Simulation is Crucial

Sally Floyd
AT&T Center for Internet Research at ICSI

January 2001
To appear as a sidebar in IEEE Spectrum, January 2001

Focusing on three key questions may perhaps best help to explain the difficulties involved in modeling and simulating the Internet to make the changes needed for its expansion. First, why are simulations necessary for developing and evaluating proposed additions to the Internet architecture?

Research on proposed changes to the Internet is ongoing. The Internet is not a static entity, but is undergoing constant change, with the development of new link-level technologies such as wireless and satellite links, new infrastructures involving Web caching and content distribution networks, and new applications such as Napster and streaming multimedia.

In addition, constant development of existing Internet protocols-like transmission control protocol (TCP), the Internet routing protocol (IP), and the domain name service (DNS)-are necessary to address these new developments in the infrastructure, to meet new threats from denial-of-service attacks, or simply to meet the challenges of scale as the Internet expands in size and functionality.

Some research on these proposed changes can be done with analysis, experiments in testbeds, or small-scale deployment in the current Internet. But simulations are essential as a tool for exploring proposals in diverse environments or in large-scale topologies. They are also needed to explore proposals in environments that have not yet been realized in the Internet, but that might be in the future.

A second question is: how much do we understand about modeling the jnternet for such research?

In some simulations-say, for exploring the ramifications of adding router A or link B to a particular topology-all that is needed is to model a specific topology, with the protocols and traffic mix characteristic of that network. Here discrete-event, analytic or hybrid simulations addressed in the main article may prove useful.

Still, for exploring a protocol to be deployed in the global Internet, there is no such thing as a typical topology or traffic mix. The global Internet is characterized by heterogeneity and change, not only in terms of topology and traffic, but also in terms of router scheduling and queue management mechanisms, transport protocols, Web-caching infrastructures, applications, and other areas.

This means that claims of a typical topology with a typical traffic mix and a typical packet drop rate are pitfalls that are better avoided, to be replaced by a goal of reflecting the heterogeneity of the global Internet.

Fortunately, several invariant properties of the Internet can be relied upon in modeling, along with a growing set of tools, such as topology generators and traffic generators. For example, traffic generators can be used to generate a traffic load characteristic of the "mice" and "elephants" of current Internet traffic, where the mice are small transfers (say, a simple Web page) and the elephants are the large ones (say, downloading a video file).

One invariant of Web traffic is that, while most of the transfers are mice, the elephants are sufficiently large that the bulk of the packets in the Internet are from the elephants. Other invariant properties that can be used in modeling the Internet include the 24-hour fluctuations characteristic of Internet traffic as well as the modeling of session arrivals from individual users as independent events.

The third key question is: which modeling issues are inherently problematic?

Obviously, no one can predict the future. Thus, we don't know whether to explore our proposed changes in an environment with or without the newest Internet protocol, version 6 (IPv6), or with or without explicit congestion notification (ECN), or with or without IP multicast traffic, or even with or without Napster. The best that we can do is to make educated guesses about other potential changes in the Internet that might interact with our own proposed change, and to investigate the potential interactions.

The other problematic issue is that of traffic dynamics and protocol interactions related to issues of scale. Clearly we can't run experiments on testbeds the size of the global Internet. Using current simulation technology, we can't simulate networks of that size, either. And, even if we could scale, we would not have the proper tools to interpret the results effectively. Addressing issues of scale remains one of the simulation and modeling issues yet to be tackled.