Network Kriging

Eric D. Kolaczyk
Department of Mathematics and Statistics
Boston University
111 Cummington Street
Boston, MA 02215
kolaczyk@math.bu.edu

ABSTRACT

It is often desirable to monitor end-to-end properties, such as loss rates or packet delays, across an entire computer network. However, active end-to-end measurement in such settings does not scale well, and so complete network-wide measurement quickly becomes infeasible. More efficient measurement strategies are therefore needed. Previous work, examining this problem from a linear algebraic perspective, has shown that for exact recovery of complete end-to-end network properties, the number of paths that need to be monitored can be reduced to approximately the number of links in the network. Here we argue that in fact measurement strategies of even greater efficiency are possible.

We begin by recasting the problem as one of statistical prediction and show that end-to-end network properties may be accurately predicted in many cases using a significantly smaller set of carefully chosen paths than needed for exact recovery. We formulate a general framework for the prediction problem – which we call 'network kriging', propose a simple class of predictors for standard quantities of interest (e.g., averages, totals, differences), and show that linear algebraic methods of subset selection may be used to make effective choice of which paths to measure. We explore the accuracy of the resulting methods both analytically and numerically, in the context of real network topologies of varying size. The feasibility of our methods derives from the low effective rank of routing matrices as encountered in practice, which appears to be a new observation of interest in its own right.

Keywords

Internet traffic, network monitoring, statistical prediction.

NetSci2006, May 22–25, 2006, Bloomington, IN, USA.
.