

Jonathan R. Stanton  
Department of Computer Science  
The Johns Hopkins University  
3400 North Charles Street  
Baltimore, MD 21218

jonathan@cs.jhu.edu  
<http://www.cnds.jhu.edu/~jonathan/>  
Phone: (410) 516-0454 (Office)  
Fax: (410) 516-6134

## Research Interests

I have worked for the last five years on the research and development of a scalable, high-performance, wide-area group messaging toolkit called Spread. The Spread toolkit was made publicly available over two and a half years ago and has become a popular package for reliable messaging. Currently, Spread is believed to have thousands of users in commercial, research, and teaching environments. Several popular applications use Spread, such as the Apache-SSL secure web server, a distributed logging service for the Apache and thttpd web servers, the native database replication in Postgres, and the Zope application server. Spread is part of FreeBSD and is included in some Linux distributions. We have recorded over 2500 distinct downloads from our web site.

About three years ago, the difficult problem of security for group-oriented systems became apparent and I, along with my adviser Professor Yair Amir and with Professor Tsudik from UC Irvine (previously from USC/ISI), started a project to create a secure group communication system.

Also, over the last five years I have helped my adviser and Professor Baruch Awerbuch to establish the Center for Networking and Distributed Systems. I have been involved in building and maintaining the computing infrastructure, selecting and working with the masters and Ph.D students, and obtaining financial resources to support research. About two years ago I helped write two DARPA proposals that were funded, and became a co-PI on both of them. I believe I have a solid start on what it takes to be a successful systems professor who leads an active lab.

**My goal is to explore how to construct reliable, high-performance distributed systems, and then actually build them.**

I want to work in the intersection of networking, distributed protocols and systems, and security research. One of the major challenges today is to make networks, and the distributed protocols that run over them, more reliable and trustworthy. This is definitely a long term task. It is also a task that will require collaboration among researchers, which for me is one of the most rewarding parts of research.

I believe one of the best ways to manage the complexity of modern systems is to build well-defined, correct, and efficient software toolkits. These toolkits provide abstractions that isolate some of the hard parts of building distributed and networked applications. Some of the interesting challenges when creating toolkits are discovering what service or semantics provide the right abstraction, designing efficient protocols to provide that abstraction, and proving both the correctness and security of the toolkit.

During the next few years, I would like to focus on the following specific research areas.

- Group communications — My Ph.D research focused on the development of a high performance, scalable group messaging system that is optimized for wide-area networks such as

the Internet, yet still provides flexible and strong semantic guarantees. The resulting system is one of the highest performing and most widely used group messaging toolkits in the world. Group communication systems simplify the design and development of reliable distributed services such as consistent replication, interactive conferencing, and multicast publishing. Strong group messaging services for Internet wide networks will be an important tool for future replicated application servers. I plan to continue research on how to make group messaging scale up to large Internet applications, what types of semantic models are most appropriate for different applications, and how group communication techniques can help solve different types of problems.

- Distributed system security — Distributed systems that function over public networks require strong security services, but strong security often has high performance costs. As part of my work adding security to the Spread messaging toolkit, I have focused on how an integration of security with the design of the system can provide strong security without a prohibitive performance impact. I am interested in working on how to address security and performance in distributed systems in an integrated way. This research topic calls for working with other researchers, especially those in security and cryptography.
- Scalable, decentralized services — The predominant architecture used today in Internet services is client-server, in which the server may be distributed but it is still centralized. The bandwidth and computational resources available today to Internet users, and concerns about privacy and the power of centralized authority, have driven a movement to develop decentralized models for Internet services. I have recently begun research on how decentralized services can be built so that they provide scalability in both users and data management and resilience to the faults and changes in resources that are inevitable in any dynamic system. I am interested in continuing this line of research in services such as directory lookups and network data storage.
- Custom network protocols — One very interesting question in networking today is how custom protocols running on an overlay network interact with the underlying physical network. As part of my Ph.D research, I developed a reliable, weakly ordered transport protocol and designed an overlay network architecture to efficiently disseminate multicast messages and the control information required to correctly deliver them. I believe the overlay network model has great potential for solving real applications need for smart, highly efficient network services. To accomplish this goal, research is needed in developing protocols and understanding how they can be used successfully and safely on the Internet.