Scalability Challenges and Solutions for Emerging Networks

Professor Kenneth P. Birman[1]
Dept. of Computer Science, Cornell University
CEO, Reliable Network Solutions, Ithaca NY and Mountainview CA

Extended Abstract

The rapid growth of computer networks, so evident during the past five years, continues unabated. Within five years, it will be common to encounter networks, dedicated to a single enterprise, which encompass tens of thousands of computing devices.

Some of these networks will host extremely sensitive applications. For example, one can easily envision medical computing networks that would connect hospital systems to monitoring devices and even drug administration devices in the homes of housebound patients. The electric power grid is evolving in ways that will demand new kinds of computer-mediated control. Military technology planners in the United States are designing a whole series of network-oriented information systems intended to tie all the sensors and databases in a battlefield environment together, creating a shared information resource that would be of extremely high value to mission planners, tacticians, and even directly available to the soldier in the field. Many of the associated platforms will have wireless links, hence the environment will be a highly dynamic one. Failures, other forms of disruption, and link degradation must be anticipated and viewed as "routine" properties of these environments.

To solve such problems, we will need dramatic advances in terms of the tools available for managing, controlling, and programming large distributed applications. Today's technology has many limitations and is clearly inapplicable to the kinds of demanding large-scale settings just enumerated. For example, existing tools suffer from a tendency to "melt down" under stress, exhibit severe sensitivity to configuration, lack adequate security, and often exhibit loads and overheads that rise with the size of the network.

Our challenge, as researchers, is to identify the design principles that would enable a new generation of solutions having the required properties. Needed are technologies that would be inherently robust, provably scalable, and sufficiently self-organizing to adapt as conditions change in the network.

At Cornell University, the Spinglass project has been successful in solving an important class of such problems. At the core of our work is a new style of gossip-based communication protocol. We are using this protocol in support of a variety of systems-programming tools. The talk will discuss two of them: Bimodal Multicast [1], a scalable reliable multicast protocol having probabilistic reliability properties, and Astrolabe [2], a virtual distributed database constructed entirely through peer-to-peer interactions among the components of a large system. Both technologies are shown to be stable under stress, arbitrarily scalable without growth in communication or processing loads, and to have real-time properties (in the sense of predictability even when faults occur) [3]. Reliable Network Solutions, a company we founded in 1998, is now offering commercial implementations of these technologies.

References

[1] Bimodal Multicast. Ken Birman, Mark Hayden, Oznur Ozkasap, Zhen Xiao. *ACM Transactions on Computer Systems*17:2 (41-88), May 1999.

[2] Scalable Domain Management and Aggregation. Robbert van Renesse, Kenneth P. Birman, Werner Vogels. Cornell University Dept. of Computer Science Technical Report, March. 2001.

[3] Fighting Fire with Fire: Using Randomized Gossip to Combat Stochastic Scalability Limits. Kenneth P. Birman, Indranil Gupta, Robbert van Renesse. Cornell University Dept. of Computer Science Technical Report, March. 2001.

URLs:
http://www.cs.cornell.edu/Info/Projects/Spinglass
http://www.rnets.com

[1] This work was supported in part by DARPA-AFRL/IFGA Grant No. F30602-99-1-0532. Speaker's contact information: Ken Birman, Dept. of Computer Science, Cornell University; Ithaca, NY 14850. 607-255-9199. ken@cs.cornell.edu

Proceedings of the IEEE International Symposium on Network Computing and Applications (NCA'01)
0-7695-1432-4/02 $17.00
IEEE COMPUTER SOCIETY