5th International Workshop on Networking Meets Databases (NetDB 2009)

Boon Thau Loo
University of Pennsylvania
3330 Walnut Street, Philadelphia, PA 19104
boonloo@cis.upenn.edu

Stefan Saroiu
Microsoft Research
One Microsoft Way, Redmond, WA 98052
ssaroiu@microsoft.com

1. INTRODUCTION

The Workshop on Networking Meets Databases (NetDB) is a venue that aims to bring together researchers from the systems and networking community and the database community. Many current research areas, such as cloud computing, data-center networking, sensor networks, network management, or social networks, raise research problems that lie at the boundary between these two communities. The workshop's goal is to foster an environment in which researchers from both communities can discuss ideas that will shape and influence these emerging research areas. The workshop encourages submissions of early work, with novel and interesting ideas. The expectation is that work introduced at NetDB, once fully thought through, completed, and described in a finished form, may be relevant to conferences such as SOSP, OSDI, SIGCOMM, SIGMOD, VLDB, NSDI, or ICDE.

Traditionally, NetDB has been co-located with either the IEEE International Conference on Data Engineering (ICDE'05, ICDE'06, ICDE'08) or the USENIX Symposium on Networked Systems Design and Implementation (NSDI'08). This year is the first time when the workshop was co-located with the ACM Symposium on Operating Systems Principles (SOSP). By co-locating with SOSP, we aimed at raising the interest from the wider systems community on topics relevant to NetDB.

2. REVIEW PROCESS

NetDB received a total of 16 submissions and many of these submissions were of very high quality. Most papers covered topics in hot and emerging research areas relevant to NetDB, such as declarative systems, online social networks, and large-scale MapReduce-like systems. As in previous years, NetDB continued to attract novel work at the confluence of the networking and database communities. All papers received three reviews, and we added a fourth review for borderline cases. We held a phone meeting that lasted an hour and a half. During the meeting, we discussed 14 papers and we accepted seven.

3. WORKSHOP DAY

NetDB 2009 was held on the final day of SOSP (October 14th, 2009) starting at 1:30pm. The workshop program included seven paper presentations, followed by a 1.5 hour panel discussion. While we intended the workshop to conclude by 7pm, the panel fostered a healthy debate that made us end the workshop closer to 8pm. Officially, NetDB had 21 registrations, although we counted more people in the audience. More details of the program can be found online on the workshop's website.

3.1 Presentations

The accepted papers covered topics in cloud computing (one paper), social networks (one paper), publish-subscribe systems (one paper), RFID monitoring (one paper), and declarative systems (three papers). The papers along with the slides used in their presentations can be found on the workshop's website. The accepted papers were:

- **Data Indexing for Stateful, Large-scale Data Processing**. Dionysios Logothetis and Kenneth Yocum (UC San Diego).
  One-line summary: This paper presents techniques to integrate indexes with stateful/incremental batch processing at large scale (systems such as MapReduce restart the complete computation when presented with data).

- **Scaling Online Social Networks without Pains.** Josep M. Pujol, Geogos Siganos, Vijay Erramilli, Pablo Rodriguez (Telefonica Research).
  One-line summary: This paper describes a mechanism for partitioning a graph among multiple servers, so that social network computations on the graph can be carried out in parallel.

- **Generating Wide-Area Content-Based Publish/Subscribe Workloads.** Albert Yu, Pankaj K. Agarwal, Jun Yang (Duke University).
  One-line summary: This paper describes a workload generator for publish/subscribe systems parameterized by data gathered from Google Groups.

- **Architectural Considerations for Distributed RFID Tracking and Monitoring.** Zhao Cao, Yanlei Diao, Prashant Shenoy (University of Massachusetts).
  One-line summary: This paper considers centralized and distributed architecture designs for RFID monitoring systems that combine inference and query processing techniques.

- **Declarative Transport: A Customizable Transport Service for the Future Internet.** Karim Mattar, Ibrahim Matta, John Day, Vatche Ishakian, Gonca Gursun (Boston University).
  One-line summary: This paper uses declarative networking techniques to specify and implement policies used in transport protocols.

- **I Do Declare: Consensus in a Logic Language.** Peter Alvaro, Neil Conway, Russell Sears, Tyson Condie, Joseph M. Hellerstein (UC Berkeley).
  One-line summary: This paper uses declarative networking techniques to implement the Paxos consensus protocol, and shows that primitives used in consensus protocol specifications map directly to simple declarative networking language constructs.

- **On the Declarativity of Declarative Networking.** Yun Mao (AT&T Labs - Research).
  One-line summary: This paper surveys recent systems that are based on the use declarative networking, in order to investigate their level of declarativity.

Each paper was allocated a 30 minute slot that included 5-7 minutes allocated to Q&A.

3.2 Panel Discussion

During the final hour and a half of the workshop, we held a panel titled "Declare your declarativity". The panelists are Fred Baker (Cisco), Joe Hellerstein (UC Berkeley), Eddie Kohler (UCLA), Arvind Krishnamurthy (University of Washington), Petros Maniatis (Intel Research Berkeley), and Timothy Roscoe (ETH Zurich). We also were able to have a small wine and beer open bar that we think helped ignite the panel discussion.

The purpose of the panel was spurred by the recent interests in the use of declarative languages to implement complex systems and networking protocols. No less than three papers accepted at NetDB 2009 covered topics in declarative systems. The appeal of declarative languages is their compactness: with just a few tens of lines, one can describe the entire semantics of a complex protocol, such as the consistency semantics of a file-system, or the TCP protocol, or the Paxos consensus protocol. Such compactness can make it easier for developers to ensure that their protocol implementations are semantically correct and they have few bugs. At the same time, there is little experience with using declarative languages for solving systems problems, and some are skeptical about their ease of use and ease of debugging.

Each member of the panel took a stand on whether systems will benefit from using declarative languages. The discussion raised a couple of interesting points:

- An argument in favor of declarative systems is that declarative implementations are brief (few number of lines of code). Some people pointed out that lines of code is not an adequate metric for measuring how good an implementation is.

- Another benefit of declarative systems is that it gives semantics to an implementation and it *could* make it easier to build systems on top of these implementations. For example, it might make it easier to run verification tools or bug findings tools. The consensus was that this might be a fruitful future direction for this research area.

- A system might be more successful when written in programming languages everyone understands.

- Many of the declarative systems implemented by the research community already have robust implementations running in industry. Instead, implementing declarative systems that industry needs might have a bigger impact.

- Any system has a configuration. Many configuration files today are already written in a declarative fashion.

- Some rule-of-thumbs used when building systems appear to be hard to capture declaratively. For example, "be liberal in what you receive and conservative in what you send" is a principle for building networked systems. It is unclear whether a declarative implementation can capture such a principle.

This brief list is far from a rigorous summary of the panelists' points of view. Instead, we encourage the reader to visit the NetDB's website where the panelists' slide decks are posted.

4. WORKSHOP INFORMATION

4.1 Program Committee

The program committee of NetDB 2009 was formed by:

- Atul Adya, Google
- Brian Cooper, Yahoo! Research
- Mary Fernández, AT&T Labs
- James Hamilton, Amazon,
- Joe Hellerstein, UC Berkeley
- Boon Thau Loo, University of Pennsylvania (co-chair)
- Sam Madden, MIT
- Ratul Mahajan, Microsoft Research
- Petros Maniatis, Intel Berkeley
- Stefan Saroiu, MSR Redmond (co-chair)
- Emin Gün Sirer, Cornell

4.2 Workshop Website

`http://netdb09.cis.upenn.edu`

4.3 Acknowledgments

We would like to thank Marc Fiuczynski, the SOSP workshop chair, for his help with the workshop organization. We are also very grateful to Microsoft Research Redmond for a generous sponsorship.