P. Krishnan (PK)'s Projects

(A Brief Summary)

April 15, 2000

I like to study real-life computing and networking systems problems, develop practical techniques to solve the problems, and implement and verify the proposed techniques. My work spans multiple areas of computer science, as described below[1]. My work exploits the synergy between theory and practice to model and develop better solutions for problems in traditional and emerging areas in computing and communications. It has evolved to emphasize on speed in implementation and projects with not only technical merit but also significant market value.

My current work (described in Sections 1.1 and 1.2) is in the area of Internet Protocol (IP) management. Specifically, I am working on IP Network Configuration, and caching techniques for improving WWW performance.

1 Network and Services Management

The problems in network management arise from issues involved in managing continuously evolving artificial complex systems. My work in network and services management is in the area of configuration and performance management.

1.1 IP Network Configuration

The IP Network Configurator is a unique tool for achieving automation in IP network configuration by not only reducing the cost of configuring an IP network but also enabling novel, compelling and advanced management applications.

The IP Network Configurator is a centralized and automated IP network configuration tool. It provides an innovative model for router configuration: rather than a device-centric view, the configurator sees the router as an intrinsic part of the network, understands the network semantics, topology and protocols being run on the network, and uses this information to determine configuration. The tool is multi-vendor, GUI-based, and policy-based. It has several industry-first features: it is an online network configuration tool that collects information from the routers directly, provides the administrator with multiple views of the network depending on routing protocols, simplifies adding a new router to the network, does semi-automatic routing protocol configuration/reconfiguration, performs version management of configuration files and enables disaster recovery.

IPNC incorporates several key research innovations including new network-aware software models for configuration, consistency checking of configurations, innovative ways to access routers, and novel algorithms to sequence changes when the changes are made to multiple routers. The tool is written almost entirely in Java, and has a few C components. (See publications [25, 26, 28], and patent [1].)

[1] All references herein refer to the attached list of publications.

1.2 Internet Management and World-Wide Web Caching

Caches are used for saving bandwidth and reducing latency in WWW browsing. In our work, we study transparent caching. We formulate the cache location problem as an optimization problem and show that, in general, it is intractable. We then solve two important special cases: the problem of a single source that wishes to minimize the average accessing delay for clients, and the problem of using caches to reduce the overall traffic in a backbone with some regular structure. We experimentally demonstrate the performance improvement using real web server data and traceroute data that we collected. (See publications [24, 9, 27] and patent [2].)

I have also worked in the area of proxy cache management, and managing co-operation in WWW proxies. We built a prototype of PROXYMAN: a proxy caching manager and agent conforming to the DMTF/DMI standard. PROXYMAN also allows cooperation between the management agents, allowing for exchange of management information between the agents. (See publications [4, 29, 30] and patent [4].)

1.3 Firewall Management

This work deals with managing rules in firewalls to optimize their performance. The work was done in the context of the Lucent Managed Firewall. (See patent [3].)

1.4 Internet Management and User Defined Protocols

User Defined Protocols provide applications with some type of quality of service information so that they can adapt to the variable (network) resources. This involves melding the network management activity of monitoring networks with providing feedback to the application. An infrastructure is provided so that the application can dynamically change its resource requirements and desired quality. (See publications [1, 8 10].)

1.5 Provisioning

Provisioning processes are an important component of any network management system. Provisioning involves all activities to implement changes to the network. We have looked at the problem of improving provisioning processes, and have developed a new and automated technique for this problem. Our method (described in publication [13]) is based on representing workflows by finite state machines augmented with the notion of attribute sets. The technique has been successfully tested on trunk provisioning workflows in telecommunication networks.

2 Prediction in Computer Systems

Prediction or anticipation of the future, and adapting to it is a central theme in many computer systems applications. The issue of optimal and practical prediction in systems was the topic of my Ph.D. thesis dissertation [31]. In my thesis, I considered typical *prediction* scenarios (i.e., the *what*, *when*, and *how much* is expected to happen) and developed online prediction algorithms for improving response time performance of databases and operating systems.

We looked at prediction in systems environments from a new perspective and proposed and analyzed the novel idea of using *data compression* techniques for prediction (publications [3,6,19,22,23], patent [6]). In addition to proving the theoretical optimality of our methods under general analysis models, we demonstrated their practical merit by using these ideas to develop good *prefetchers* for object-oriented database applications. Apart from being accurate, the predictors need to be efficient in the memory and time needed to perform their actions.

Our techniques have a broad range of applications beyond prefetching. *Estimating selectivity* is crucial in databases to determine the best search order during query optimization. Based on our intuition that data compression and prediction are closely related, we developed a novel method for estimating alphanumeric selectivity (publication [11]). This technique is currently being used by IBM in some of their database products. We have also developed intelligent learning algorithms for effective device management in mobile machines (publications [2,16]); also see Section 3.

3 Mobile Computing

Mobile computing is an important area of computing and communications research. People are increasingly using small but powerful portable devices to remain connected at all times. I have studied data access issues in mobile computing, and developed an efficient storage hierarchy for mobile machines (publications [2, 7,16,17,18,20,21,32], patents [5,7]); in particular, we have looked at the power/response-time tradeoffs of disk spindown policies and using flash memory in portable computers. Analytical modeling of such systems gives new insight into the tradeoffs and leads to interesting new learning algorithms for disk spindown and related problems (publications [2, 16]).

I have also studied the problem of locating mobile phones. This issue is crucial to achieving truly mobile communication. Fast location of mobile phones is important from both the customer's point of view and the service provider's point of view. We have studied for the first time the issue of queuing delays in mobile phone location, and developed important design criterion for mobile phone search (publications [5,12]). We have developed new models for mobile computing, and studied load balancing issues in improving the performance as seen by mobile users (publications [14,15]).

In [1,8], we study techniques to reduce power during communication by mobile applications by intelligently controlling the state of the wireless communication device. Careful experimentation of a novel transport/application-level protocol demonstrates the benefits of our techniques.

Publications

Journal Publications

[1] R. Kravets and P. Krishnan. "Application-Driven Power Management for Mobile Communication," to appear in ACM/URSI/Baltzer Wireless Networks (WINET) special issue of best papers from MOBICOM '98.

[2] P. Krishnan, P. Long, and J. S. Vitter. "Adaptive Disk Spindown via Optimal Rent-to-Buy in Probabilistic Environments," in *Algorithmica* **23(1)**, January 1999, pp. 31–56.

[3] P. Krishnan and J. S. Vitter. "Optimal Prediction for Prefetching in the Worst Case," in *SIAM J. on Computing*, **27:6**, December 1998, pp. 1617–1636. (This is an expanded version of [19].)

[4] P. Krishnan and B. Sugla. "Utility of Cooperating Web Proxy Caches," in the *Elsevier J. on Computer Networks and ISDN Systems* **30**, 1998, pp. 195–203. (This appears in the *Proceedings of the 1998 World Wide Web Conference*.)

[5] D. Goodman, P. Krishnan and B. Sugla. "Minimizing Queuing Delays and Number of Messages in Mobile Phone Location," in the *ACM/Baltzer J. on Mobile Networks and Applications*, **1:1**, 1996, pp. 39–48.

[6] J. S. Vitter and P. Krishnan. "Optimal Prefetching via Data Compression," in *J. of the ACM*, **43:5**, September 1996. (This is an expanded version of [23].)

[7] F. Douglis, P. Krishnan and B. Bershad. "Adaptive Disk Spindown Policies for Mobile Computers," in *Computing Systems*, **8:4**, 1995. (This is an expanded version of [17].)

Refereed (Conference) Publications

[8] R. Kravets and P. Krishnan. "Power Management Techniques for Mobile Communication," in *Proceedings of the 1998 ACM/IEEE International Conference on Mobile Computing and Networking*, (MOBICOM '98).

[9] P. Krishnan, D. Raz, Y. Shavitt. "Transparent En-route Cache Location in Regular Network Topologies," DIMACS Workshop on Robust Communication Networks: Interconnection and Survivability, November 1998.

[10] R. Kravets, K. Calvert, P. Krishnan, and K. Schwan. "Adaptive Variation of Reliability," in *Proceedings of the 1997 High Performance Networking Conference* (HPN '97).

[11] P. Krishnan, J. S. Vitter and B. Iyer. "Estimating Alphanumeric Selectivity in the Presence of Wildcards," in *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, Montreal, Canada, June 1996, pages 282–293.

[12] D. Goodman, P. Krishnan and B. Sugla. "Design and Evaluation of Paging Strategies for Personal Communications," in *Proceedings of the Conference on Multiaccess, Mobility and Teletraffic for Personal Communications*, B. Jabbari, P. Godlewski and X. Lagrange eds., Kluwer, Paris, France, May 1996, pages 131–144.

[13] P. Krishnan, S. Nag and B. Sugla. "An Automated Technique for Designing and Improving Provisioning Processes," in *Proceedings of the 1996 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Kyoto, Japan, April 1996, pages 329–338.

[14] B. Awerbuch, Y. Azar, E. F. Grove, M. Y. Kao, P. Krishnan and J. S. Vitter. "Load Balancing in the L_p Norm," in *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*, Milwaukee, WI, October 1995.

[15] E. Grove, M. Kao, P. Krishnan and J. S. Vitter. "Online Perfect Matching and Mobile Computing," in *Lecture Notes in Computer Science:955, Proceedings of the 4th International Workshop on Algorithms and Data Structures*, Kingston, Ontario, August 1995.

[16] P. Krishnan, P. M. Long and J. S. Vitter. "Learning to Make Rent-to-Buy Decisions with Systems Applications," in *Proceedings of the 12th International Conference on Machine Learning,* July 1995.

[17] F. Douglis, P. Krishnan and B. Bershad. "Adaptive Disk Spindown Policies for Mobile Computers," in *Proceedings of the Second USENIX Symposium on Mobile and Location-Independent Computing*, April 1995.

[18] F. Douglis, R. Cáceres, F. Kaashoek, P. Krishnan, K. Li, B. Marsh and J. Tauber. "Storage Alternatives for Mobile Computers," in *Mobile Computing*, Chapter 18, Imielinski and Korth eds., Kluwer, 1995.

[19] P. Krishnan and J. S. Vitter. "Optimal Prediction for Prefetching in the Worst Case," in *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms*, January 1994.

[20] F. Douglis, P. Krishnan and B. Marsh. "Thwarting the Power Hungry Disk," in *Proceedings of the 1994 Winter USENIX Conference*, San Francisco, January 1994.

[21] B. Marsh, F. Douglis and P. Krishnan. "Flash Memory File Caching for Mobile Computers," in *Proceedings of the 27th Hawaii Conference on Systems Sciences*," IEEE, January 1994.

[22] K. Curewitz, P. Krishnan and J. S. Vitter. "Practical Prefetching via Data Compression," in *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, Washington, D.C., May 1993.

[23] J. S. Vitter and P. Krishnan. "Optimal Prefetching via Data Compression," in *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*, San Juan, Puerto Rico, October 1991.

Other Publications

[24] P. Krishnan, D. Raz, Y. Shavitt. "The Cache Location Problem," submitted for journal publication.

[25] P. Krishnan, T. Naik, and R. Sequeira. "Sequencing of Configuration Operations for IP Networks," to be submitted.

[26] S. Das, P. Krishnan, S. Mazumdar, T. Naik, G. Ramu, B. Sugla, G. K. Swanson. "IP Network Configurator: Information Model and Architecture," in preparation.

[27] P. Krishnan, D. Raz, and Y. Shavitt. "Transparent En-route Caching in WANs?," Work in progress presentation at the Fourth International Web Caching Workshop, March 1999.

[28] B. Sugla and P. Krishnan. "IP Network Configurator: Advancing the State of the Art in IP Network Deployment and Operations," whitepaper, November 1998.

[29] P. Krishnan and B. Sugla. "PROXYMAN: Proxy Caching Manager With Advanced Features," in *Lucent Netdays 97*, September 29–October 1, 1997.

[30] P. Krishnan and B. Sugla. "Proxy Caching Manager With Enhanced Features," Bell Laboratories, Restricted Document, August 1996.

[31] P. Krishnan. "Online Prediction Algorithms for Databases and Operating Systems," Ph. D. thesis, Brown University, Providence, RI, May 1995.

[32] P. Krishnan, F. Douglis and B. Marsh. "StoreSim: A Storage Sub-System Simulator," Matsushita Information Technology Lab. Technical Report MITL-TR-103-94, Confidential.

Patents

1. P. Krishnan, S. Mazumdar, T. Naik, G. Ramu. "Method to Sequence Changes for IP Network Configuration," filed, October 1999 (provisionally filed, April 1999).

2. P. Krishnan, D. Raz, Y. Shavitt. "Methods to Place Caches in a Wide Area Network," filed, March 1999.

3. P. Krishnan, D. Raz, B. Sugla. "Firewall Optimization and Management," filed, January 1999.

4. P. Krishnan and B. Sugla. "Technique For Effectively Managing Proxy Servers In Intranets," filed, November 1998.

5. "Adaptive Disk Spin-down Method for Managing the Power Distributed to a Disk Drive in a Laptop Computer," (with F. Douglis, B. Marsh and B. Bershad), United States Patent 5,493,670, issued February 20, 1996.

6. "Online Background Predictors and Prefetchers for Locality Management," (with J. S. Vitter and K. M. Curewitz), United States Patent 5,485,609, issued January 16, 1996.

7. "Method for Managing the Power Distributed to a Disk Drive in a Laptop Computer," (with F. Douglis and B. Marsh), United States Patent 5,481,733, issued January 2, 96.