Brad Nelson Karp

Office Address:
417 S. Craig St., Suite 300
Pittsburgh, PA 15213
USA
+1 412-297-4111
bkarp@cs.cmu.edu, brad.n.karp@intel.com
http://www.cs.cmu.edu/~bkarp/

Home Address:
2 School House Ln.
Pittsburgh, PA 15232
USA
+1 412-621-8851

Education

Harvard University, Cambridge, Massachusetts.
Doctor of Philosophy in Computer Science. Graduated November 2000.
Thesis title: *Geographic Routing for Wireless Networks*.
Master of Science in Computer Science received June 1995.
Advisor: H. T. Kung.

Yale University, New Haven, Connecticut.
Bachelor of Science in Computer Science. Graduated May 1992.
Thesis title: *Massively Parallel Algorithms for Adaptive Function Approximation in One and Two Dimensions*.

Research Interests

Networked systems and their algorithms: Internet worm quarantine; wireless and mobile networks; sensor networks; Internet architecture.

Research Experience

Autograph: 2003–present. Led two Carnegie Mellon Ph.D. students in the design and building of Autograph, a distributed system that automatically derives the signatures of novel Internet worms. A signature—a byte pattern that is found uniquely in a particular worm's payload—may be used to filter worm traffic, and thus halt a worm's spread. To do its work, Autograph analyzes traffic that crosses an edge network's DMZ in two steps: suspicious flow selection, and content analysis. The former stage chooses a subset of network traffic using heuristics (*e.g.,* identification of traffic sources that port scan) that focus the attention of the system on traffic that is likely to be malicious, to reduce generation of false positive signatures, and to avoid the computational cost of analyzing innocuous traffic, which in the common case predominates. The latter stage finds the most frequently occurring payload substrings within the suspicious traffic pool, and proposes those substrings as signatures. Autograph instances distributed at edge networks throughout the Internet share their observations; this information sharing significantly speeds Autograph's generation of signatures after a worm's initial release. An evaluation of Autograph on real DMZ traces from multiple sites reveals the system generates accurate signatures quickly: it would have generated a signature for the Code-RedI-v2 worm before 2% of vulnerable Internet hosts had become infected, without generating any signatures that cause false positives. A prototype of the system is in continuous operation on Intel Research Pittsburgh's network. Ongoing work targets generation of signatures that match *polymorphic* worms, which vary their payload on every infection attempt.

Further details available at http://www.cs.cmu.edu/~hakim/autograph/.

Open DHT: 2003–present. As project co-PI, led two UC Berkeley Ph.D. students in designing, building, and deploying Open DHT, a publicly accessible distributed hash table (DHT) service. A DHT allows hosts distributed across the Internet to form an overlay network that provides a simple hash-table-like functionality, through which *key-value pairs* may be stored distributedly across those hosts. While a wide variety of DHT-based applications have been proposed by computer scientists in the past four years, research in the area has had little impact, as measured in novel systems that have found use by a broad community of users. This dearth of deployed systems persists because researchers build prototype applications that incorporate a purpose-built DHT; no attention is paid to sharing and reuse of the DHT infrastructure. Building and maintaining a widely distributed overlay is too high a barrier to entry for any one application to succeed. To spur adoption of DHTs and drive research on real systems built upon them, we have built and deployed Open DHT, a *generic, reusable DHT service* that is broadly useful as a building block for distributed applications. The system supports put and get operations on key-value pairs, and runs 24 hours per day, 7 days per week, on a collection of over 200 infrastructure hosts distributed across five continents. The system is public in that any Internet host may use it. Open DHT incorporates novel approaches to authenticating data in a distributed system shared by mutually distrusting users; to fair storage allocation for soft-state storage among Internet-scale client populations; and to supporting a shared routing layer among heterogeneous applications. It is in use today in over 15 projects, by researchers at institutions including MIT, UC Berkeley, HP Labs, UCL, Intel Research, and UCLA.

Further details available at `http://opendht.org/`.

GPSR, CLDP, and DCS: 1998–present. Created and evaluated Greedy Perimeter Stateless Routing (GPSR), a scalable routing algorithm for wireless and sensor networks that uses the *positions* of routers and a packet's destination to make forwarding decisions. In link-state, distance-vector, and ad-hoc routing algorithms, each router relies on state concerning the full path between that router and each network destination. In contrast, by using geography for forwarding decisions, GPSR requires state describing only a router's *immediate neighbor* routers. As a result, GPSR uses far less memory at each router than do LS, DV, and ad-hoc routing algorithms, and finds routes more robustly after topology changes, while generating significantly reduced routing protocol message traffic. This work continues with the extension of GPSR by CLDP (the Crossing Link Detection Protocol), through which GPSR works *provably correctly* on *arbitrary connected network graphs, regardless of topology*; and with the application of GPSR to *data-centric storage (DCS) on sensor networks*, through which sensors may store and answer queries over sensed data in an energy-efficient fashion.

Further details available at `http://www.cs.cmu.edu/~bkarp/gpsr/`.

Reordering-Robust TCP: 2001–2003. With a student intern advisee, co-designed extensions to the TCP reliable transport protocol to improve its robustness on network paths that reorder packets. SACK TCP, the previous state-of-the-art TCP, performs poorly on such paths because it mistakes reordering for loss, and incorrectly sends more slowly. Reordering-Robust TCP uses measurements of the degree of reordering on a network path to adapt TCP's retransmission algorithm dynamically, such that the sender proactively avoids such spurious window reductions. The result is significantly increased throughput on networks that reorder. This work relaxes the constraint that the Internet deliver packets in order, and thus enables work on desirable network systems that violate this constraint, such as multi-path routing.

APRL and HUMR: 1997–1999. Designed and built a multi-hop wireless mobile network on commodity Metricom Ricochet radios. Created HUMR, an application-level UNIX driver for peer-to-peer IPv4 and IPv6 communication using Ricochet radios, that are ordinarily used only for PPP connections with Metricom's commercial wireless service. HUMR uses hints provided by the radios' firmware to detect acquisition and loss of neighbors. Created APRL, Any-Path Routing without Loops, a distance-vector routing system that *provably* does not acquire looping routes. Demonstrated the combined system running on mobile networks of dozens of hosts for the Air Force.

SABRA: 1995–1996. Built SABRA, an ATM ABR with EPRCA rate-based flow control system for Intel hosts with a CMU/Harvard/Intel-designed ATM network interface and the Harvard/Nortel CreditNet ATM switch. The research community widely expected ABR to require a dedicated ASIC on host interfaces. Motivated by the importance of software's flexibility in the evolution of flow control protocols, I built an optimized ABR system where the host CPU performs ABR in software. A SABRA host could sustain the full OC-3 ATM link rate; no ASIC was needed.

Mobile IP: 1993–1994. Collaborated in the design and building of the Harvard Mobile IP system, one of the first working Mobile IP implementations, when Mobile IP was still being defined. The Harvard system supports secure "short-cut" routing, through which wired hosts communicate directly with a roaming mobile host, rather than through the mobile host's home agent.

CreditNet: 1992–1995. As a member of a team comprised of H.T. Kung's lab at Harvard and a group at Nortel, co-designed the CreditNet switch, a multi-gigabit ATM switch with hardware support for per-circuit credit-based flow control. Built the first host-side N23 credit-based flow control system. This system assumes minimal host adapter hardware, and performs all credit cell processing on the host CPU. A modest 133 MHz Alpha CPU running this code can saturate an OC-3 ATM link.

Work Experience

Carnegie Mellon University, Computer Science Department, Pittsburgh, Pennsylvania: September 2002–present. Adjunct Assistant Professor. Co-taught graduate course in sensor networks; gave frequent guest lectures in graduate networking courses; co-advised two CMU Ph.D. students on the Autograph project.

Intel Research Pittsburgh, Pittsburgh, Pennsylvania: September 2002–present. Staff Researcher. Framed and led the Autograph project. Together with my co-PI, Sylvia Ratnasamy of Intel Research Berkeley, framed and led the Open DHT project. Advised Ph.D. student interns from Berkeley, Carnegie Mellon, and Princeton. Designed and ran Ph.D. student intern selection process for the lab.

ICSI Center for Internet Research, Berkeley, California: October 2000–September 2002. Research Scientist. Developed algorithms for improving TCP's robustness to packet reordering; developed techniques for geographic provisioning of wireless networks; contributed to design of a distributed hash table for sensor networks, built on GPSR.

AT&T Center for Internet Research at ICSI, Berkeley, California: June 1999–April 2000. Research Intern. Developed GPSR, a scalable geographic routing system for wireless networks.

Digital Equipment Corporation, Littleton, Massachusetts: June 1996–August 1996. Research Intern. Designed, simulated, and evaluated open-loop and best-effort traffic scheduling algorithms for input-buffered switch architectures.

USC/Information Sciences Institute, Arlington, Virginia: June 1995–August 1995. Research Assistant. Built a software ATM Forum ABR/EPRCA system for the NetBSD UNIX kernel running on the Intel Pentium processor with an experimental Intel ATM interface. Analyzed performance of resulting system.

Consultancy

Consultant in networked systems and algorithm design for firms including Nortel (1997), Ungermann Bass Networks (1994), Kendall Square Research (1993–1994), and Bell Communications Research (1992).

Teaching and Advising

Co-instructor (2003): Sensor Networks (Carnegie Mellon CS 15-829, graduate course).
Co-instructor (1996): Operating System Network Performance (Harvard CS 248, graduate course).
Teaching Fellow (1993): Operating Systems (Harvard CS 161/261, graduate and undergraduate course).

Advisor (2004): Three Ph.D.-student research summer interns at Intel Research.
Advisor (2003): Two Ph.D.-student research summer interns at Intel Research.
Advisor (2001): Three Ph.D.-student research summer interns at ICSI.
Co-advisor (1998): Undergraduate Computer Science thesis, graded *summa cum laude* by Harvard's Computer Science Department.

Resident Tutor in Computer Science (1995–1999): Dunster House. Advised and assisted Harvard undergraduate Computer Science students in thesis topic selection, course work, problem sets, *&c.*

Program Committees

ACM/USENIX WORLDS Workshop 2006 (affiliated with OSDI), **PC co-chair**.

ACM MobiCom Conference 2005.
IEEE INFOCOM Conference 2005.
IEEE NET+DB Workshop 2005 (affiliated with ICDE).

ACM MobiCom Conference 2004.
ACM/USENIX WORLDS Workshop 2004 (affiliated with OSDI).

IEEE RTSS Conference 2003 (Sensor Networks Track).
ACM WSNA Workshop 2003 (affiliated with MobiCom).

Grantsmanship

Co-PI, NSF Grant 0330178, SENSORS: Robust and Efficient Data Dissemination for Data-Centric Storage, with co-PIs S. Shenker, UC Berkeley and R. Govindan, USC, $430,000 over 2 years, awarded September 2003.

Awards and Prizes

Henry Dunster Tutor Prize (1997) for excellence in advising Harvard undergraduates.

Best Student Paper Award (1994) for Secure Short-Cut Routing for Mobile IP in *Usenix '94*.

Publications

Newsome, J., Karp, B., and Song, D., Polygraph: Automatically Generating Signatures for Polymorphic Worms, to appear in *Proceedings of the IEEE Symposium on Security and Privacy (Oakland 2005)*, May 2005.

Kim, Y.-J., Govindan, R., Karp, B., and Shenker, S., Geographic Routing Made Practical, to appear in *Proceedings of the Second USENIX/ACM Symposium on Networked System Design and Implementation (NSDI 2005)*, May 2005.

Gummadi, R., Kothari, N., Kim, Y.-J., Govindan, R., Karp, B., and Shenker, S., Reduced State Routing in the Internet, to appear in *Proceedings of the Third ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets 2004)*, November 2004.

Nath, S., Ke, Y., Gibbons, P., Karp, B., and Seshan, S., A Distributed Filtering Architecture for Multimedia Sensors, in *Proceedings of the First IEEE Workshop on Broadband Advanced Sensor Networks (BASENETS 2004)*, October 2004.

Kim, H.-A. and Karp, B., Autograph: Toward Automated, Distributed Worm Signature Generation, in *Proceedings of the 13th Annual Usenix Security Conference (Usenix Security 2004)*, August 2004.

Karp, B., Ratnasamy, S., Rhea, S., and Shenker, S., Spurring Adoption of DHTs with OpenHash, a Public DHT Service, in *Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS 2004)*, Springer-Verlag Lecture Notes in Computer Science Hot Topics Series, February 2004.

Gibbons, P., Karp, B., Nath, S., Ke, Y., and Seshan, S., IrisNet: An Architecture for a Worldwide Sensor Web, in *IEEE Pervasive Computing, Special Issue on Sensor and Actuator Networks*, IEEE Press, October-December 2003.

Zhang, M., Karp, B., Floyd, S., and Peterson, L., RR-TCP: A Reordering-Robust TCP with DSACK, in *Proceedings of the 2003 IEEE International Conference on Network Protocols (ICNP 2003)*, November 2003.
(Extended version published as ICSI Technical Report TR-02-006, July 2002.)

Ratnasamy, S., Karp, B., Shenker, S., Estrin, D., Govindan, R., Yin, L., and Yu, F., Data-Centric Storage in Sensornets with GHT, A Geographic Hash Table, in *Mobile Networks and Applications (MONET)*, Kluwer Academic Publishers, August 2003.

Tolia, N., Kozuch, M., Satyanarayanan, M., Karp, B., Bressoud, T., and Perrig, A., Opportunistic Use of Content-Addressable Storage for Distributed File Systems, in the *Proceedings of the Usenix 2003 Technical Conference*, June 2003.

Shenker, S., Ratnasamy, S., Karp, B., Govindan, R., and Estrin, D., Data-Centric Storage in Sensornets, in the *Proceedings of the First ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets 2002)*, October 2002.

Ratnasamy, S., Karp, B., Yin, L., Yu, F., Estrin, D., Govindan, R., and Shenker, S., GHT: A Geographic Hash Table for Data-Centric Storage, in the *Proceedings of the First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA 2002)*, September 2002.

Karp, B., Challenges in Geographic Routing: Sparse Networks, Obstacles, and Traffic Provisioning, in the *Proceedings of the DIMACS Workshop on Pervasive Networking*, May 2001.

Karp, B., *Geographic Routing for Wireless Networks*, Doctoral Disseration in Computer Science, Harvard University Division of Engineering and Applied Sciences, October 2000.

Karp, B. and Kung, H. T., Greedy Perimeter Stateless Routing for Wireless Networks, in *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, August 2000.

Gaynor, M., Karp, B., and Kung, H. T., A PC-Based ATM Link Delay Simulator, in *Proceedings of the 1998 Summer Computer Simulation Conference (SCSC '98)*, July 1998.

Karp, B., Mankin, A., Kung, H.T., Demirtjis, A., and Edwards, B., An Implementation Study of ABR/EPRCA, *ATM Forum Contribution 96-587*, April 1996.

Blackwell, T., Chan, K., Chang, K., Charuhas, T., Karp, B., Kung, H.T., Lin, D., Morris, R., Seltzer, M., Smith, M., Young, C., Bahgat, O., Chaar, M., Chapman, A., Depelteau, G., Grimble, K., Huang, S., Hung, P., Kemp, M., Mahna, I., McLaughlin, J., Ng, M.T., Vincent, J., Watchorn, J., An Experimental Flow-Controlled Multicast ATM Switch, in *Proceedings of the First Annual Conference on Telecommunications in Massachusetts*, October 1994.

Blackwell, T., Chan, K., Chang, K., Charuhas, T., Gwertzman, J., Karp, B., Kung, H.T., Li, W.D., Lin, D., Morris, R., Polansky, R., Tang, D., Young, C., Zao, J., Secure Short-Cut Routing for Mobile IP, in *Proceedings of the Usenix Summer 1994 Technical Conference*, June 1994.

Karp, B. and Moulin, P., Implementation of Multiresolution Regression Splines on the MasPar: Function and Image Estimation in Parallel, *Bell Communications Research Internal Technical Memorandum*, February 1993.

Karp, B. and Bischof, C., Increasing the Granularity of Parallelism and Reducing Contention in Automatic Differentiation, *Argonne National Laboratory Technical Memorandum ANL/MCS-TM-142*, November 1990.

Personal

Languages: French and German (moderate proficiency).
Non-technical interests: Choral Music (Tenor I in the Harvard and Yale Glee Clubs and the UC Berkeley Chamber Chorus); Jazz Piano (playing and listening); 19th- and 20th-Century Fiction; Fine Food (cooking and eating); Theatre; Skiing; Improvisational Comedy (former musical director of Yale's Viola Question group).