Metcalfe's Law: A misleading driver of the Internet bubble

Bob Briscoe[1], Andrew Odlyzko[2], and Benjamin Tilly[3]

[1] Networks Research Centre, BT Research
B54/130 Adastral Park
Martlesham Heath, Ipswich, IP5 3RE
United Kingdom
bob.briscoe@bt.com
[2] Digital Technology Center, University of Minnesota
499 Walter Library, 117 Pleasant St. SE
Minneapolis, MN 55455, USA
odlyzko@umn.edu
http://www.dtc.umn.edu/~odlyzko
[3] ben_tilly@operamail.com
Preliminary version, February 11, 2006

1 Introduction

Metcalfe's Law states that the value of a communication network is proportional to the square of the size of the network. The name was coined by George Gilder in the early 1990s, who proclaimed the importance of this rule in influencing the development of the *New Economy.* Soon afterwards, Reed Hundt, the then-chairman of the Federal Communications Commission, claimed that Metcalfe's Law and Moore's Law "give us the best foundation for understanding the Internet." A few years later, Marc Andreessen attributed the rapid developments of the Web, and in particular the growth in AOL's subscriber base and of the time AOL members spent online, to Metcalfe's Law. This "law" (which is just a general rule-of-thumb, not a physical law) has attained a sufficiently exalted status that a recent article in SPECTRUM classed it with Moore's Law as one of the five basic "rule-of-thumb 'laws' [that] have stood out" and passed the test of time. What we demonstrate is that Metcalfe's Law is false (and so is Reed's Law, which says the value of a network of size n is proportional to 2^n. We also propose $n\log(n)$ as a more appropriate approximation to the value of a network of size n. The slower growth rate of $n\log(n)$ as opposed to the n^2 of Metcalfe's Law helps explain the failure of the dot-com and telecom booms, and has implications for network growth and network interconnection.

Metcalfe's Law dates back to around 1980. In the words of Bob Metcalfe, the inventor of Ethernet (as well as a successful entrepreneur, tech industry pundit, and more recently venture capitalist), who stated it first,

> The original point of my law (a 35mm slide circa 1980, way before George Gilder named it, in 1995) was to establish the existence of a cost-value cross-over point – critical mass – before which networks don't pay. The trick is to get past that point, to establish critical mass, which is why it has been so important for generations of Ethernet to be backward compatible.

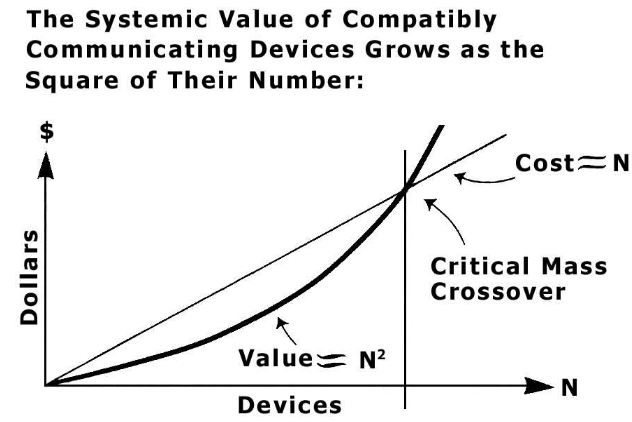

Fig. 1. Bob Metcalfe's original slide, from around 1980, that gave rise to Metcalfe's Law. Reproduced by permission of Bob Metcalfe.

Metcalfe's historic slide is shown in Fig. 1.

Metcalfe's Law was one of the popular mantras of the Internet boom, together with "Internet time," "network effects," "first-mover advantage," and "build it and they will come." It was especially popular among engineers, since it seemed to provide a quantitative explanation for the other mantras, and to justify the high hopes for the dot-com and telecom startups. The mad rush for growth and the neglect of profits seemed natural when one had assurance that once that magical cross-over point was reached, profits would escalate rapidly. And profits seemed bound to occur. For although Metcalfe's Law claims to describe the value of the network, not necessarily the profits that a service provider can extract from it, yet with the value growing quadratically in the number of users, it seemed that there was plenty of room to get bountiful profits from that value.

Reflecting on the Internet bubble of the late 1990s, we see that there was some validity to many of the Internet mantras of that time. Yahoo!, eBay, and Google are very successful dot-coms that exploit the power of the Internet to provide attractive services that do yield great profits. Yet these successes are few compared to the dreams of the dot-com era, and they have come far more slowly. The rapid growth of AOL that Andreessen attributed Metcalfe's Law came to a screeching halt, and the company has struggled in the last few years. The $n \log(n)$ valuation for a network provides a heuristic explanation for the dynamics that led to this disappointingly slow growth. On the other hand, since this growth is faster than the lines growth that Sarnoff's Law prescribes for the value of a broadcast network, it helps explain the dot-com successes we have witnessed.

Our $n \log(n)$ valuation of a communication network should not be taken too seriously. It is just a rough heuristic, and ignores a variety of group-forming phenomena that take place in networks, especially networks as versatile as the Internet. But if we search for a parsimonious description of network's value, something as simple to state as Metcalfe's Law, then $n \log(n)$ appears the best choice. It is supported by several quantitative heuristics, and fits in with observed developments in the economy.

2 The falsity of Metcalfe's Law and Reed's Law

The foundation of Metcalfe's Law is the observation that in a general communication network with n members, there are $n(n-1)/2$ connections that can be made between pairs of participants. If all those connections are equally valuable (and this is the big "if" that we will discuss in more detail below), the total value of the network is proportional to $n(n-1)/2$, which, since we are dealing with rough approximations, grows like n^2.

Reed's Law is based on the further insight that in a communication network as flexible as the Internet, in addition to linking pairs of members, one can form groups. With n participants, there are 2^n possible groups, and if they are all equally valuable, the value of the network grows like 2^n, the statement of Reed's Law.

The fundamental fallacy underlying Metcalfe's and Reed's laws is in the assumption that all connections or all groups are equally valuable. The defect in this assumption was pointed out a century and a half ago by Henry David Thoreau. In *Walden,* he wrote:

> We are in great haste to construct a magnetic telegraph from Maine to Texas; but Maine and Texas, it may be, have nothing important to communicate.

Now Thoreau was wrong. Maine and Texas did and do have a lot to communicate. Some was (and is) important, and some not, but all of sufficient value for people to pay for. Still, Thoreau's insight is valid, and Maine does not have as much to communicate with Texas as it does with Massachusetts or New York, say.

In general, connections are not used with the same intensity (and most are not used at all in large networks, such as the Internet), so assigning equal value to them is not justified. This is the basic objection to Metcalfe's Law, and it has been stated explicitly by many observers, including Bob Metcalfe hinself. Some users (those who generate spam, worms, and viruses, for example) actually subtract from the value of a network.

There are additional scaling arguments that suggest Metcalfe's and Reed's laws are incorrect. For example, Reed's Law is implausible because of its exponential (in the precise mathematical sense of the term) nature. If a network's value were proportional to 2^n, then there would be a threshold value m such that for n below $m-50$, the value of the network would be no more than 0.0001% of the value of the whole economy, but once n exceeded m, the value of the network would be more than 99.9999% of the value of all assets. Beyond that stage, the addition of a single member to the network would have the effect of almost doubling the total economic value of the world. This does not fit general expectations of network values and thus also suggests that Reed's Law is not correct.

Metcalfe's Law is slightly more plausible than Reed's Law, as its quadratic growth does not lead to the extreme threshold effect noted above, but it is still improbable. The problem is that Metcalfe's Law provides irresistible incentives for all networks relying on the same technology to merge or at least interconnect. To see this, consider two networks, each with n members. By Metcalfe's Law, each one is (disregarding the constant of proportionality) worth n^2, so the total value of both is $2n^2$. But suppose these two networks merge, or one acquires the other, or they come to an agreement to interconnect. Then we will effectively have a single network with $2n$ members, which, by Metcalfe's Law, will be worth $4n^2$, or twice as much as the two separate networks. Surely it would require a combination of singularly obtuse management and singularly inefficient financial markets not to seize this obvious opportunity to double total network value by a simple combination. Yet historically there have been many cases of networks that resisted interconnection for a long time. For example, a century ago in the U.S., the Bell System and the independent phone companies often competed in the same neighborhood, with subscribers to one being unable to call subscribers to the other. Eventually interconnection was achieved (through a combination of financial maneuvers and political pressure), but it took two decades. In the late 1980s and early 1990s, the commercial online companies such as CompuServe, Prodigy, AOL, and MCIMail provided email to subscribers, but only within their own systems, and it was only in the mid-1990s that full interconnection was achieved. More recently we have had (and continue to have) controversies about interconnection of IM systems and peering of ISPs. Thus the general conclusion is that the incentives to interconnect cannot be too strong, and so Metcalfe's Law cannot be valid.

We have several quantitative arguments for the $n \log(n)$ rule-of-thumb valuation of a general communication network of size n. Some are based on information locality (Bradford's Law of scattering from library science), others on geographical locality phenomena (the numerous "gravity models" that describe intensity of transportation and telecommunication interactions). They are available in a longer version of the paper available at the second author's Web site (http://www.dtc.umn.edu/∼odlyzko). Here we present just one simple argument, based on Zipf's Law.

3 The ubiquitous Zipf's Law and the value of networks

Zipf's Law is yet another of the empirical rules that describe well many real-world phenomena, even though their origins are not well understood. It says that if we order some large collection by size or popularity, the 2nd one will be about half of the first one in the measure we are using, the 3rd one will be about one third of the first one, and in general the one ranked k-th will be about $1/k$ of the first one. As one example, Wikipedia's entry for Zipf's Law cites a large body of English language text in which the most popular word, "the," "accounts for nearly 7% of all word occurrences," the second-place word, "of," makes up 3.5% of such occurrences, and the third-place "and" accounts for 2.8%. Although Zipf's Law was originally formulated by the linguist George Kingsley Zipf to apply just to word frequencies, it has been found to describe a far wider range of phenomena, such as wealth or income distributions, populations of cities, and popularity of blogs.

Now let us suppose that the incremental value that a person gets from other people being part of a network varies as Zipf's Law predicts. Let's further assume that for most people their most valuable communications are with colleagues, friends, and family, and the value of those communications is relatively fixed - it is set by the medium and our makeup as social beings. Then each member of a network with n participants derives value proportional to $\log(n)$, for $n \log(n)$ total value.

4 Other applications of Zipf's Law

Zipf's Law helps explain, in a quantitative form, several additional phenomena. In particular, Chris Anderson has noted (in the October 2004 issue of *Wired*) that online merchants, such as Rhapsody in music, and Amazon in books, derive a substantial fraction of their sales from the great number of items in their inventory that are bought relatively infrequently. His conclusion (stated in the subtitle of his article) is that online merchants should "[f]orget squeezing millions from a few megahits at the top of the charts. The future of entertainment is in the millions of niche markets at the shallow end of the bitstream." This conclusion is easy to justify using Zipf's Law, provided that we have enough items in our collection. By Zipf's Law, if value follows popularity, then the value of a collection of n items is proportional to $\log(n)$. If we have a billion items, then the most popular one thousand will contribute a third of the total value, the next million another third, and the remaining almost a billion the remaining third. But if online music stores such as Rhapsody or iTunes carry 735,000 titles while the traditional brick-and-mortar record store carries 20,000 titles, then the additional value of the "long tails" of the download services is only about 33% larger than that of record stores. Note that this view does not deny that blockbusters like "Titanic" or "Lord of the Rings" will continue to be very popular and gather huge revenues. It just says that they will form a decreasing fraction of the total value that society derives from the growing volume of material on the Internet.

Zipf's Law also explains another, closely related phenomenon, namely the seriously misleading preoccupation with "content," meaning material prepared by professionals for consumption by the wide public. The second author's "Content is not king" paper, published in the February 2001 issue of *First Monday,* shows that historically society has valued connectivity far more than content. Your phone call with your brother is of little interest to the billions of people in the world, but it matters a lot to you, and the collective value of all such calls outweighs the value of all of Hollywood's creations. And again, as the universe of users grows, total value grows.

5 Value of interconnection

Our proposed $n \log(n)$ valuation for a network of size n helps explain why network interconnection often requires time, effort, and in many cases regulatory pressure to achieve. It also suggests an explanation for why large networks often refuse to interconnect with smaller ones.

When Sarnoff's Law holds, and the value of a network grows linearly in its size, as it does for a broadcast network, there is no net gain in value from combining two networks. Mergers and acquisitions in such situations are likely driven by other factors, such as scale efficiencies in operations, or attempts to increase bargaining power in purchases of content. On the other hand, when the value of a network grows faster than linearly in its size, then generally (subject to some smoothness assumptions we will not discuss), there is a net gain from a merger or interconnection, as the value of the larger network that results is greater than the sum of the values of the two constituent pieces. We next consider how that gain is distributed among the customers of the two networks.

Let us assume that Metcalfe's Law holds, and we have two networks, call them A and B, with m and n customers, respectively, and that on average the customers are comparable. Then interconnection would provide each of the m customers of A with additional value n (assuming the constant of proportionality in Metcalfe's Law is 1), or a total added value of

$$mn$$

for all the customers of A. Similarly, each member of B would gain m in value, so all the customers of B would gain total value of

$$nm$$

from interconnection. Thus gains to customers of A and B would be equal, and the two ISPs should peer, if they are rational. However, the incentives are different if our $n \log(n)$ rule for network valuation holds. In that case, each of the m customers of A would gain value $\log(m + n) - \log(m)$ from interconnection, and so all the customers of A would gain in total

$$m(\log(m + n) - \log(m)).$$

On the other hand, the total gain to to customers of B would now be

$$n(\log(m + n) - \log(n)).$$

If m and n are not equal, this would no longer be the same as the total gain to to customers of A. As a simple example, if $m = 2^{20} = 1,048,576$, and B has 8 times as many customers as B, so $n = 2^{23} = 8,388,608$, then (again taking logarithms to base 2) we find that interconnection would increase the value of the service to A's customers by about 3,323,907, while B's customers would gain about 1,425,434. Thus the smaller network would gain more than twice as much as the larger one. This clearly reduces the incentive for the latter to interconnect without compensation. This is a very simplistic model of network interconnection, of course, and it does not deal with other important aspects that enter into actual negotiations, such as geographical span of network, and balance of outgoing and incoming traffic. All we are trying to do is show that is that there may be sound economic reasons for larger networks to demand payment for interconnection from smaller one, a very common phenomenon in real life.

6 Conclusions

Metcalfe's Law and Reed's Law both significantly overstate the value of a communication network. In their place we propose another rough rule, that the value of a network of size n grows like $n \log(n)$. This rule, while not meant to be exact, does appear to be consistent with historical behavior of networks with regard to interconnection, and it captures the advantage that general connectivity offers over broadcast networks that deliver content. It also helps explain the failure of the dot-com and telecom ventures, since it implies network effects are not as strong as had been hoped for.