from *SIAM News*, Volume 32, Number 10

Akamai Technologies: A Mathematical Success Story

By Alan Edelman

What a success story for mathematics! No, it's not a proof of the Riemann hypothesis, but a multibillion-dollar mathematical event. It's the story of a research problem that blossomed into one of the hottest recent IPOs on Wall Street. And it's all due to the tremendous impact of the Internet.

The mathematicians are Tom Leighton (of the applied mathematics group and the Laboratory for Computer Science at MIT) and Danny Lewin (often referred to as "Tom's graduate student" but a visionary in his own right). Together, they created Akamai Technologies. Akamai is now a large company, with hundreds of employees (many of whom are SIAM members; Leighton himself is a SICOMP editor).

Like any good high-tech company, Akamai is run by someone with business experience. George Conrades, the CEO, is a former senior vice president of IBM and more recently a venture partner at Polaris Venture Partners; he joined Akamai last spring. Jonathan Seelig, another Akamai founder, recognized the opportunity while at MIT's business school; it is probably fair to say that he became a business school sensation.

Akamai is very much the talk of MIT these days, and the news has spread rapidly among mathematicians everywhere. Some are thrilled with the great news and others are inspired, while some at MIT are concerned about the implications for academics. At least five faculty members in MIT's applied mathematics group are working at Akamai, on a part-time basis or more.

A Limosine Service for the Web

Akamai speeds up the delivery of Web pages with its "FreeFlow" service. On entering Akamai's Web site, once you've passed the flashy "Shockwave" graphic, you will see the slogan: "Internet Content Delivery: Guaranteed." Mathematically, if T is the download time for a Web page, they strive to make T as small as possible, and they guarantee that $T < \infty$. Backing up this claim are mathematical cleverness and programming talent. The policy is that if Akamai fails to perform, the customer does not pay. No excuses.

I recently began working at Akamai part time, while on sabbatical from the mathematics department at MIT. I had never expected to find the atmosphere of MIT reproduced in industry—people working around the clock to exceed ordinary standards of quality. (Akamai has talented people from many places, not only MIT, but I think it is fair to say that the MIT culture and the Akamai culture are indistinguishable.) Readers will see that my enthusiasm level is high as I write this; it has been great fun going to Akamai.

Understanding Akamai requires some understanding of the nature of the Internet. The Internet is much like the streets of Boston, or perhaps any old, unplanned city. Connections were based at first on local needs, and thoroughfares were superimposed later. Routes are constantly clogged or unavailable due to repairs. Similarly, no central authority built the Internet; that's why it is so widely available today, and it's also why it is so messy and chaotic. Akamai is something like a limousine service that guarantees a short ride no matter what the traffic conditions.

It is easy to tell if a Web page is "Akamaized." My favorite example is *The New York Times*. Go to www.nytimes.com and view the source. (As many of you know, the official term for a Web address is a URL, or Universal Resource Locator, sometimes called a Universal Record Locator or Uniform Resource Locator. Most URLs begin with the prefix http:// and most browsers assume the http:// prefix by default.) If you search the source for Akamai, as I have just done, you will find a long URL: http://a836.g.akamaitech.net/7/836/967/3cee073de96781/graphics.nytimes.com/images/nytlogo12.gif. A few clicks away is the familiar newspaper masthead logo:

The New York Times

The URL seems terribly long, but of course nobody ever types it. (Internally, an Akamaized URL is known as an ARL.) Just for fun, you can grab the last three fields from the above URL: graphics.nytimes.com/images/nytlogo12.gif. Again, with just a few clicks, you will see the same image. So why did my image come from the long Akamaized URL rather than the shorter one?

The shorter URL produces the image "live" from New York, which might not be the most convenient location for every user. The New York server could also slow down or die in the event of an overload of requests or some other failure. But Akamai servers, located around the world, grab the image once from New York and serve it to readers worldwide from convenient locations near them. Akamai customers can choose to Akamaize some or all of their content.

Even more fun is the Star Wars trailer: http://www.apple.com/quicktime/trailers/fox/episode-i/480.html. The long version of the URL for the quicktime movie that is being downloaded is http://a1912.a.kamai.com/7/1912/52/3847dda752238/download.akamai.com/apple/menac480.mov.

If you try to download the movie, you will probably notice the a.kamai on the bottom of your browser. A large number of well-known content providers, including CNN, Yahoo, and ESPN, are now Akamaized.

Dynamically Changing Traffic Patterns

Many readers may be familiar with the simple idea of "mirroring." For two servers with exactly the same content—one in the U.S. and the other in Europe, for example—mirroring reduces the need for ongoing trans-Atlantic communication. More sophisticated is the notion of caching—remote servers fetch a page from the original server only the first time it is requested; thereafter, the "caches" serve the requests until the page expires or the least popular material is removed from the cache to make space. In summary, mirroring is concerned with geographic locality (really network locality), while caching makes it possible to handle large demand. (Mirroring would be enough in a world of low demand for Web pages. Caching would be enough in a tiny world with high demand.)

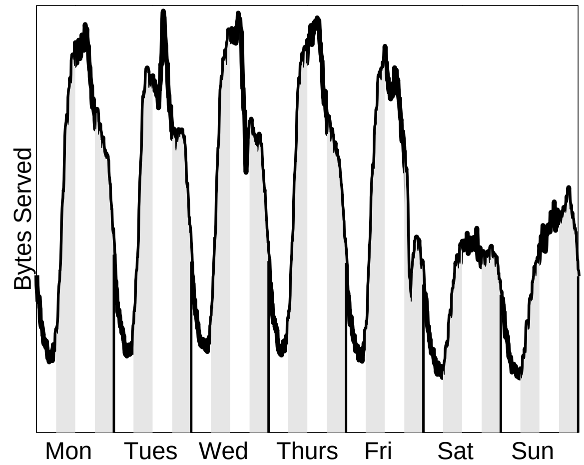

Figure 1. *Number of bytes served to an Akamai customer during a typical week.*

Mirroring and caching are still not enough in a world of dynamically changing traffic patterns. Returning to our U.S./Europe mirror example: During the morning hours in Europe, when people in the U.S. are still asleep, a European mirror could well be overloaded. In this situation, new requests might best overflow to an American server. This is but one example of the problems that must be solved because of the dynamic nature of the Internet. The real world of the Internet is very messy: Messages get lost for inexplicable reasons, servers go up and down, content becomes more popular and then less popular, traffic can surge, lull, and then surge again.

Akamai has a network of 1475 servers (and the number is growing) spread around the entire planet. These servers and the entire Akamai system are dedicated to optimizing the combined issues of network locality, user demand, Internet traffic, and hardware reliability. On arriving at Akamai's headquarters, in Cambridge, Massachusetts, a visitor sees a command center that makes this author think of something like NASA's Mission Control. There, the Internet is monitored and warning lights activate when trouble arises.

The system details are proprietary, but roughly speaking, a request for a Web page is mapped to a server. This server is likely to be geographically nearby, but Internet traffic and general bandwidth issues are really the deciding factor. If the Web page is in the cache, it is served; otherwise, the first step is to fetch the page. Dynamic routing occurs based on immediate knowledge of real-time Internet monitoring and the requirement to serve fresh content.

Measuring the Internet can also be a challenging problem. I suspect that it is easier than obtaining data about automobile traffic patterns, because a lot can be done with software. We can already make predictions: For most American content providers, demand is lowest at 5 in the morning, Eastern time. By then, most Americans are probably asleep and Europeans have not yet reached their lunch breaks. The cycle of lunch breaks tends to start the surge. The issue of measurement fascinates me. Probably *SIAM News* readers would not be surprised to learn that I consider Matlab the right data playground for Internet data exploration.

Just for fun, a graph of a typical week for one of Akamai's customers is shown in Figure 1. Every five minutes, we take a count of the number of bytes served. Each of the vertical (shaded and white) regions represents six hours; a day begins at midnight, Eastern Standard Time. The daily human cycles are easily observed. We can match one weekday against another, or one week against the next, to predict the daily trends and the weekly growth.

"Consistent hashing" is the mathematical idea underlying the caching of Web pages. "Hashing" itself refers to the fast generation of a relatively short identifier for a larger piece of text. The hash value serves as a pointer to the item itself. If every cache had a copy of every item, there would be no room left on disk. If the items are distributed, then a quick addressing scheme is needed to locate the right cache. Any scheme used should require a small number of changes when new objects need to be cached and old objects are deleted. This is the mathematical problem that began the whole story. The problem was originally posed by Tim Berners-Lee, founder of the World Wide Web, whose office in MIT's Laboratory for Computer Science is on the same floor as Tom Leighton's.

Will Academia Suffer?

I would like to close with some observations about industry and academia. A large number of technical employees at Akamai are MIT students and faculty. Many work part time, although Akamai is so invigorating that it is easy to spend more time there. Concerns have been voiced that students, hooked by Akamai, are neglecting their studies or social life. Similar concerns are expressed about faculty. Many of us are aware that industry in general is attracting many of the best minds, and academia does not always seem to be able to compete.

Some gloom-and-doom predictors believe that academia will suffer because of the allure of Internet start-ups. Indeed, many talented people are leaving academia, and classes are harder to staff. The five mathematicians and two computer scientists who are consulting or on temporary leave at Akamai retain their commitments and relationship with MIT. I believe that we will all return with the benefit of a real-world applied mathematics and computer systems experience that will enhance our research and teaching. Applied mathematics will only have gained.

Further Reading

D.R. Karger, E. Lehman, T. Leighton, M. Levine, D. Lewin, and R. Panigrahy, Consistent hashing and random trees: *Distributed caching protocols for relieving hot spots on the World Wide Web*, Proceedings of the 29th ACM Symposium on Theory of Computing, ACM, ACM Press, May 1997, 654–663.

D. Karger, A. Sherman, A. Berkheimer, B. Bogstad, R. Dhanidina, K. Iwamoto, B. Kim, L. Matkins, and Y. Yerushalmi, *Web caching with consistent hashing*, Eighth World Wide Web Conference, Toronto, Canada, May 1999, http://www8.org.

Alan Edelman is an associate professor of applied mathematics and a member of the Laboratory for Computer Science at MIT.