade
Systems:	CMPSCI 653	Fall 2004	A
	CMPSCI 677	Spring 2006	I
Theory:	CMPSCI 601	Spring 2005	A-
	CMPSCI 611	Fall 2004	A-
AI:	CMPSCI 683	Fall 2005	B+
	CMPSCI 689	Fall 2006	(in progress)

**Synthesis:**

- Title: Vulnerabilities of Prefix-preserving Anonymization of IP Traces
- Advisors: Don Towsley and Gerome Miklau
- Abstract: Accurate network measurement through trace collection is critical for advancing network design and for maintaining secure, reliable networks. While the publication of network traces is critical to network research, it often threatens the privacy of network users. Trace anonymization has emerged as a common technique used to balance the benefits of trace publication and threats to privacy. In this work we study the practical security of prefix-preserving IP trace anonymization, which has been adopted by the research community and used to protect a number of publicly-available traces. We describe a systematic attack which allows an adversary to use public information about the traced network to break the anonymization and recover the actual IP addresses of anonymized hosts in the trace. We demonstrate the effectiveness of this attack on real network traces, measure the impact of the attack for varying degrees of public information, and suggest solutions to resist the attack.
- Final Approval Date (expected): December 2006.

**Reference Letters Requested:**

1. Professor Don Towsley
2. Professor Jim Kurose
3. Professor Gerome Miklau

**Personal Statement:**

The ultimate goal of my PhD. is to become a faculty researcher at a top university. This means that I approach the curriculum with a strong bias for research. My research attention during these two years at UMass has been focused on network monitoring and distributed source coding for wireless networks. While realizing that research progress is usually delta incremental, I feel that I can say that my current works are relevant contributions to the networks field. My research breadth spans through information theory, anonymization techniques, performance evaluation, estimation theory and networks.

My research achievements during my first year were limited by my classes workload and my adaptation to a new country and life. Nevertheless, during the first weeks at UMass and in collaboration with my advisor Don Towsley, professor Jim Kurose and Weifeng Chen (networks student), we developed a method used to measure network one-way delay differences that does not require Internet end-hosts to be instrumented. It uses an IP field (IPID) and a characteristic of Windows operating systems to do the measurement. Our finds were published at PAM'05 along with finding from other colleagues.

During my first year, a fair amount of my time was also devoted to the study of stochastic geometry and its applications.

During my second year I developed an interest in the use of information theoretic principles for network monitoring. In collaboration with my advisor Don Towsley and Tao Ye from Sprint, we looked at the problem of estimating network flow level metrics from sampled packets at Internet core routers. This work has received very good feedback from the many researchers I had a chance to talk to. Positive feedback also came from the IMC'06 reviewers, where it was published. An excerpt of one of the reviewers: "I wish IMC had more of this kind of foundational research. How many of our "measurement" papers will still be worth reading in 10 years time". Beyond its academic impact, this work brought three other developments: (1) A one year scholarship from Sprint (tuition, fees and stipend). (2) Invitation for an internship in the following year. (3) Sprint wants to patent the estimation method we developed.

My second year also featured a work in collaboration with Gerome Miklau, Weifeng Chen (now a former student) and my advisor Don Towsley on the issue of breaking prefix-preserving anonymization. This type of anonymization is widely used by the network community to anonymize network traces collected at gateway routers because of its ability to anonymize while keeping IP address structure information, which is beneficial for routing research. In this work we devise an attack and show that prefix-preserving anonymization is not safe. This work is yet to be submitted, and will also be part of my synthesis project.

During this third year, aside from continuing efforts over network monitoring and prefix-preserving anonymization, my interests have also broaden to include distributed source coding over wireless networks. On the network monitoring side, the continuation of my previous work cannot be publicly disclosed, at Sprint's request. Sprint shows a strong interest over its patent. On the prefix-preserving anonymization side, after showing previous scheme to be ineffective, we focus our attention on devising an anonymization scheme that is at the same time strong against attacks and useful for routing research. This development is based on data base anonymity concept of k-anonymity. My latest research topic, on source coding over wireless networks, is ambitious and promising. It uses the physical broadcast nature of the

up-link wireless channel to overcome the deficiencies of previous distributed source coding methods. I am currently reviewing the literature in this area.

I feel that I'm ready to start my PhD. thesis research and, with the completion of my required classes, I intend to be much more productive next year.

---

Advisor's Signature

---

Student's Signature