Chapter 1

Introduction

As the Internet grows larger, measuring and characterizing its dynamics grows harder. Part of the difficulty is how quickly the network changes. Depending on the figure of interest, the network between 80% and 100% each year, and has sustained this growth for well over a decade. Furthermore, the dominant protocols and their patterns of use can change radically over just a few years, or even a few months [Pa94b, CBP94].

Another difficulty, though, is the network's incredible—and increasing—heterogeneity. It is more and more difficult to measure a plausibly representative cross-section of its behavior. It is this latter concern that we attempt to address in this work. In this chapter, we develop the context for the rest of the study, by discussing different types of traffic studies, and how research efforts of those types have addressed heterogeneity problems. Our study falls in the category with perhaps the greatest heterogeneity difficulties, that of the “end-to-end” performance of entire paths through the network.

Our work has two distinct parts: a study of end-to-end routing behavior in the Internet (Part I), and a study of end-to-end Internet packet dynamics (Part II). These two are united by the common measurement framework used to gather the data analyzed in each part (described in Chapter 4). In addition, some of the results used in each part are incorporated into analysis in the other part. However, in many ways the two parts are distinct and self-contained. A reader particularly interested in one or the other topic might profitably just read the relevant part. Consequently, we defer an overview of each part to later chapters (Chapter 2 of Part I, and Chapter 9 of Part II). We summarize both parts, and what we perceive as the themes of the work, in Chapter 17, at the end of Part II. For the remainder of this introduction, we give an overview of the general problem of measuring large networks,

We can classify measurement studies into several basic types. Each faces the problem of heterogeneity to varying degrees, as follows.

Exhaustive studies analyze properties of a significant fraction of the entire network. Examples are Kleinrock's study of the ARPANET's behavior on time scales of hours to days [Kl76]; the series of “ping” experiments conducted by Mills to evaluate the effectiveness of the TCP retransmission-timeout algorithm [Mi83]; Claffy et al.'s study characterizing traffic on the T1 NSFNET backbone [CPB93b]; and Chinoy's study of the dynamics of routing information within the NSFNET backbone [Ch93]. While these studies can convincingly characterize the full range of behavior one might expect to observe from the network, they become impractical as the network grows in size.[1]

Site studies characterize the aggregate traffic patterns observed for entire sites. They focus on the connections sizes, durations, and interarrival times. An early site study, by Danzig and colleagues, identified large heterogeneities in the traffic “mix” at each site, meaning that the proportion of total traffic (total connections, total packets, or total bytes) due to different applications varies greatly [DJCME92]. Our subsequent work extended this finding to the characteristics of the connections made by each type of application. We found that the distributions of a particular application's connection sizes and durations varied greatly from site to site [Pa94a], in agreement with much earlier findings, in a different communications context, by Fuchs and Jackson [FJ70].

Another type of study focuses on *server* behavior, for services that are distributed over the Internet. The heterogeneity issues faced by these studies vary greatly, depending on the service. For example, Danzig and colleagues analyzed requests arriving at a “root” name server, finding a variety of performance problems [DOK92]. However, there are only a handful of root name servers. Because clients divide their requests between them, studying a single server yields results plausibly representative for all of the servers. On the other hand, a recent study of World Wide Web servers had to grapple with the issue of differences among the various servers studied [AW96], and did so by developing its central theme around the search for behavioral “invariants” among the six Web servers analyzed.

Related to server studies are *client* studies, which analyze the different ways in which clients access servers. From a heterogeneity standpoint, client studies are more difficult than server studies, since usually there are many more clients than servers. One approach is to study the behavior of all clients located at a particular site [CB96]. Doing so, however, incurs problems similar to those of site studies: it is difficult to gauge the generality of the findings. These problems, however, can be tempered based on the nature of the service. For example, we might expect request from one site's Web clients to more closely resemble those of another site's Web clients, than for the site's aggregate traffic to resemble that of another site.

Another type of study analyzes the aggregate traffic seen on network *links*. These studies have focussed on the dynamics of packet arrivals on the link [FL91, LTWW94, PF95, WTSW95], the characteristics of packet “flows” [JR86, He90, CBP95], or on traffic patterns over particularly singular links, such as the trans-Atlantic link connecting the U.S. and the U.K. [CW91, WLC92].

For link studies of *local area* networks [JR86, FL91, LTWW94, WTSW95], heterogeneity presents less of a problem than for those of *wide area* networks, because the latter encompass a much broader range of traffic sources and path characteristics than the former. Some wide area link studies attempt to address heterogeneity issues by analyzing traces from multiple links. However, gathering link traces is difficult (and becoming more so due to security, privacy, and business concerns), and such studies have not to date analyzed more than two dozen or so traces.[2]

Our work falls in still another class, that of *end-to-end* studies. These studies concern

[1] At the time of the later of the first two studies, the Internet comprised about 600 hosts. As of this writing, it comprises about 16 million hosts [Lo97]. At the time of the Claffy study, the backbone consisted of 15 nodes and two dozen links. Today, it is much larger, though sources of accurate statistics on its size have virtually disappeared with the commercialization of the Internet infrastructure.

[2] Most site studies are conducted as link studies, too, since an extremely convenient way to capture an entire site's Internet traffic is to monitor what is usually a single link connecting the site to the rest of the Internet. Some server studies can also be conducted in the context of a link study, by analyzing all of the server requests seen on a highly aggregated link [EHS92, DHS93].

how the network performs from the perspective of an end user. To users, a network like the Internet is simply a black box that somehow forwards packets between their host and hosts with which they wish to communicate. End-to-end studies face extreme heterogeneity problems because they strive to characterize the behavior of *paths* through the network. Not only does the Internet contain millions of distinct paths, but the dynamics of each path reflects the concatenation of the dynamics of each forwarding element along the path, and hence can be highly complex.

The few studies to date of end-to-end packet dynamics—Mogul's look at TCP dynamics such as ack compression [Mo92], Bolot's analysis of patterns of packet loss and delay [Bo93], and Claffy et al.'s characterization of one-way latencies [CPB93a]—have all been confined to measuring a handful of Internet paths, because of the great logistical difficulties, *and also analysis difficulties*, presented by large-scale measurement.[3] Consequently, it is hard to gauge how representative these end-to-end findings are for today's Internet.

As a result, even basic Internet path questions such as "how often do routes change?" and "how often are packets dropped?" remain unanswered in any sort of general way. It is towards answering these questions that we now embark.

[3] Mogul's study was actually conducted as a link study. Doing so let him observe behavior from a fairly large number of Internet paths, albeit ones that all had the single link in common. A drawback of this approach, however, is that it is difficult to infer from the perspective of a link the full end-to-end behavior as perceived by the endpoints, an issue we discuss further in § 10.4.