

# CALL FOR PAPERS

## Measuring the Internet's Vital Statistics

### A Special Section of ACM SIGCOMM COMPUTER COMMUNICATION REVIEW

There are a number of measurements of the Internet's basic properties that are tremendously useful for researchers to know. A solid understanding of key Internet properties is useful for:

- creating quality models to evaluate innovative new protocols, algorithms and architectures
- understanding the overall context with which a new system will inevitably have to cope if/when deployed
- gaining an appreciation for the variety and breadth of situations one may encounter when measuring the Internet

Yet, as the Internet has grown, many of these measurements are no longer widely circulated -- often because they are viewed as operational rather than of research interest. Assumptions based on dated information about the operational Internet are used every day by researchers. Unreliable assumptions about the Internet's key properties may lead to everything from wasted time in appreciating the breadth of scenarios one must take into account in measurement analysis to simulation studies that are not of practical import because of the inaccurate models.

The goal of this special section is to begin the practice of periodically publishing measurement studies of the Internet that concisely provide reliable, easily accessible information for researchers to use and build on. Our aim is to complement traditional measurement venues that emphasize new measurement techniques and evaluations of new protocols and architectures, by updating the community's working knowledge of the basic properties.

Examples of measurements that would be of interest are:

- a breakdown of the types of bad/defective/broken DNS queries received at the typical root server;
- how frequently the TCP/UDP checksum and the MAC-layer CRC disagree;
- the distribution of traffic among different applications (perhaps, measured at different points in the network);
- the ratio of local traffic to global traffic;
- how often the forward and return paths of a TCP connection differ;
- how often traffic is reordered;
- the variation in BGP route prefixes advertised over the course of a typical minute, hour, day, week, and month;
- distributions of round-trip times experienced in the network

Our expectation is that each paper will be short: a description of the methodology by which the data was captured and where it was gathered, a presentation of results, and where possible, a comparison of the current results with those of prior years (and other researchers). The focus of the papers should be on the data, rather than on developing new methodologies. We encourage authors to collaborate to illustrate information from multiple vantage points in the network. In addition, we especially solicit papers that are coupled to public release of measurement datasets (even if anonymized in some form).

Our aim is to initiate a cycle whereby the community constantly updates our collective understanding of the Internet's basic properties. Therefore, following the publication of this special section, CCR and Sigcomm will endeavor to continuously publish papers on the Internet's basic properties to keep the community's perspective fresh.

#### **Submission Information:**

Please submit papers to the special section by email to [ccr-ivs@lists.csail.mit.edu](mailto:ccr-ivs@lists.csail.mit.edu). For further information and submission details please see <http://www.acm.org/sigcomm/ccr/ivs>

Please address questions about content or submission procedure to [ccr-ivs@lists.csail.mit.edu](mailto:ccr-ivs@lists.csail.mit.edu)

#### **Schedule:**

Submissions due November 1, 2004  
Acceptance decisions December 3, 2004  
Publication January, 2005

#### **Special Section Editors:**

Mark Allman, ICIR  
Craig Partridge, BBN