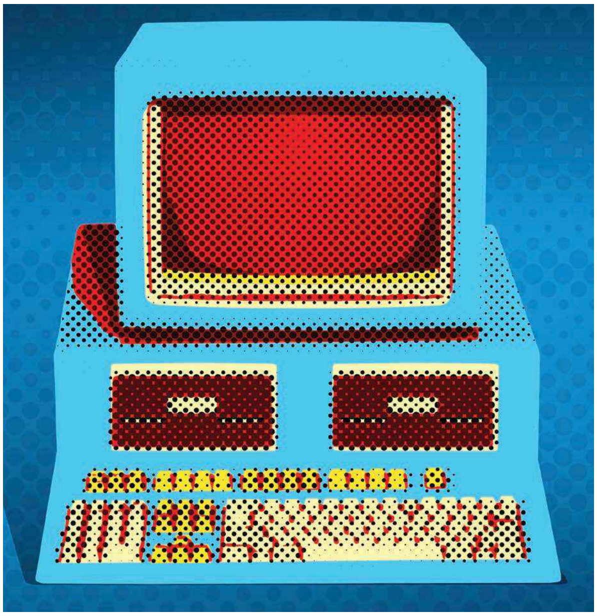

Metcalfe's Law after 40 Years of Ethernet

Bob Metcalfe, *The University of Texas at Austin*

Critics have declared Metcalfe's law, which states that the value of a network grows as the square of the number of its users, a gross overestimation of the network effect, but nobody has tested the law with real data. Using a generalization of the sigmoid function called the netoid, Ethernet's inventor and the law's originator models Facebook user growth over the past decade and fits his law to the associated revenue.

On May 22, 2013, Ethernet industry leaders gathered at the Computer History Museum in Mountain View, California, to celebrate Ethernet's 40th birthday. More than that, we gathered to praise Ethernet's many unsung heroes, to gather lessons from its various innovation successes, and to catch up on progress in what has become a $100 billion industry.

As part of that celebration, we also revisited the "law" bearing my name. Metcalfe's law states that the value of a network grows as the square of the number of its users: $V \sim N^2$. This law started as a high-concept Ethernet sales tool in the early 1980s. It entered public discourse in the mid-1990s when George Gilder, who had earlier helped to popularize Moore's law in his book *Microcosm*,[1] championed my law in his sequel *Telecosm*.[2]

The US is now at great expense "gigafying" the Internet—upgrading Internet infrastructure to enable next-generation applications in a wide range of areas including education, energy, and healthcare (http://us-ignite.org). The network effect generally, and Metcalfe's law in particular, are often invoked to justify such large investments.

Critics of my law have argued that N^2 greatly overestimates the network effect. Not only is the law wrong, they say, it is dangerous, because quantifying the network effect is central to society's major infrastructural investment decisions. Such decisions—thanks in part to Metcalfe's law—went famously awry during the dot-com bubble in the late 1990s. Today, detractors say, my law is causing stock markets to grossly overvalue Internet-based companies such as Google, LinkedIn, Facebook, Twitter, and Snapchat.

However, nobody (including me) has ever made the case for or against Metcalfe's law with real data. On this occasion of the 40th birthday of Ethernet, I revisit my law—which is not to say revise it. Using a generalization of the sigmoid function called the netoid, I have modeled Facebook user growth over the last decade and fitted Metcalfe's law to associated revenue, a surrogate of Facebook's network value.

Furthermore, I argue that regardless of how precise a predictor Metcalfe's law is, it remains an important agenda setter in innovation.[3]

BIRTH OF ETHERNET

Ethernet was born on May 22, 1973, in a memo I circulated at the Xerox Palo Alto Research Center (PARC)

Published by the IEEE Computer Society

0018-9162/13/$31.00 © 2013 IEEE

Metcalfe's Law Stands the Test of Time

Sumi Helal, *University of Florida*

Metcalfe's law as originally stated did more than estimate the value of a network in terms of the number of its nodes: it drew rebuttals and a quest for a deeper understanding of this value, leading to this "revisit" article.

Metcalfe explains his original, theoretic growth model of networks using a practical counterpart that captures user adoption of Internet services. Maximum adoption rate, speed of adoption, and peak adopting population are the three parameters that shape this process, which he models as a sigmoid and calls the netoid. Validating the original law against real adoption data obtained from Facebook using a netoid clearly shows that the law has been remarkably accurate. Explaining and modeling network growth by utilizing user adoption of network services reflects the value transformation that has occurred over the years as the Ethernet became commoditized and the Internet permeated the globe.

But on whose behalf are we assessing network value? In Metcalfe's exemplary validation study of the netoid, Facebook was the value receiver, marking a critically important and broad category of value recipients—social networks. As such, it would be prudent to give further consideration to other value receiver categories. Perhaps we should make sure that the netoid is applicable and adequate to all categories. For instance, the user, whose adoption directly fuels the value of social networks, is surely driven by a perceived value as well, so how do we assess the network's value to its users? How can LinkedIn convince users it attempts to recruit of the value of its service to them, given the monthly premium? Knowing the answer would be helpful both to the user and to LinkedIn. The user is evidently aware of the value concept, and unconsciously applies Metcalfe's law (by owning multiple nodes—PCs, iPads, smartphones, and so on) for access anytime, anywhere.

One critical aspect of the Internet that is difficult to ignore, in analyzing value receivers, is content. Several stakeholders find a proven value in content—the businesses who create it to increase profits, the users who create it to share, and search engines who index it and facilitate search for advertisement profits. Google, for instance, would not reap the value of the network without content. Understanding the interplay between physical network growth and content growth could prove to be useful in deepening our understanding of Metcalfe's law.

Perhaps another value receiver category would include cloud service providers and users. If we look at clouds as super network nodes packed with incredible values, then the value of the overall network to the user obviously increases (higher value than a network without super nodes!). In fact, users are guaranteed a value simply equal to the difference between the high cost of ownership and the cloud's incremental service fees. The network value to cloud providers, on the other hand, does not seem all that simple to estimate or ponder, and may entail other processes in lieu of or addition to adoption (for example, transactions).

All in all, revisiting Metcalfe's law is a great renewal of the importance of understanding the broad value of the technologies we create. Understanding network value will not only help us assess it more reliably but also provide guidance into how we might create such value and grow it or predict its growth more sensibly.

Sumi Helal is a professor in the Computer and Information Science and Engineering Department and directs the Mobile and Pervasive Computing Laboratory at the University of Florida. He is also a Computer *editorial board member. Contact him at helal@cise.ufl.edu.*

showing how a local area network (LAN) might work. The memo included the sketch in Figure 1.

After cofounding 3Com Corporation in 1979—in large part to commercialize Ethernet—I served as vice president of sales and marketing from 1982 to 1984. Following announcement of the IBM PC in August 1981, we recruited six former minicomputer salespeople to sell Ethernet adapter cards to IBM PC owners. Since minicomputers cost something like $30,000 and our Ethernet adapters cost about $1,000, our commissioned sales force initially set out to sell 30-node Ethernets to early PC owners. Trouble was, PCs were new and few people knew what a LAN was for. In our profession's lingo, our sales cycle went to infinity.

Try to imagine the early 1980s mindset. With coax cable in hand, I was told repeatedly that nobody would ever install new wiring to carry Ethernet packets among computers—they would have to be carried, if at all, over existing ubiquitous electricity power lines. Today, of course, there's an IEEE standard for delivering

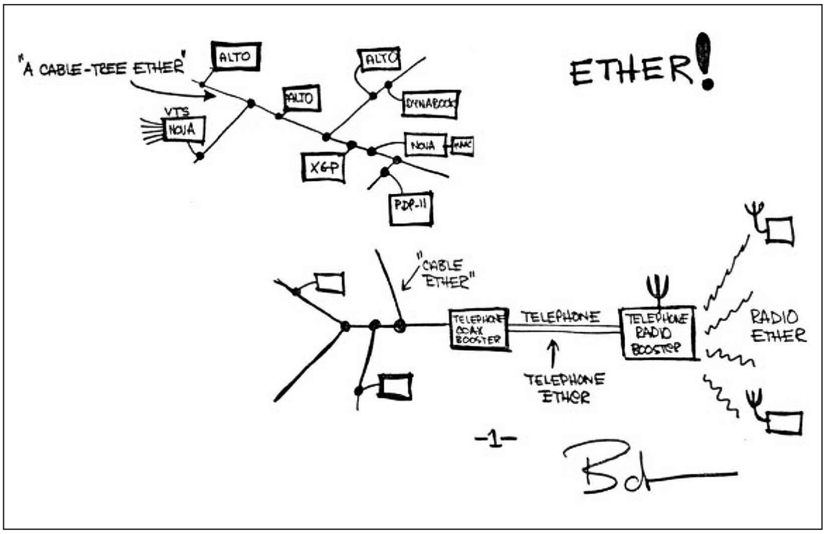

Figure 1. Diagram of Ethernet that appeared in a May 22, 1973, memo written by the author at the Xerox Palo Alto Research Center. Image provided courtesy of Palo Alto Research Center Inc., a Xerox Company.

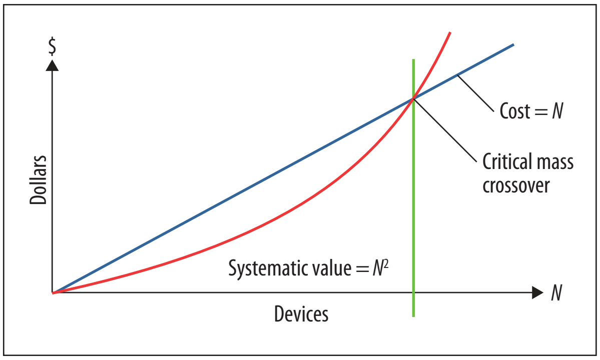

Figure 2. Graph showing that the systemic value of a network is proportional to the square of the number of devices connected to the network.

electric power over Ethernet cables. The cables, as it were, have been turned.

During a 1983 offsite meeting at Lake Tahoe, 3Com's sales and marketing team decided to lower the cost of trying Ethernet by offering customers a three-node starter kit with cables, connectors, and a diskette full of software. The $3,000 price was below the level most companies required for capital acquisition approval from upper management; adventurous early adopters, God bless them, could put kits on their expense accounts. Our salespeople were initially resistant, because the commission on a kit sale was small and they had mouths to feed, but then pioneering PC owners starting buying the kits in startling numbers.

Customers saw value in the starter kits because they allowed three PC users to share a printer and a hard disk, which in those days were quite expensive. In 1983, few businesses could afford the costly 10-Mbyte hard disks IBM offered on their revolutionary PC XTs. Apple's LaserWriter, released two years later, ran $7,000—I know, I bought one. Our starter-kit customers hoped to amortize the cost of a printer and hard disk over three PCs on a standalone three-node Ethernet. And what's more, they could send LAN email for free!

So, months later, 3Com's salespeople returned to our starter-kit customers to sell them another 30 Ethernet interface cards. But, while admitting that the kits performed as promised, customers told us that their three-node Ethernets were not all that useful. There just wasn't enough to say on an email network with three users. Uh oh.

As far back as 1972, I had experienced the high value of LAN-connected desktop PCs. Why didn't 3Com's customers have that same experience with their PC LAN starter kits in 1983?

THE NETWORK EFFECT AND METCALFE'S LAW

It was at this point that I came up with what, 15 years later, would come to be called Metcalfe's law.

During a presentation to the 3Com sales force, I projected a 35-mm slide with the graph in Figure 2. I argued that if a network is too small, its cost exceeds its value; but if a network gets large enough to achieve critical mass, then the sky's the limit. I claimed that the systemic value V of a network is proportional to the square of the number of compatibly communicating devices: with N nodes each connecting to $N - 1$ other nodes, V would be proportional to the total number of possible connections, $N \times (N - 1)$, or approximately N^2.

Armed with this quantification of the benefits of the network effect, the 3Com sales force went back out into the field determined to get their customers' Ethernets above critical mass, which we hoped to be about 30 nodes.

The rest of the story, as they say, is history. Many 3Com customers believed us and added more adapter cards to their trial Ethernets. Our company eventually grew from selling hundreds of Ethernet adapter cards per month to millions. 3Com went public on NASDAQ in March 1984. The company's revenue peaked at $5.7 billion in 1999 and, in 2010, 3Com was acquired by Hewlett-Packard.

Of course it wasn't only Metcalfe's law that convinced 3Com customers to start buying Ethernet adapter cards in volume. The price of the cards was rapidly falling, and Ethernet applications were growing beyond printing and disk sharing. Today the cost of Ethernet is nearly zero and included in the base price of PCs. Ethernet LANs have become the standard packet plumbing of the Internet, the medium through which most users access email and the World Wide Web. If you count Wi-Fi as wireless Ethernet, and I do, more than a billion Ethernet ports are shipped annually in all types of wired and wireless devices, from desktop PCs to mobile phones.

For the last 30 years, as the Internet has grown beyond all expectations, various observers have questioned the validity of Metcalfe's law. Nobody denies the all-important network effect in the growth of the Internet and its applications; they mainly wonder if squaring the number of a network's nodes estimates a network's value too highly. It's a fair question, especially considering that my law initially dealt with networks of 30 nodes, while today the Internet has an estimated 2.4 billion.

METCALFE'S COMPARED TO OTHER NETWORK "LAWS"

Metcalfe's law isn't the only—or even the first—network law. David Sarnoff, the "Father of American Television" who led RCA from 1930 to 1970, stated that the value of a broadcast network grew linearly in proportion to the number of viewers: $V \sim N$. The big difference between Sarnoff's broadcast networks and the Internet is, of course, that Internet users can derive value by communicating back to the broadcaster and with one another.

David P. Reed[4] postulated that a "group-forming network" can give rise to 2^N subnetworks—thus, $V \sim 2^N$. Obviously, 2^N quickly gets to be quite a bit larger than N^2, so Reed's law is a notable exception among complaints about Metcalfe's law overestimation.

The main critique of Metcalfe's law was best expressed by Andrew Odlyzko, Bob Briscoe, and Benjamin Tilley in an *IEEE Spectrum* article that described my law as both "wrong" and "dangerous."[5] Suspecting that not all network connections are of equal value and referring to Zipf's law, Odlyzko and his colleagues countered—in what I will call Odlyzko's law—that the growth in value of a network is approximately $N \times \ln(N)$.

However, Odlyzko's law suffers from the same two main problems as my own. First, like N^2, $N \times \ln(N)$ goes to infinity with N, and network values, like trees, do not grow to the sky. Odlyzko's law has network value growing less than Metcalfe's law, but still without bound. Second, neither of us has attempted to validate our law with data from real networks. Until now.

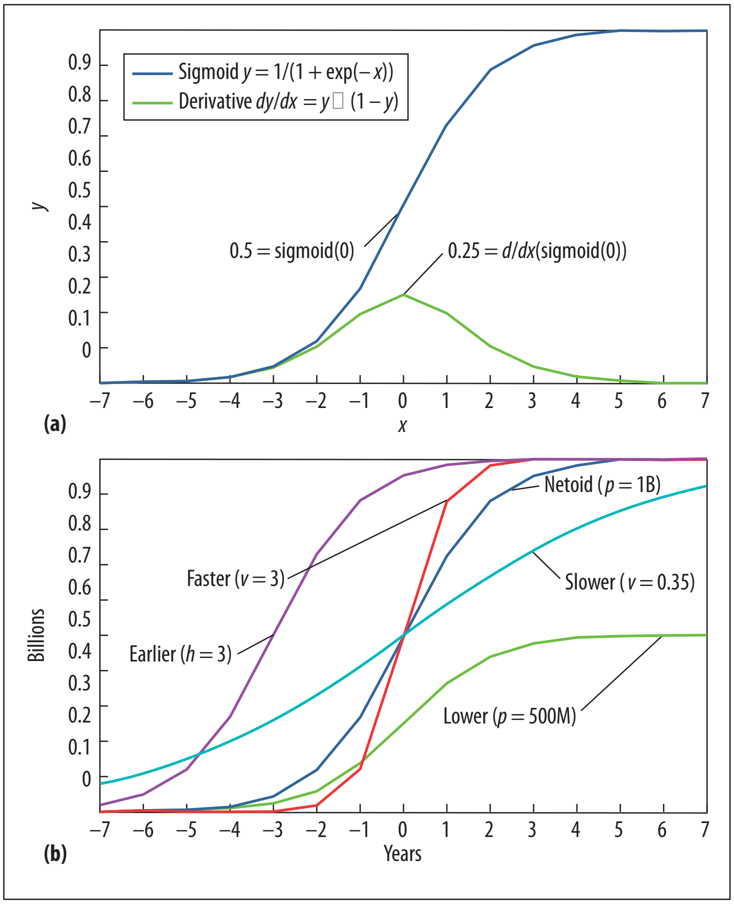

Figure 3. Generalizing the sigmoid function to the "netoid" function. (a) Sigmoid function computed and plotted with Python. (b) Netoid function variations (*h*, *v*, *p*) computed and plotted with Python.

REVISITING METCALFE'S LAW

To address the various critiques of Metcalfe's law, I'll put the law in context and attempt to fit it with real data. In short, the context for $V \sim N^2$ will be looking at N as a function of time. I won't try to limit network value as a function of N, but instead limit N as a function of time. And the real data will come from the last 10 years of Facebook.

Facebook, of course, epitomizes Reed's group-forming network. The company continues reporting exponential growth into billions of users, but what about its groups? Do groups of Facebook friends tend to have an infinite number of users or to approach some asymptote?

Enter Dunbar's number. Anthropologist Robin Dunbar theorized that humans have a cognitive limit on the number of people with whom they can maintain stable social relationships; he put the figure at about 150,[6] though others have suggested that the limit may range from 100 to 230.

At the end of 2012, Facebook had about 1.06 billion users and 150 billion friend connections (http://goo.gl/w9Li05). This implies an average of 141 friends per user, which is amazingly close to Dunbar's number. Of course, some of Facebook's newest users might not have finished growing their friends networks, and new Facebook tools could increase the number of friends that a user's cognition can sustain. So maybe the Facebook friends asymptote exceeds Dunbar's number.

Netoid function

Network laws rely on various functions: linear, quadratic, logarithmic, exponential. To estimate the growth of Facebook users, and that of individual groups of friends, let's apply the sigmoid function, a particular S-curve adoption function. The sigmoid models a population's growth from 0 percent at time minus infinity to 100 percent at time plus infinity. The sigmoid adoption rate peaks at time 0.0 with a population fraction of 50 percent.

Using the Python programming language, I'll first generalize the sigmoid to what I call the "netoid" function, as Figure 3 shows. Metcalfe's law is defined as follows:

```
def Metcalfe(n,c=1.0): ## V~N^2 ~ C*N*(N-1)
    if n<=1.0: return 0
    return c*n*(n-1.0)
```

The sigmoid function is defined as follows:

```
def sigmoid(x): return 1.0/(1.0+math.exp(-x))
```

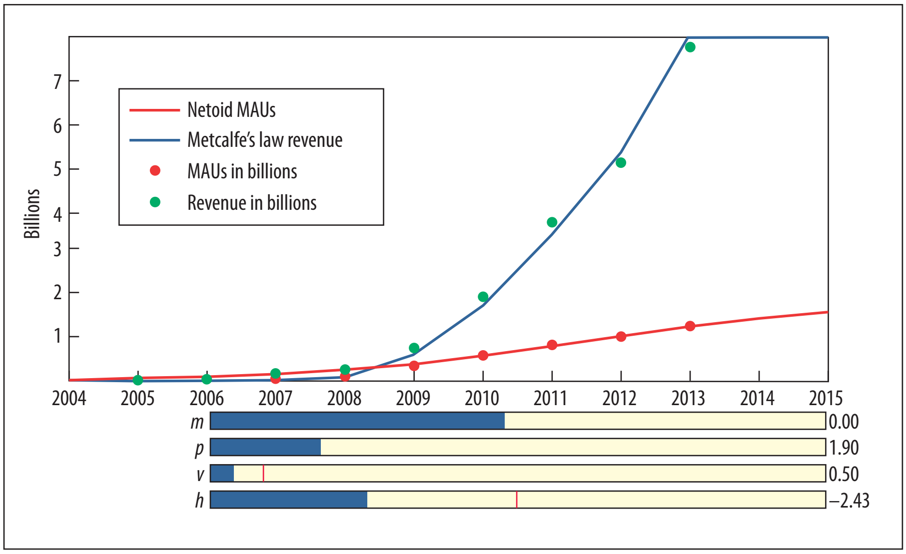

Figure 4. The netoid can be closely fitted to Facebook user growth data, measured in terms of monthly average users (MAUs), and Metcalfe's law can be closely fitted to Facebook's associated revenue data.

Generalizing the sigmoid to the netoid results in the following:

```
def netoid(t,h=0.0,v=1.0,p=1.0):
  ## "netoid" is s-curve sigmoid adoption
from zero to population
  ## t = moment in time
  ## h = center of sigmoid s-curve
  ## v = virality, how fast adoption
  ## p = peak, population, asymptote
  return float(p)/(1.0+math.exp(-float(v)*
(float(t)-float(h)))
```

The netoid has the same S-curve shape as the sigmoid. Its slope (the adoption rate) is proportional to the product of the fraction of the population already adopted times the fraction awaiting adoption. It peaks when adoption is 50 percent. The adoption rate is driven by the number of adoptions so far and limited by the number of those awaiting adoption.

The netoid offers three parameters: *h*, the point in time at which the growth rate is maximum, when the population is half the peak; *v*, the virality or speed with which adoption occurs; and *p*, the peak value, which the netoid approaches asymptotically. In short, the netoid can model when and how fast adoption will occur, and how large it will get.

Fitting the netoid to Facebook data

With these three parameters, the netoid can be fit to a wide range of real datasets.

I first plotted Facebook user growth data from 2004 to 2013 for monthly average users (MAUs) in billions. I fiddled in Python with the three slider parameters in the netoid until I achieved a good visual fit to this data, as Figure 4 shows. I observed that the asymptote (*p*) for the fitted netoid is about 2.5 billion Facebook users. This is today's number of Internet users, which is what you'd expect. Facebook hasn't yet reached 50 percent peak adoption and the Internet is still growing, probably along its own netoid.

Next, I plotted Facebook annual revenue for the last 10 years. I attached a Python slider to the Metcalfe's law function and once again got a pretty good visual fit to this data. If we take revenue as a surrogate for Facebook network value, then here we have the first application of Metcalfe's law to real data, albeit data before the netoid S-curve flattens Facebook growth. Of course, Facebook creates much more value than is captured and monetized by Facebook selling ads.

A more elaborate model of network growth and value would involve *nesting* netoids. Facebook groups seem to be netoiding toward Dunbar's number, while Facebook's number of users is netoiding toward the number of Internet users. The number of Internet users is likely also netoiding toward the population of Earth, which is estimated to reach 10 billion by 2100, also perhaps along a netoid.

Netoids can also be compounded into *cohorts*. As people join Facebook, they each presumably begin a Dunbar-number-limited netoid process of finding friends. Thus, at any given time, the total number of Facebook users is the growing sum of many netoids, in cohorts according to when they started.

So, what's driving the netoid parameters *h*, *v*, and *p*? What determines when a network takes off, how viral its growth will be, and how large it will eventually become? Trying to answer these questions takes us back to lessons from Ethernet's growth.

ETHERNET'S PROMISE: BANDWITH ELASTICITY

At the Computer History Museum celebration of Ethernet's 40th birthday, many people asked what the word "Ethernet" means today.

Strict constructionists define Ethernet as a 2.94-Mbps CSMA/CD (carrier sense multiple access with collision detection) LAN with 8-bit addresses running over 0.4-inch coaxial cable at Xerox PARC in the 1970s. Others more generously declare that Ethernet refers to the collection of various IEEE 802 standards. Some equate it with IEEE 802.3—that is, wired LAN as opposed to wireless LAN, or Wi-Fi (IEEE 802.11). Others say Ethernet is simply the idea of a PC LAN.

What I say—and I think I'm in charge—is that Ethernet has become an innovation brand. Brands are promises, so what does the Ethernet brand promise?

Ethernet is based on open de jure standards, particularly IEEE 802. Among the Internet protocol layers, it is the native-mode packet plumbing. Standardization maximizes interoperability and accelerates cost reduction through consolidated investment and price competition. Unlike open source implementations, Ethernet implementations are owned by their developers. Therefore, competition among Ethernet suppliers is fierce, which motivates them to listen to customers, drive prices down, and continue to innovate, while keeping their products interoperable. Ethernet standards evolve rapidly after market engagement, but that evolution is constrained by the promise of backward compatibility. Metcalfe's law is premised on maintaining installed bases and leveraging the network effect, not starting network growth all over again with each new generation.

Another promise of the Ethernet brand is "Build it, and they will come." Consider, for example, data transmission speed. Ethernet started at 2.94 Mbps in 1973 but now runs at speeds up to 100 Gbps. IEEE is in the process of standardizing 400 Gbps, with terabit Ethernet around the corner. With past Ethernet speedups, many questioned whether the new speed was too high, unjustified by known applications. But each time, Ethernet kept its promise, with faster speeds followed by unanticipated applications and, after them, many new users. Today's Ethernets are plumbing the way from the megabit Internet to the gigabit Internet.

In parallel with the gigabit Internet, the interconnected mobile Internet is also growing. Of particular note is the current market battle between LTE (Long Term Evolution), a mobile Internet technology, and Wi-Fi, an IEEE 802 gigabit Internet technology. Both carry packets to and from back-haul gigabit Ethernets.

So, after 40 years, Ethernet continues to exhibit bandwidth elasticity: the more bandwidth we supply, the more the world demands. When will this end? Is there a Dunbar number of Ethernet speed, above which some limit on human intelligence makes Ethernet overkill? Not yet.

Metcalfe's law implies a critical mass point in network size, after which network value begins to exceed its cost. That critical mass point is roundly given by the ratio of the cost of the network to the value of network participation. In the Internet, this ratio has been going rapidly to zero. Why?

The asymptotes of various network netoids are moving according to Moore's law, which states that the number of transistors on an integrated circuit doubles every two years or so. Metcalfe's law depends on Moore's law in two ways. Faster and cheaper semiconductor processors and memory are enabling more valuable applications that demand ever-larger bandwidths. Meanwhile, faster and cheaper network ICs are driving down the cost of networking.

Moore's law is expected to continue for another 15 years. We've heard predictions like this before, but since Ethernet's bandwidth elasticity depends on the continuation of Moore's law, let's hope that Moore's law doesn't soon hit one of its netoid asymptotes, such as the speed of light, the optical limits of lithography, quantum effects at smaller feature sizes, or overheating.

Of course, Moore's law isn't an inevitable law of nature—it's more a self-fulfilling prophecy that relies on continuing investment decisions at many levels among semiconductor scientists and engineers, chipmakers, and device makers.

And so too should we continue investing in Internet/Ethernet technology and thereby increase freedom and prosperity. Build it, and they will come.

References

1. G. Gilder, *Microcosm: The Quantum Revolution in Economics and Technology*, Touchstone, 1989.
2. G. Gilder, *Telecosm: The World after Bandwidth Abundance*, Touchstone, 2000.
3. B. Metcalfe, "There Oughta Be a Law," *The New York Times*, 15 July 1996.
4. D.P. Reed, "That Sneaky Exponential—Beyond Metcalfe's Law to the Power of Community Building," 1999; www.reed.com/dpr/locus/gfn/reedslaw.html.
5. B. Briscoe, A. Odlyzko, and B. Tilly, "Metcalfe's Law Is Wrong," *IEEE Spectrum*, 1 July 2006; http://spectrum.ieee.org/computing/networks/metcalfes-law-is-wrong.
6. "Primates on Facebook," *The Economist*, 26 Feb. 2009; www.economist.com/node/13176775.

***Bob Metcalfe** is Professor of Innovation and Murchison Fellow of Free Enterprise at The University of Texas at Austin. An Internet Society Hall of Fame Pioneer and National Inventors Hall of Famer, he received both the IEEE Medal of Honor and the National Medal of Technology for his leadership in the invention, standardization, and commercialization of Ethernet. Metcalfe received a PhD in computer science from Harvard University. He is a member of the National Academy of Engineering and a Life Trustee of his alma mater, the Massachusetts Institute of Technology. Contact him at bob.metcalfe@utexas.edu.*