New Developments and Applications in Experimental Design
IMS Lecture Notes – Monograph Series (1998) Volume 34

OPTIMAL MONITORING OF COMPUTER NETWORKS

VALERII FEDOROV[1] AND DEBORAH FLANAGAN[1]

Oak Ridge National Laboratory

We apply the ideas from optimal design theory to the very specific area of monitoring large computer networks. The behavior of these networks is so complex and uncertain that it is quite natural to use the statistical methods of experimental design which were originated in such areas as biology, behavioral sciences and agriculture, where the random character of phenomena is a crucial component and systems are too complicated to be described by some sophisticated deterministic models. We want to emphasize that only the first steps have been completed, and relatively simple underlying concepts about network functions have been used. Our immediate goal is to initiate studies focused on developing efficient experimental design techniques which can be used by practitioners working with large networks operating and evolving in a random environment.

1. Introduction. In most cases a computer network can be represented as a graph with a given number of nodes (vertices, sites) and edges (links, communication channels). Possible objectives of experiment(s) may include evaluating such quantities as the following: delays on a given subset (subset of interest) of edges, processing times at a given subset of nodes, traveling times from one subset of nodes to another, etc. Existing software and hardware allow measuring [see, for instance, Paxson (1997) for details and a comprehensive list of references] a large variety of network performance indicators, so that in general our "measurement" is a vector. Types of measurement strategies may be very different. For instance, a meter can be installed at any chosen node to measure input and output flows; a measurement software or hardware device can be placed at a host node, and a preselected set of nodes or edges can be monitored; a practitioner can cooperate with others (i.e., there are a few host nodes) to monitor a network. Thus, if we have an opportunity to plan (design) experiments, we may look for the best subset of host nodes where devices must be allocated, find the most informative subset of nodes and edges to be monitored by the given host, or in the most general setting select the most effective team of host nodes (sites) and match them with

---

[1] Research supported by the Oak Ridge National Laboratory Directed Research and Development Fund.

Received September 1997; revised March 1998.

AMS 1991 subject classifications. 62K05.

Key words and phrases. Correlated observations, design of experiments.