

INDEX: A Platform for Determining how People Value the Quality of their Internet Access

INDEX Project Report #98-010P

Björn Rupp Richard J. Edell Harish Chand Pravin P. Varaiya

May 1998

© 1998 IEEE. All rights reserved. *Do not redistribute this document electronically.* This document may be found on the web at:
<http://www.INDEX.Berkeley.EDU/reports/98-010P>

Abstract

The continuing exponential growth of the Internet and the emergence of new time-critical applications have led to the integration of a large number of different services on the Internet. In the process, the question of how to efficiently allocate bandwidth as a scarce resource has become a crucial issue for the continued proliferation of these new services. Future growth depends on the division of services into quality-differentiated market segments and the pricing structure of each segment. Successful growth requires service providers to offer combinations of quality and price that match user need. But to do this providers must understand the structure of user demand. Such understanding is lacking at present.

This paper describes a platform designed to obtain a basic understanding of how individuals value Internet usage when offered different Quality of Service choices. The Internet Demand Experiment (INDEX) has two main objectives: (a) Measurement of user demand for Internet access as a function of Quality of Service (QoS), pricing structure, and application; and (b) Demonstration of an end-to-end system that provides access to a diverse group of users at attractive price-quality combinations. The data being collected is expected to reveal the correlation between user application and service demand, how demand varies with user experience, and up to what extent users form discrete market segments. This paper gives an overview of both the technology employed at INDEX and the goals of the experimental design.

The INDEX Project is supported by the National Science Foundation, Cisco Systems, SBC Communications and the California State MICRO Grant program.

INDEX: A Platform for Determining how People Value the Quality of their Internet Access

Björn Rupp* Richard Edell* Harish Chand† Pravin Varaiya*

*Department of Electrical Engineering & Computer Sciences †Department of Economics
University of California at Berkeley

Abstract

The continuing exponential growth of the Internet and the emergence of new time-critical applications have led to the integration of a large number of different services on the Internet. In the process, the question of how to efficiently allocate bandwidth as a scarce resource has become a crucial issue for the continued proliferation of these new services. Future growth depends on the division of services into quality-differentiated market segments and the pricing structure of each segment. Successful growth requires service providers to offer combinations of quality and price that match user need. But to do this providers must understand the structure of user demand. Such understanding is lacking at present.

This paper describes a platform designed to obtain a basic understanding of how individuals value Internet usage when offered different Quality of Service choices. The Internet Demand Experiment (INDEX) has two main objectives: (a) Measurement of user demand for Internet access as a function of Quality of Service (QoS), pricing structure, and application; and (b) Demonstration of an end-to-end system that provides access to a diverse group of users at attractive price-quality combinations. The data being collected is expected to reveal the correlation between user application and service demand, how demand varies with user experience, and up to what extent users form discrete market segments. This paper gives an overview of both the technology employed at INDEX and the goals of the experimental design.

1 Motivation

In recent years, the Internet has undergone a dramatic transformation from a computer network dominated by traditional, mostly text-based applications and a comparatively small, coherent user community to a universal platform for ever more users and services. This was not without its consequences. While traditional applications like electronic mail or file transfers can react in an elastic fashion to deviations in available bandwidth, new time-critical applications like Internet telephony and video conferencing cannot, thereby causing their employment to be severely limited as soon as network congestion leads to high packet delays and packet drops. With the explosion of demand for Internet services, higher speed access, and new applications, this situation continues to worsen. A single “best effort” service quality seems to become increasingly inappropriate for a network serving a wide variety of users and applications. Currently, users who occasionally need high bandwidth are either forced to lease over-provisioned dedicated lines,

risk the vagaries of the performance of “best effort”-quality shared resources, or forego the desired application altogether. When demand for Internet access varies among the population (as indicated by population-projectable data as in [CommerceNet/Nielsen 1997]), quality differentiation, along with proper economic incentives, can increase the overall value of the network by making available resources when needed for high value applications. The division of services into quality-differentiated market segments and the design of appropriate pricing structures for each segment is crucial for further proliferation of Internet services. Successful growth requires service providers to offer combinations of quality and price that match user need. But to do this providers must understand the structure of user demand. While there have been many pricing proposals in recent literature (for a short overview of different approaches, see [Shenker et al. 1996]), such understanding of user demand is lacking at present.

INDEX — the Internet Demand Experiment — is a real-world market trial seeking to provide this information and measure how individuals value Internet usage when they are offered different Quality of Service choices. INDEX has two main objectives: (a) Measurement of user demand for Internet access as a function of quality of service (QoS), pricing structure, and application; and (b) Demonstration of an end-to-end system that provides access to a diverse group of users at attractive price-quality combinations. The experiment will provide Internet access over ISDN lines to a group of about 150 users from the Berkeley campus community for a two-year period. Users select network services from a menu of QoS-price offerings and pay for their usage. It is important to stress that while the subjects’ basic Internet access (in particular, the ISDN line and access equipment) is greatly subsidized, each choice on these QoS menus has a real economic cost which the subjects pay out of their own pockets. This is necessary in order to achieve incentive compatibility, i.e. given the incentive schedule as represented by their active menu, users pick the option that corresponds to their true valuation of the network resources in question. The menu changes in certain intervals in order to measure demand for a wide range of combinations of QoS, price and user characteristics. The data being collected is expected to reveal the correlation between user application and service demand, how demand varies with user experience, and up to what extent users form discrete market segments. The data will also allow to test hypotheses about the structure of the market for variable-quality ATM services. In addition, the experiment demonstrates a single system that offers variable service quality-price combinations that meet the needs of a diverse user population, an automated billing system that also gives the user control over service selection, and

a remotely operated network monitoring and management system.

This paper gives an overview of the INDEX Project’s scope and describes both the technology employed and the goals, timing and structural details of the experimental design.

2 Experimental Setup

2.1 INDEX Access Network Provision

The INDEX access network provides IP service over dedicated, 128kbps ISDN lines in order to establish a predictable and stable QoS between the subjects’ homes and the INDEX Project Network Operations Center¹. For this purpose, INDEX loans a pre-configured Cisco 762 ISDN router to each subject participating in the experiment and installs an ISDN phone line at their home. The 128kbps basic rate interface lines coming from the subjects’ homes are then multiplexed over ISDN primary rate lines at the Pacific Bell central office before they reach the INDEX Project Network Operations Center. In contrast to common industry practice, the overall available bandwidth is not reduced in the multiplexing process and the whole network is heavily overprovisioned to make sure that none of the subjects experience deteriorations of their selected quality level due to potential bottlenecks at the INDEX access network.

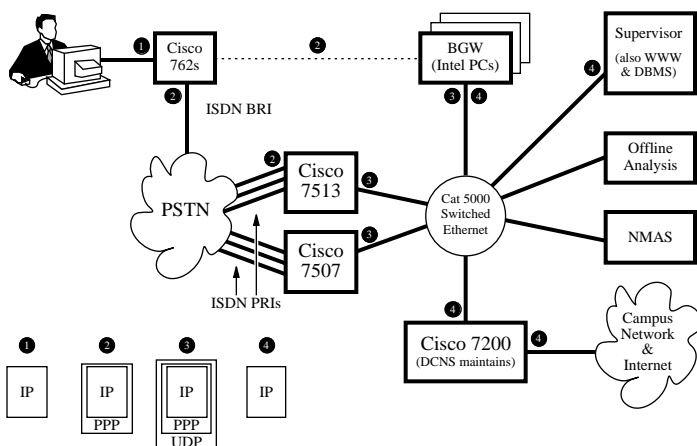


Figure 1: INDEX Network – Transport Layer

At the INDEX NOC, all connections are through either a Cisco 7507 or 7513 Internet router. These routers distribute all user traffic over a set of Billing Gateways specifically designed to meter usage and selectively adjust the service quality of individual connections. The user may select a service quality from the currently active menu of choices at any time. Connections are aggregated by user so that the quality for this bundle can then be controlled accordingly. All outbound packets are forwarded to a Cisco 7200 router that is directly connected to the UC Berkeley 100Mbps FDDI backbone.

¹It should be noted that although the current experimental setup is oriented towards providing service over ISDN lines, the INDEX network architecture is flexible enough to allow us to expand the experiment to demonstrate ADSL or CATV access using cable modems at a later stage.

2.2 User Interaction, Accounting and Billing

INDEX uses a locally developed system for user interaction and metering individual subject usage. The user interacts with this system by means of the “Control Center”, a Java application running on the user’s computer. For the subjects, this is the central application enabling them to select different Qualities of Service and control their usage of network resources. Apart from functions for login and authentication, it consists of a small window informing the user about the current experiment, the price schedule currently in effect and the actual choices. The subjects can choose a service quality by the click of a button and change their Quality of Service even during the active session. The Control Center also provides usage feedback by displaying a summary of charges for either the current session, the current day or the current month.




Figure 2: INDEX User Interface (“Control Center”)

The Control Center application communicates user choices and selected quality levels as control data going through a Billing Gateway to a “supervisor” process. This supervisor process then orders the Billing Gateway to treat this user’s connections according to the selected quality level. The Billing Gateway in turn meters the traffic and reports back to the supervisor process.

User traffic is monitored and recorded at a fairly detailed level for both billing purposes and subsequent offline analysis. The database contains records for each TCP connection. Apart from an anonymized user ID, time stamp, selected QoS/price information and a variety of TCP control data types, they include information about connection length, the amount of inbound and outbound traffic for the connection, source and destination IP addresses, port numbers, and other data describing the type of user activity. It is important to collect data at this level of detail in order to not only record at what time users change their QoS choices, but also to infer what parameters influence these decisions and what the reasons for these changes are. Such detailed records are able to reveal, for instance, what applications are running at the time of a QoS change and what types of hosts and network services are involved.

2.3 Network QoS Emulation

After a subject has chosen a desired quality level, the QoS must be adjusted (i.e. degraded) accordingly. As the current Internet infrastructure does not permit controlling QoS, the INDEX Billing Gateways do not only control and measure network usage, but they are also capable of selectively degrading the performance of certain TCP connections (e.g., all connections on behalf of a given subject). User quality choices map to entry points of an internal “emulated network” composed of different elements including leaky bucket, random router and packet delay or packet drop. This behavior can be altered quickly in response to subject choices or experimentally controlled random processes.

While it would be desirable to give users complete control over end-to-end performance of their connections, this is unfortunately impossible while the single-quality “best effort” paradigm still prevails. As a consequence, the quality offered by INDEX cannot exceed the baseline undegraded QoS. However, much of our subject pool is accustomed to a congested, 14.4kbps modem pool for accessing the campus network. In addition, many of the services that our subjects seek to access are in fact campus services for which there are little other sources of degradation. Therefore, we believe that INDEX is adequately capable of controlling the QoS delivered to the subjects.

3 Experimental Design

3.1 Purpose and Objectives

INDEX seeks to answer the question of how people value the quality of their Internet access. There are three basic aspects to this question. The first and most fundamental is what dimensions of QoS truly matter for overall end user’s perception of service quality. While there are indeed many dimensions of QoS that can or could be varied on the supply side, little is known about along which dimensions users feel performance to be most significantly affected by these parameters and up to what extent they are willing to sacrifice certain service characteristics in order to improve others. Although these internal valuations are very likely to be heavily dependent on the type of application the user is running, empirical data capable of accurately verifying and quantifying these hypotheses is still missing at present. One important goal of the experimental design therefore is to *identify the key parameters for user’s perceptions of QoS* and to quantify the correlation between application type and service demand. This information will substantially aid network service providers in future decisions on which aspects of QoS to optimize and when and where to expand the network in order to provide service options that satisfy users’ needs.

The second task is to *measure the economic value which individual users place on different resource levels for each of the QoS dimensions identified*. The results from these investigations should support decisions on what kind of pricing structure to choose and what economic incentives the price structure should provide. For example, congestion-based prices require users and network administrators to know the value lost due to congestion. The degree to which congestion needs to be discouraged or the markup charged

to priority service depends on the degradation of user value due to user contention. If the perceived network degradation from congestion is substantial, then congestion-related pricing (such as time-of-day, traffic-based, or priority-based charges) can rationalize the allocation of network resources and increase the value of the network.

Apart from gaining a better understanding of how to quantify valuation of different dimensions of QoS, survey data (i.a. [CommerceNet/Nielsen 1997, Kehoe/Pitkow/Morton 1997]) also suggests that demand for Internet access varies among the population. The data also indicates that there is a correlation between user experience and intensity of network usage. If users are indeed considerably heterogeneous in their consumption of network resources, information about the exact nature of the price elasticity of demand will help to differentiate types and levels of service through pricing. Therefore, INDEX also includes experiments involving nonlinear tariffs to *determine up to what extent users form discrete market segments*. Such multi-part tariffs involve price discrimination in the sense that different bundles of homogeneous output are sold at different prices. If substantial variation exists among users, and if users are sufficiently sophisticated in their decisions, then self-selecting tariffs differentiated on the basis of quantity and quality may segment the market on the basis of willingness to pay.

3.2 Subject Population

INDEX will recruit about 150 subjects affiliated with the University of California at Berkeley (students, faculty, staff). Participation in the Project is highly attractive for this group of experimental subjects because the campus modem pools are highly congested while there exists at the same time a lack of other options offering service similar to INDEX. The availability of this subject pool offers several advantages: Firstly, it provides the required geographically concentrated pool of diverse users for such an experiment. It also allows for a better QoS control than otherwise possible, ensures continued participation and reduces costs of setting up the experiment.

The recruiting process involves several steps: The project advertises by electronic means such as posting to newsgroups as well as traditional means of Berkeley newspaper advertisements and articles. The advertisements direct interested people to the INDEX WWW server to learn more about the project. Potential subjects who wish to participate are required to complete an on-line screening survey. The INDEX screening survey collects basic contact information, residence location and nature of university affiliation to verify eligibility and aid in ISDN service planning. Prospective subjects must reasonably expect that their affiliation with UC Berkeley will continue for at least two years. We also plan to facilitate participation from individuals affiliated with UC San Francisco in the near future. After prospective subjects have completed the screening survey, they are invited to complete an extensive demographic survey. At this stage, taking the survey is optional, however completing it is required for all participating subjects before they take part in the experiment. After evaluating the screening survey results, we invite selected persons to participate. If they agree, they receive their access equipment and

can begin using their INDEX-provided Internet access after the ISDN line has been installed at their home and they have signed an informed consent form.

Subjects are recruited with the goal of obtaining a suitable variation in, e.g., field of study, expected computer usage, travel distance to campus, and demographic characteristics. Nevertheless, it is evident that the sample, like any other, remains biased. In order to overcome this sampling bias and be able to extrapolate the results, significant sample-specific characteristics influencing the results will be identified and analyzed to determine the exact structure of the INDEX Project's demographic base. Before receiving INDEX provided Internet access, each subject has to complete a detailed demographic survey. Many of its questions are taken from a representative, population-projectable study conducted by Nielsen Media Research [CommerceNet/Nielsen 1997] in early 1997. Apart from general demographical data (i.e. income, age, gender, household characteristics etc.), the INDEX Demographic Survey asks subjects about their recent Internet usage, computer sophistication and related data. By comparing the responses to the INDEX Demographic Survey with the Nielsen data, it is possible to extrapolate findings from the INDEX Project to the general U.S. population².

3.3 Revealed Preference Experiments

To achieve the objectives outlined above and investigate the effects of alternative price and QoS combinations, a series of experiments will be conducted. These experiments infer preferences from data based on actual choices for which the subjects have to make economic decisions immediately affecting them – therefore providing reliable incentives to accurately represent their preferences. The experiments are preceded by a free trial period of several weeks. This serves as a “control” for later INDEX experiments and allows the subjects to gain familiarity with their new Internet access, the billing interface and experimental procedures. After the trial period is over, users begin paying for their usage. Depending on the experimental setting, the fee structure will change either every Monday on a weekly basis or daily at 4 a.m. Pacific Standard Time. Higher qualities of service will incur a higher fee, but the level and ratios depend on the specific sub-experiment and will also change over the course of each experiment. Whereas nonexperimental studies are forced to rely on cross-sectional variation in prices and demand to infer the price elasticity, varying the prices during the individual sub-experiments allows for measuring the demand response for each participant. Fairly detailed data will be collected on the characteristics of the subjects' usage. As described in section 2.2, each time the subject uses the ISDN connection, data will be collected on the time and length of the session, the speed of the connection, the price in effect, and the amount of data transferred and applications used during the session.

The experiments can be partitioned into two groups. The first group, representing the first four experiments, examines how the

²We also plan to expand the experiment to conduct general population experiments and test CATV and ADSL access in the future. This is dependent upon cooperative arrangements with CATV and ADSL service providers which have not yet been negotiated. These future experiments will be based on a different ISP Resource Model and are not described in this paper.




Figure 3: Traditional ISP Resource Model

dimension of resources within the traditional ISP resource model (as depicted in Figure 3) affect end user valuation. The second set of experiments, experiments five through nine, examine users' response to alternative price structures. In the majority of these later experiments, users will compare alternative pricing schemes with a flat rate scheme. The following sections describe the individual sub-experiments in detail.

3.3.1 Network Resource Valuation Experiments

(1) – Variable Bandwidth. Do individuals value connection speed sufficiently to pay higher prices for high speed connections? How does the elasticity of demand depend on application and demographics? Does demand exhibit habit formation? This experiment will address these questions and examine how users value the speed of their connection to the Internet cloud. It isolates the last link in Figure 3 and permits the accurate measurement of the price elasticity of demand for connection speed to the Internet cloud.

Design Details. QoS dimension controlled: Bandwidth (six connection speeds $A = \{8, 16, 32, 64, 96, 128 \text{ kbps}\}$). Duration: 6 weeks. Price structure: Five weeks with weekly price changes, one week with daily price changes – strictly increasing prices within certain limits.

(2) – Variable Asymmetric Bandwidth. Since the early days of the 1200/75 bps modems, the question of whether individual end users value bandwidth for incoming traffic more than for outgoing traffic has been discussed intensely. With the advent of ADSL and CATV proposals that feature different fixed data rates for incoming and outgoing traffic, this issue warrants further research. Do individuals value the speed of their connection *from* the Internet cloud differently than the connection speed *to* the Internet cloud? If yes, what ratios of incoming vs. outgoing bandwidth are deemed appropriate? Up to what extent do these ratios depend on the type of application run by the user? This experiment seeks to answer these questions.

Design Details. QoS dimension controlled: Bandwidth (six connection speeds for both incoming and outgoing traffic). Duration: 6 weeks. Price structure: Five weeks with weekly price changes, one week with daily price changes – prices from the first experiment will be cut in half and applied separately towards each direction of all data flows.

(3) – Access Reliability. Congestion in dialup access to the network is a significant problem in many networks. The expected waiting time for a free line can be significant. As a result, once

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.