

# JOHN SINCLAIR OTTO

Ph.D. Candidate  
AquaLab Distributed Computing Research Group  
EECS Department, Northwestern University

2145 Sheridan Road  
Evanston, Illinois 60208  
(847) 467-3250

905 Sherman Ave Apt B4  
Evanston, Illinois 60202  
(262) 424-6965

<http://www.cs.northwestern.edu/~jot836/>

[jotto@eecs.northwestern.edu](mailto:jotto@eecs.northwestern.edu)

## RESEARCH INTERESTS

---

Networking; distributed computing; content delivery; CDN brokering; peer-to-peer systems.

## EDUCATION

---

- **Ph.D., Computer Science** (December 2012)  
**Northwestern University**, Evanston, Illinois  
Advisor: Prof. Fabián E. Bustamante  
GPA: 3.793/4.0
- **M.S., Computer Science** (November 2011)  
**Northwestern University**, Evanston, Illinois  
Advisor: Prof. Fabián E. Bustamante
- **B.S., Computer Science** (June 2007)  
**Northwestern University**, Evanston, Illinois  
GPA: 3.735/4.0 (*cum laude*)

## PROFESSIONAL EXPERIENCE

---

- **Graduate Research Assistant.** AquaLab Group, EECS Dept., Northwestern 2007 to present
  - Design, execute, and publish research pertaining to systems, networking, and content delivery.
  - Assist in writing and editing NSF grant proposals.
  - Mentor junior graduate, masters, and undergraduate students regarding research, writing, and career decisions.
- **Intern.** AT&T Labs Research; Florham Park, NJ Summer 2011
  - Evaluated traceroute-based approach for identifying and locating network points of presence.
  - Studied usage of CDN brokering techniques and policies employed by Netflix (application-based CDN brokering) and Cotendo (CDN brokering as a service).
- **Intern.** Telefónica Research; Barcelona, Spain Summer 2010
  - Designed and evaluated a system for load-aware allocation of elastic computation resources.
  - Analyzed call detail record datasets to infer aggregate patterns of social connections and travel.
- **Teaching Assistant.** EECS Department, Northwestern
  - Operating Systems Fall 2007, 2008, 2009
  - Introduction to Computer Systems Spring 2007
  - Introduction to Networking Winter 2007
  - Responsible for grading assignments, and holding office hours to answer students' questions.
  - Teach weekly recitation sessions and (except for Networking) the full class 1-2 times per quarter.

- **Undergraduate Research Technician.** Evanston Northwestern Healthcare: Research Institute
  - Genetic analysis software development for Gejman Laboratory April 2005 to June 2007
  - Learned and applied statistical concepts and genome analysis tools and techniques.
  - Developed LSMatchMaker, a tool for large-scale genetic association analysis.
  - Used this tool to perform single- and multi-point linkage analysis on genome scan data in search of a genetic association for schizophrenia.
- **Undergraduate Summer Fellow.** Intelligent Information Lab, EECS Department, Northwestern
  - Intelligent Information Systems Development June to September 2006
  - Developed natural language processing tools for manipulating news story text (News at Seven).
  - Developed a mobile application with a barcode reader to provide users with product information.
- **Undergraduate Research Assistant.** O’Halloran Lab, Chemistry Department, Northwestern
  - Analysis of ZIP/CDF proteins in *Plasmodium falciparum* September 2003 to June 2005
  - Self-taught use of BLAST and MEME tools to search for related proteins based on highly-conserved sequences in ZIP/CDF proteins to search for a novel malaria treatment approach.

## PUBLICATIONS

---

### Under Submission

- “Temporal Rate Limiting: cloud elasticity at a flat fee.” John S. Otto, Rade Stanojevic, and Nikolaos Laoutaris. Under submission to NetEcon. 2012.

### Conference Publications

- “Distributed Systems and Natural Disasters – BitTorrent as a Global Witness”. Zachary S. Bischof, John S. Otto, and Fabián E. Bustamante. In Proc. of ACM CoNEXT Special Workshop on the Internet and Disasters. 2011.
- “On Blind Mice and the Elephant – Understanding the Network Impact of a Large Distributed System”. John S. Otto, Mario A. Sánchez, David R. Choffnes, Fabián E. Bustamante, and Georgos Siganos. In Proc. of ACM SIGCOMM. 2011. (Otto, 2011a)
- “Crowdsourcing ISP Characterization to the Network Edge”. Zachary S. Bischof, John S. Otto, Mario A. Sánchez, John P. Rula, David R. Choffnes, and Fabián E. Bustamante. In Proc. of ACM SIGCOMM Workshop on Measurements Up the STack (W-MUST). 2011. (Bischof, 2011)
- “Distributed or Centralized Traffic Advisory Systems – The Application’s Take”. John S. Otto and Fabián E. Bustamante. In Proc. of IEEE SECON. June 22-26, 2009. (Otto, 2009a)
- “Down the Block and Around the Corner – The Impact of Radio Propagation on Inter-vehicle Wireless Communication”. John S. Otto, Fabián E. Bustamante, and Randall A. Berry. In Proc. of IEEE ICDCS. June 22-26, 2009. (Otto, 2009b)

### Posters & Demos

- “Dasu – ISP Characterization from the Edge”. Demo. Mario A. Sánchez, John S. Otto, Zachary S. Bischof, and Fabián E. Bustamante. ACM SIGCOMM. 2009.
- “Turning the Corner in the City—The Impact of Radio Propagation on Inter-vehicle Wireless Communication”. J. S. Otto, F. E. Bustamante and R. A. Berry. HotMobile. 2009.
- “Distributed Systems on Vehicular Ad-Hoc Networks; from Simulation to Deployment & Back”. J. Otto, R. Letcher, D. Choffnes, O. Soyannwo and F. E. Bustamante. Motorola Undergraduate Research Fellowship session. 2006.
- “LSMatchMaker: An Integrated Environment for Association Analysis”. J. S. Otto, R. S. Gejman, J. Duan, J. O’Connell and M. Martinez. World Congress on Psychiatric Genetics XIII. 2005. And International Genetic Epidemiology Society Conference. 2005.

## INVITED LECTURES

---

- October 2009 – “Down the Block and Around the Corner: The Impact of Radio Propagation On Inter-vehicle Wireless Communication,” University of Illinois at Chicago, Computational Transportation Science (IGERT Program).

## RESEARCH PROJECTS

---

### **Studying the Internet from the edge.**

To gain a complete picture of the impact of Internet systems (e.g. BitTorrent, DNS), it is often necessary to capture the view from the edge of the network. Monitoring systems from inside the network can miss significant characteristics that are only visible from the end user’s perspective. For instance, the paths that peer-to-peer data take through the network frequently include peering links, many of which are not visible from the core of the network.

We leverage the diverse vantage points in peer-to-peer systems as a platform for collecting information about the network from the edge. One use of this knowledge, for example, is to characterize networks and their links, allowing us to infer their roles and underlying business models. Although peerings in Internet eXchange Points (IXPs) and the use of similar services (e.g. Remote Peering) are increasing, information about these networks and maps of participants are often incomplete and kept updated by hand (e.g. PeeringDB). We are developing methods for identifying and mapping these networks using traceroute measurements made between peers from the edge of the network.

### **Peer-to-peer and its impact on ISPs.**

The unpredictable patterns of connections made by P2P software can generate large volumes of potentially costly cross-ISP traffic. By mapping BitTorrent data flows to the AS paths the data takes through the network, we are able to infer the volume and cost of BitTorrent to ISPs, from eyeball ASes up to Tier 1 networks (Otto, 2011a).

### **Trends in content distribution: CDN Brokering and public DNS.**

The Domain Name System, which is a global distributed database of names, provides a foundation for nearly every system running on the Internet. Content Delivery Networks, for example, often use the location of a user’s DNS resolver (frequently operated by the user’s ISP) as a proxy for determining the user’s location so that they can redirect the user to a nearby replica server.

We are examining several aspects of the DNS system to evaluate its impact on, for example, CDNs and identify trends in DNS usage and CDN performance. How does the increasing popularity of public DNS resolvers (e.g. Google Public DNS), whose DNS resolvers are not necessarily close to the user, impact the performance of DNS-based CDN redirections? (Otto, 2012).

### **Challenges in Vehicular Networking.**

Vehicular-based systems provide exciting platforms that can enable, for example, scalable pervasive monitoring of urban areas. However, such environments present several challenges in terms of communication and data dissemination.

We evaluate inter-vehicular communication characteristics in urban environments to more accurately model the performance of systems in simulation (Otto, 2009b). Using our more-realistic model of urban communication, we examine the trade-offs between distributed and centralized data dissemination approaches to inform system design (Otto, 2009a). Finally, we explore the use of low-cost pollution sensors mounted on vehicles for urban air quality monitoring. This approach offers a scalable and cost-effective solution that provides improved spatial and temporal resolution relative to traditional pollution monitoring systems.

## SKILLS AND EXPERTISE

---

- Large-scale data analysis: MapReduce, Hadoop & Dumbo (Python interface)
- Cloud computing (Amazon EC2) and research testbeds (PlanetLab)
- Programming Languages: Python, Java, MySQL, Javascript, bash, C/C++, HTML, PHP, Perl

## HONORS AND AWARDS

---

- Royal E. Cabell Graduate Fellowship, Northwestern 2007 to 2008
- Summer Undergraduate Research Fellowship, Cog-Sci Dept., Northwestern Summer 2006
- Motorola Undergraduate Research Fellowship, Motorola Spring 2006
- Weinberg College Summer Research Grant, Chemistry Dept., Northwestern Summer 2004

## COMMUNITY SERVICE

---

- Leader, *Graduate Students in EECS* at Northwestern (2008 to 2010)
- Coordinator, *Graduate Research Seminar in CS & CE* at Northwestern (2008 to 2010)