An Economically Scalable Internet

Abdelsalam "Solom" Heddaya
InfoLibria

Over the past decade, many observers have claimed that the Internet brings the information revolution's components together in a way that will rival the industrial revolution's effects on human productivity and quality of life. Replication is central to this comparison: The industrial revolution enabled large-scale replication of increasingly complex physical objects, and the information revolution—primed by the Internet—fulfills a similar function for information objects. Yet this comparison highlights a key Internet deficiency: Economically, it does not scale well.

In manufacturing technology, the per-unit cost decreases as the total production rate increases. Successful products create a virtuous cycle—the more popular they get, the cheaper and more profitable they become. This cycle emerges from an industrial architecture that becomes more efficient in per-unit resource costs as the resources required for production increase. Doubling production resources more than doubles production capacity.

Transportation and other delivery systems evolved to match this economy of scale as these production efficiencies made their way to the marketplace. Success in this environment rewarded both the producer with higher profits and the consumer with lower prices.

BANDWIDTH RESOURCES

But the Internet's architecture does not scale economically this way. The recurring costs of Internet bandwidth—not the capital costs of hardware and software—tend to dominate the costs of Web hosting. Roughly speaking, a Web site that requires 500 Mbps costs 10 times more than one requiring 50 Mbps. In this business environment, a site's popularity does not translate into producer profits and consumer savings.

Thus, from 1980 to 2000, Internet backbone traffic grew by a factor of 100 million, but the bandwidth cost for serving a single streaming advertisement in 2000 remained higher than the revenue it generated.

Similarly, the bandwidth cost of delivering one hour of a 1-Mbps video stream was roughly US$5—far more than service providers could expect consumers to pay, even though much cheaper than long-distance telephone network bandwidth. More significantly, the delivery costs applied almost equally to popular and unpopular content.

When investors saw the Internet's popularity soar among the general population in the 1990s, they wanted to invest in resources that they expected would make the Internet incredibly cheap—at least for popular content and services. Examples of such investments include the deployment of huge data centers and massive long-distance fiber capacity. There is little evidence, however, that these investments realized fundamental economies of scale, even for successful Internet services.

QoE VERSUS QoS

Nor did the end users' *quality of experience* reach a high enough or consistent enough level to convince them to pay for new services. Internet QoE remains low and unreliable. I use the term QoE rather than quality of service because QoS is not necessary for QoE, and QoE is sufficient for successful service. For example, downloading a two-hour movie to a home entertainment system or PC for later viewing can yield high QoE for the movie-viewing experience, even if the available bandwidth is of low quality in terms of packet loss, packet delay, and bandwidth variation.

The Internet has proven its system scalability but has yet to demonstrate significant economic scalability.

To complicate matters, the Internet's commercialization and rapid growth during the 1990s spawned many scaling, quality, and security solutions that ran afoul of some dearly held architectural principles of the Internet community. The Internet Engineering Task Force—the key standards-setting body entrusted with maintaining the network's integrity—became the forum for many heated debates.

The architectural debate is best captured by Marjory S. Blumenthal and David D. Clark in "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World" (*ACM Trans. Internet Technology*, vol. 1, no. 1, Aug. 2001, pp. 70-109). Regardless of the merits of

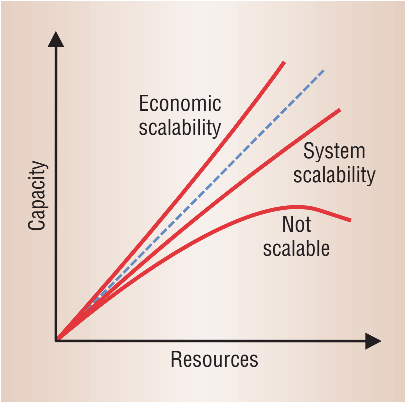

Figure 1. Economic scalability must improve the performance-price ratio superlinearly.

particular proposals or good intentions of all parties, the intensity of these debates may well have detracted from the community's ability to focus on fundamental issues.

SCALABILITIES

We can describe economic scalability by measuring system efficiency and capacity as a function of required resources. Efficiency represents the amount of resources needed to deliver a unit of service—such as a video stream or e-mail message—or to complete an e-commerce transaction. Capacity represents the maximum rate of service that a system can handle. One architecture has better economic scalability than another if its efficiency and capacity functions grow faster.

As Figure 1 shows, we can use the same metrics to capture the related concept of *system scalability*. A system is considered scalable if capacity continues to grow, even if slowly, as more resources are put into the system. A system is *not* scalable if there is a point at which adding resources yields no increase in capacity.

The difference between system and economic scalability lies in expectations. Computer scientists and engineers have defined ideal system scalability to be constant efficiency and a linear rate of change in capacity, illustrated by the dotted line in Figure 1. An ideally scalable system would grow in capacity in direct proportion to the amount of resources dedicated to it, but no more. Even though certain systems, in practice, exhibit greater scalability, the professional community has generally discounted such cases on the assumption that they did not properly account for all resources.

Vipin Kumar and others proposed one of the best metrics for system scalability, the isoefficiency function (*IEEE Parallel and Distributed Technology*, vol. 1, no. 3, Aug. 1993, pp. 12-21). This metric characterizes the increase in raw workload a system requires to maintain constant efficiency as more resources are added to it. One system is more scalable than another if it has a slower-growing isoefficiency function.

While the isoefficiency framework may be adaptable to the analysis of economic scalability, its name suggests that ideal scalability keeps efficiency constant, which makes economic scalability impossible. A simple computing resources model—processors with fixed computing rates—actually presupposes this impossibility. Thus, we need a richer model to capture the full range of resources contributing to the cost of an Internet-delivered service unit—from cache storage to cable rights-of-way. The model could then support development of an Internet architecture that exhibits economic scalability, not just system scalability.

SUPERCOMPUTING LESSONS

In the late 1980s and early 1990s, the supercomputing industry promised, but failed to deliver, economy of scale. Supercomputers were supposed to solve large problems at a lower unit cost than smaller machines could.

One explanation for the industry's failure focuses on Moore's law, which improved processor and memory component technologies at a rate of 30 to 60 percent per year. A supercomputer that used two-year-old chips suffered a factor-of-two disadvantage in efficiency compared with a smaller system employing the latest chips.

As a result, the supercomputers that survived either use the latest chips in clusters with minimal custom components or run sequential code that offsets system inefficiencies by savings in the human software development effort. Architectures that required special programming and couldn't use the latest chips are now defunct.

The supercomputing industry's experience suggests two strategies to achieve economic scalability for the Internet, one *temporal* and the other *spatial*.

The temporal strategy simply calls for waiting. Moore's law applies in varying degrees to all of the Internet's constituent technologies, except perhaps the costs of cable right-of-ways and satellite launches. By waiting, the Internet becomes cheaper. But we might have to wait a decade or more for the cost of a video stream to drop by the estimated factor of 100 needed to deliver one hour of 1-Mbps video for US$0.05 instead of the current $5.

The spatial strategy concentrates on improving efficiency for a given service as it grows in popularity. This strategy enhances the Internet's commercial potential by promising lower costs for more popular content and applications. By comparison, the temporal strategy improves efficiency for all types of content, applications, and services, whether they are popular or not.

IMPLEMENTATION TECHNOLOGIES

Bandwidth expansion devices and multicast content delivery networks are two technologies that can help.

Bandwidth expanders

Given that bandwidth represents the dominant cost factor in any Internet-based application or service, numerous products evolved over the past decade to expand bandwidth. Caching and compression are key technologies underlying many of these products.

We can analyze a bandwidth expander's effect on a single link. This analysis adopts the viewpoint of an ISP selling bandwidth over an access link, either to a server in the context of a hosting service or to a business or home.

The exponential growth in total Internet backbone traffic for the past 20 years reflects both more links and more capacity in individual links. For simplicity, assume that these two factors contribute equally to the growth—about 40 percent per year each.

If a bandwidth expander can double a link's capacity, an ISP can postpone upgrading the link for two years. The link's growth rate, however, will remain unaffected. Two years after installing the expander, the ISP will still have to double the link bandwidth and the bandwidth expander's capacity as well.

A skeptic might conclude that eliminating one doubling of a link's bandwidth has negligible effect on Internet efficiency, since the costs saved in the first two years of deploying a bandwidth expander are offset by the costs of running, managing, and upgrading it forever.

However, the technology underlying bandwidth expanders also improves with time according to Moore's law, assuming that the expander's software does not squander the hardware enhancements. Because Moore's law matches the growth in individual link bandwidth by doubling hardware performance every two years at no extra cost, ISPs that incorporate bandwidth expanders can upgrade the device's capacity at a fixed per-link recurring cost. If that recurring cost is cheaper than the cost of supporting links that are twice as large, the network will be that much more efficient than competing networks that lack expanders.

Bandwidth expanders can thus save the difference between the cost of bandwidth and the cost of the devices themselves. They do not, however, fundamentally improve the functional relationship between capacity (bandwidth delivered) and resources (raw bandwidth required).

Multicast-enabled distribution

Multicast technology has a stronger impact on network efficiency, albeit at the severe cost of requiring all receivers to receive the same transmission at the same time, and sometimes even at the same bandwidth. Multicast routing creates a distribution tree of network links in which the root is the source of the information being multicast and the leaves are the receivers. The resources used are routers (tree nodes) and communication links (tree edges).

Bandwidth represents the dominant cost factor in any Internet-based application or service.

For simplicity, assume that the multicast tree is binary, although the conclusions apply in any case. If we double the multicast delivery rate by doubling the number of receivers, we need to add a new level in the tree. The resources everywhere else in the multicast tree remain constant. The number of added leaf links and nodes approximately equals the number of interior links and nodes, thus doubling the tree's size and cost to double its capacity. So, if all resources are equally valuable, the efficiency remains constant after doubling the delivery rate.

In practice, however, links located nearer to the multicast source tend to be physically much longer, which increases their cost relative to leaf links at the network's edge. Deploying long-distance cables requires more digging, the lasers that drive them are expensive, and so is maintenance—especially if they are marine cables or satellite links. Adding resources only at the network's edge doubles a tree's capacity with far less than twice the resources. This makes multicast an economically scalable technology.

Unfortunately, multicast places constraints on the receivers to receive the same content at the same time, which renders it impractical for the vast majority of Internet applications and content. Multicast schemes that relax this requirement reduce their economic scalability in proportion to the nonoverlap of reception by different receivers.

A better solution is to deploy special storage and computation nodes near the network's edge and have them specialize in receiving multicasts of popular content and applications, replicating the material, and then delivering it on demand over the last hop or two in the multicast tree. This basic architecture describes a multicast-enabled content distribution and delivery network.

Multicast-enabled CDNs are economically scalable. Consider the extreme case in which every local area network contains a copy of a given piece of content that has been multicast to a local content-delivery server. In this case, doubling the audience simply doubles the infrastructure's efficiency because no additional resources are required—assuming that bandwidth and server costs on the LAN are negligible.

Any technology capable of a similar resource-utilization pattern can yield similar benefits. Distributed caching, in which *every* router contains a cache, is a case in point.

Multicast-enabled content distribution and delivery networks, as well as distributed caching, provide clearer economies of scale than localized bandwidth-enhancing technologies such as caching and compression on a single link. The horizon is open, however, for exploring technologies that let the Internet deliver its economic potential to the world at large. ■

***Abdelsalam "Solom" Heddaya** is chief technology officer and cofounder of InfoLibria. He is also an adjunct associate professor of computer science at Boston University. Contact him at a@heddaya.net.*

Editor: Sumi Helal, Computer and Information Science and Engineering Dept., University of Florida, P.O. Box 116125, Gainesville, FL, 32611-6120, helal@cise.ufl.edu