

Is Metcalfe's Law Way Too Optimistic?

Joe Weinman

Is the value of a network really proportional to the square of its size? Or does each new node add limited value?

In July, 2005, Rupert Murdoch's News Corp. paid \$580 million for Intermix Media, the owner of MySpace.com, a social network of teens and musicians. In October, 2006, Google bought YouTube.com, a peer-to-peer video clip-sharing network of prosumers [1, see references, p. 27] (producers who are also consumers) for \$1.65 billion. And, two months later, AT&T paid \$86 billion to complete ownership of the wireless and wireline networks of BellSouth and Cingular.

Closing these deals required agreement between buyer and seller regarding fair market value. But what is the best way to value networks such as these? Cash flow? Assets? Revenues? Profitability? Subscribers? Average Revenue Per User (ARPU)?

And, however valued, as a network grows in size, does its value increase? If so, in what proportion relative to its size?

These are important questions, because networks are ubiquitous, whether they are telecommunications networks, transportation networks, neural networks, logistics networks, financial net-

works or social networks. Valuing them, both statically and accounting for future growth, is important for developing intelligent business strategies and making wise investment decisions.

Connectivity Value

One way to define a network's size is the number of nodes that it has, typically denoted by the variable n . These nodes may be Internet websites, MySpace users, cellular voice and data service subscribers, retail outlets, neurons, etc.

Consider a fully connected network, where every node can communicate or interact with every other node. An example of this would be a plain old telephony network, which allows each customer to call any of the others.

There are a variety of methods for valuing such a network, e.g., the value of collaboration and/or subset-forming ability, as Reed [2] has argued (see "Reed's Law"). But let us focus on the connectivity value of a network, which we define as being the sum, for each pair of nodes, of the value of connectivity between these nodes. For example, for a telephony network comprising three nodes, Alice, Bob and Charlie, the *connectivity value* of the overall network would be the value Alice and Bob derive from chatting, plus the value that Alice and Charlie get, plus the value that Bob and Charlie get.

Of course, defining the connectivity value of a

Joe Weinman is responsible for emerging services and business development for a large global telecommunications service provider. He has a BS and MS in Computer Science from Cornell University and UW-Madison respectively, and has completed Executive Education at the International Institute for Management Development in Lausanne. He has been awarded 10 U.S. and international patents, and is a frequent industry speaker globally. The views expressed are his own. He can be reached at joeweinman@gmail.com.

Reed's Law

This article is focused on *connectivity* value. Reed [2] has proposed that collaborative "group-forming value" supersedes connectivity value in importance, since the group-forming value of a network with n nodes is 2^n , because there are 2^n subgroups (subsets) of a set of n nodes.

Although beyond the scope of this article, Odlyzko and Tilly [5] point out the unlikelihood of this formula, since at some point, "addition of a single member to the network would have the effect of almost doubling the total economic value of the world."

While there is merit to considering not just connectivity value but also collaboration value among sets of nodes, it is likely that rapidly diminishing, convergent values of groups, coupled with intrinsic limits of consumption of the value generated by groups also creates a value-limiting effect. For example, Usenet, Yahoo! and Google each have thousands of groups. It is unlikely that every subscriber to these online services has joined all the available groups, finds them equally valuable or has enough time or attention span to fully digest whatever value may exist in them □

network in this way begs the question of what the connectivity value of each pair of nodes is. This can be fairly rigorously defined, depending on the network, as being the expected net present value of the stream of transactions between the pair of nodes. In other words, the aggregate value of the stream, adjusted for the time value of money and the likelihood that this value will actually be realized.

For example, in an on-line auction network, the value of CrazyCat being “connected” to SurferDude is the expected value of the stream of purchases that CrazyCat makes from SurferDude, recognizing that SurferDude selling \$100 worth of goods today is worth more than SurferDude selling \$100 a year from now, due to the time value of money. Also, if there is only a 10 percent chance that CrazyCat will actually buy (and that his payment will clear), then the expected value of that connection and its associated transaction stream is only \$10. For non-economic connectivity, e.g., calling a family member, we can convert the emotional or social value to a financial amount by standard techniques used in market research.

However, for the purposes of our discussion here, we don’t need to worry about the exact quantitative value of a connection. We will just assume that we can calculate it somehow...let’s call it v .

Robert Metcalfe, an inventor of Ethernet, is widely credited for what is now called [3] Metcalfe’s Law [4], namely that the value of a network is of order n^2 , that is, it grows as the square of the number of nodes. Using this formula, if a network doubles in size, it doesn’t just double in value, but quadruples. Under this hypothesis, combining or expanding networks can be a value creation engine, since, if a network becomes 10 times larger, it becomes worth 100 times as much.

Odlyzko and Tilly [5], together with Briscoe [6], provide support for a substantially more conservative calculation, namely that the value of a network grows in accordance with the formula $n \log(n)$. Without addressing their calculations in detail just yet, their conclusion would be that a network that doubles in size is worth, maybe not four times as much, but a little more than twice as much.

Linear Connectivity Value And Its Root Causes

However, many things in the world just grow linearly in value. For example, two ounces of gold are only worth twice as much as one ounce of gold. But gold is merely an aggregation, not a network. So what about for networks?

Linear value growth has traditionally been assumed, in accordance with Sarnoff’s Law, for the somewhat limited case of a broadcast network, where a single node, such as a television station, broadcasts content to all the other nodes (TV viewers). For example, a cable network might be valued at about \$4,000 per subscriber. So a net-

work with 100,000 subscribers would be worth \$400 million, and combining two of them adds linearly to create a network worth \$800 million.

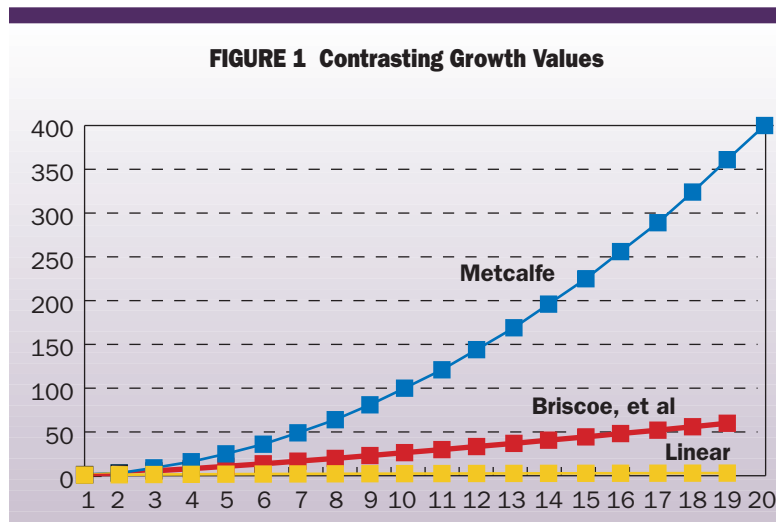
Surprisingly, however, value growth may be just linear, that is, of order n , not only for broadcast networks, but even in fully connected networks. In other words, there may be cases in which, when a fully connected network doubles in size, it just doubles in value. While such a perspective may appear somewhat Luddite, if not downright foolish, it is worth understanding if, when, and why it may be true, because n^2 , $n \log(n)$, and n are dramatically different calculations of value, as is shown in Figure 1 for a network growing from two nodes to 20. The top line shows value growth according to Metcalfe, the middle according to Briscoe et al., and the bottom shows linear growth.

These different kinds of value calculations would drive different investment, merger, and acquisition behavior, since, if this chart showed \$1,000 increments, the 20-node network would be worth somewhere between \$20,000 and \$400,000 depending on which “law” you use...a pretty wide spread!

Two mechanisms that can constrain connectivity value growth to be linear are *convergent value distributions and intrinsic limits of consumption*.

■ **Convergent value distributions**—A series, that is, the sum of a sequence of numbers, may be either *convergent* or *divergent*. For example, $1 + 1 + 1 + 1 \dots$ **diverges**, that is, gets larger and larger without limit, as you keep adding ones. On the other hand, $1 + 1/2 + 1/4 + 1/8 \dots$ **converges**: no matter how long you keep adding terms, in this case increasing powers of one-half (which decrease in size), the sum will never get larger than 2. When the distribution of pair-wise connectivity values derived from each connection is a **convergent** series, this places a constant bound on the value received by each node, which in turn causes the total connectivity value of the network to be linear in the number of nodes. It is not a

Value growth may just be linear, even in fully connected networks



Convergent value distributions and intrinsic limits of consumption can constrain the rate of value growth

requirement that each node have an identical value curve of pair-wise connectivity value. But, as a simplifying assumption, we consider the distribution of connectivity values to a “typical” specific node, then multiply by n to determine the total network value.

■ **Intrinsic limits of consumption of value**—Regardless of how much value might be “in” the network, bottlenecks in discovering, accessing and acquiring, or consuming that value may prevent all of the value from being realized by each node or user. Time is one limit, but so is capacity, e.g., the storage capacity of a user’s portable music player or personal video recorder.

Consider a telephony network. When it has 10 users on it, an individual can call them all, every day. When it has 100 users on it, she can still do that, although it becomes a little more difficult. When it has several *billion* users on it, this is impossible.

Moreover, even if one could call all several billion subscribers, say through some sort of Voice over IP War Dialer, it would be difficult to derive value from it, except in a *Guinness Book of World Records* kind of way. Hopefully, one enjoys—and perhaps needs to—chat with family members and work colleagues. There are a few times in one’s life where a call to the appliance repair person or the IRS auditor is warranted, but once a calling party goes beyond the several thousand people that she knows, or maybe the several tens of thousands of people that she might end up wanting or needing to know, there is no additional value from the other several billion users (see “Real-World Behavior”)

In simple terms, either you just don’t want to talk to everyone on the planet (unequal connection value in a convergent distribution), or even if you did, where would you find the time (intrinsic limits of consumption)? In practice, these two drivers may be interdependent: if you had more time, maybe you would talk to more people, and if there were more interesting people to talk to, maybe you would make the time.

Ultimately, the connectivity value of a network can be traced back to its value to each of its nodes or users, which also roughly corresponds to their willingness to pay, either in hard currency, or in time spent interacting over the network. If each user can hit an upper bound in value derived from the network—whether due to a convergent value distribution, an intrinsic limit of consumption or both—and this upper bound is insensitive to further network growth, then the value derived by each user is (bounded by) a constant. The total value derived from the network is then bounded by the number of nodes or users (i.e., n) times this constant. This in turn makes the total value of the network proportional to n , in other words, linear.

The strategic implications of this are potentially far reaching, because they suggest that after reaching a certain critical point, there no longer is

Real-World Behavior

A recent study [8] by the head of the Strategy Unit at Harvard Business School, Dr. Pankaj Ghemawat, demonstrates that, in the real world, interconnections of all sorts are extremely localized, whether they are direct investments, patents, immigration or management research. Oh yes, and even phone call revenues, which Ghemawat’s studies show are over 90 percent in-country. In other words, not only don’t we, on average, call many people—on average, we also don’t call very far away.

If Metcalfe’s Law were correct, then the reverse should be true. Since there are about 200 countries, if all connections were of equal value, then on average, 99.5 percent of calls should be out-of-country, and an even higher percentage of revenues should be, assuming international dial rates are higher than domestic. Even a populous country such as the U.S. only has about 5 percent of the world’s population. Metcalfe’s Law would predict that 95 percent of U.S.-originated calls should be international, which is far from the case □

value creation through consolidation. The actual point of diminishing returns for real networks will vary, but ultimately depends on where a distribution begins to converge or an intrinsic limit of consumption of value is reached.

Assessing Value Distributions

Figure 2a shows a fully connected network with 11 nodes. To aid in our analysis, we will consider such a network first from the perspective of a particular “typical” node (shown at the bottom), and its connections to the other 10 nodes. Let’s reconsider Metcalfe’s Law. It’s really stating that for a fully connected network, its value is its connectivity value, its number of nodes is n and connections is $[n*(n-1)]/2$, which is of order n^2 . Then, assuming that the value of each connection is some constant V (a significant assumption, but not necessarily correct for all networks), as shown in Figure 2b, the total connectivity value of the network would be proportional to Vn^2 , which, since V is a constant, is order n^2 .

Now, there is no argument that the number of logical connections in a fully connected mesh network is of order n^2 . But is the value of each connection identically V ?

Briscoe, Odlyzko and Tilly suggest that it is not, and mention that even Metcalfe calls that assumption into question. They propose that the value of pair-wise connections at each node drops off in proportion to a harmonic series: 1, 1/2, 1/3, 1/4, 1/5, and so on out to $1/n$, as shown in Figure

2c, p. 22. They point out that a number of real-world distributions, such as word frequency in the English language, have just this sort of distribution. Therefore, they argue, each connection is not worth some constant value V . Instead, the most important one is worth V , the second one is worth $V/2$, the next one is worth $V/3$, etc. Therefore, the

total connectivity value of the network is $n * (V + V/2 + V/3 + V/4 + V/5 + \dots V/n)$, which turns out to be mathematically equivalent to $n * V \log(n)$. Since V is a constant, the total connectivity value of the total network is order $n \log(n)$.

There are many other possible distributions of value, with varying implications as to the total

FIGURE 2a Fully Meshed Network Viewed From A Select Node


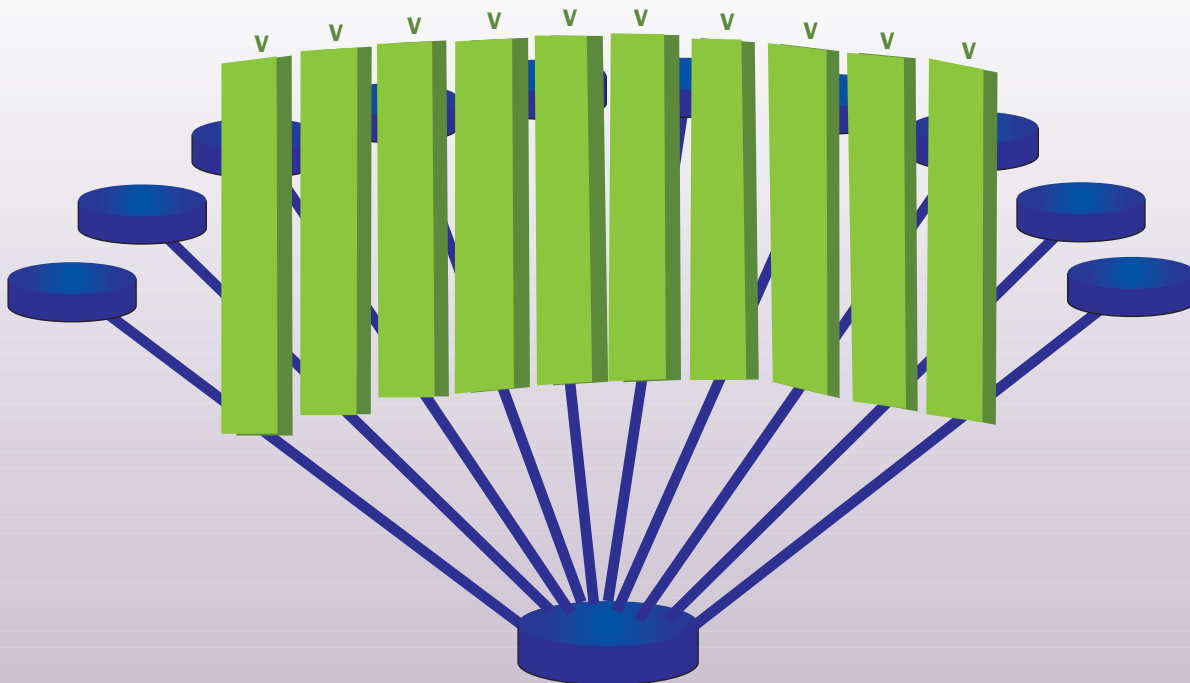


FIGURE 2b Equal-Value Connections



connectivity value of a network. For example, Figure 2d shows uniformly distributed values. Its mean is $1/2V$. Since we have n nodes generating n^2 connections, each of whose expected value is $1/2V$, this leads us to the same order of total connectivity value, so for these distributions Metcalfe's Law would correctly assess the total value of the network.

The problem with assuming these types of dis-

tributions that support either Metcalfe's Law or Briscoe, Odlyzko and Tilly's analysis is that they do not appear to match real world characteristics of many types of networks.

Long Tails Make A Big Difference

Let's return to the telephony example. There are roughly 3 billion cell phone users in the world. Metcalfe's Law assumes that each user gains

FIGURE 2c Harmonic Series Of Connection Values

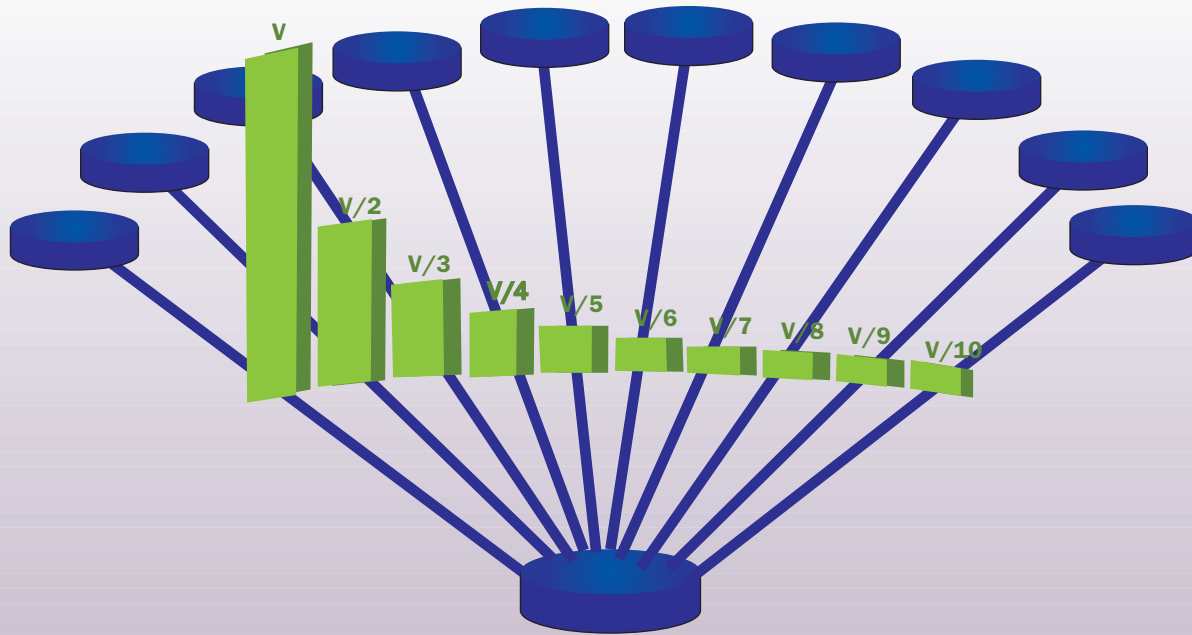



FIGURE 2d Uniformly Distributed Connection Values



equal value from each of those 3 billion. This seems overly optimistic.

Briscoe *et al.* assume in effect, that, from your perspective, you have a favorite person to call, and you talk to them say, 100 minutes a month. Then, following the harmonic series, you talk to your second favorite person $100/2 = 50$ minutes a month. Then, you talk to your third favorite person $100/3 = 33.3$ minutes a month. This is not implausible.

However, after, say, 50 people or 100 people, or even 1,000 people, you no longer call anyone “lower down” on the list. You probably do not speak to the 10,000th person for $100/10,000$ of a minute each month (which is .6 seconds), and certainly not the 10 millionth person for .6 milliseconds.

In fact, you just don’t talk to them at all. And, in the same way that Long Tails [7] can add substantial value, their absence can eliminate it.

This simple observation turns out to have significant implications. Even though .6 milliseconds either way wouldn’t appear to make much of a difference to the overall value of the network, it impacts the distribution of the long tail; in fact, it eliminates it. A harmonic series, proposed by Briscoe, Odlyzko and Tilly as characteristic of connection value distributions, “diverges,” that is, it will not only get larger, but will get larger than any predetermined limit. Although it gets larger more and more slowly, it will eventually break through any predetermined “barrier.” This means that a network which follows the assumptions of Briscoe, *et al.*, as it grows, will continue to have

value that grows faster than the size of the network, not unlike Metcalfe’s Law. Both analyses rely on potentially unrealistic assumptions about the distribution of value of individual connections, and an unbounded ability to consume that value, regardless of how large the network gets.


On the other hand, if the distribution of value is correctly modeled by a convergent series—either a finite-length, convergent infinite or truncated infinite series—whose value is insensitive to the size of the network; or if there is an intrinsic limit of consumption of value at each node (as in our cellular network example above), then the maximum value that may be consumed at each node is bounded by a constant. And, if this limit is a constant—let’s call it C —and unvarying regardless of the size of the network, then the maximum value of a network is Cn . Under these circumstances, the value of a network is just proportional to n . This is true regardless of the ordering of the sequence.

There are a number of ways in which the sum of connectivity value derived at each node can be bounded by a constant. One, as shown in Figure 3a, is where there is a predetermined maximum number of connections from each node with equal value, and all the rest have no value. To restate a critical point, this maximum number must be insensitive to the size of the network for this logic to apply.

Or, there may be an infinite convergent series—one with a long (but extremely thin!) tail of value—whose sum still is bounded by a constant. The series $1 + 1/2 + 1/4 + 1/8 + 1/16 + 1/32 + 1/64 \dots$ has just this property, as shown in Figure

Prior analytical approaches may not correctly model the characteristics of many real-world networks

FIGURE 3a Truncated Sequence of Equal Value Connections



3b. Its sum will never be larger than 2. Although this exact distribution of connection values is unlikely in the real world, it is included here for mathematical completeness.

If such a convergent series is truncated, as in Figure 3c, it remains convergent. Or, if a divergent series such as the harmonic series is truncated at a fixed point, also insensitive to the size of the network, as in Figure 3d, its total is also a constant.

As we've discussed, not only the value distribution, but also intrinsic limits of consumption may

play a role in limiting the rate of value growth. Figure 4 illustrates this. Even if we assume that there is plenty of value in the network, potentially unbounded as the network scales, there may be a bottleneck based on intrinsic limits of consumption causing limited value to be received by the user.

Value Maximization Strategies

Even if a network does demonstrate linear value growth, that does not mean that it is impossible to

FIGURE 3b Convergent Series Of Connection Values




FIGURE 3c Truncated Convergent Series Of Connection Values




FIGURE 3d Truncated Harmonic Series Of Connection Values

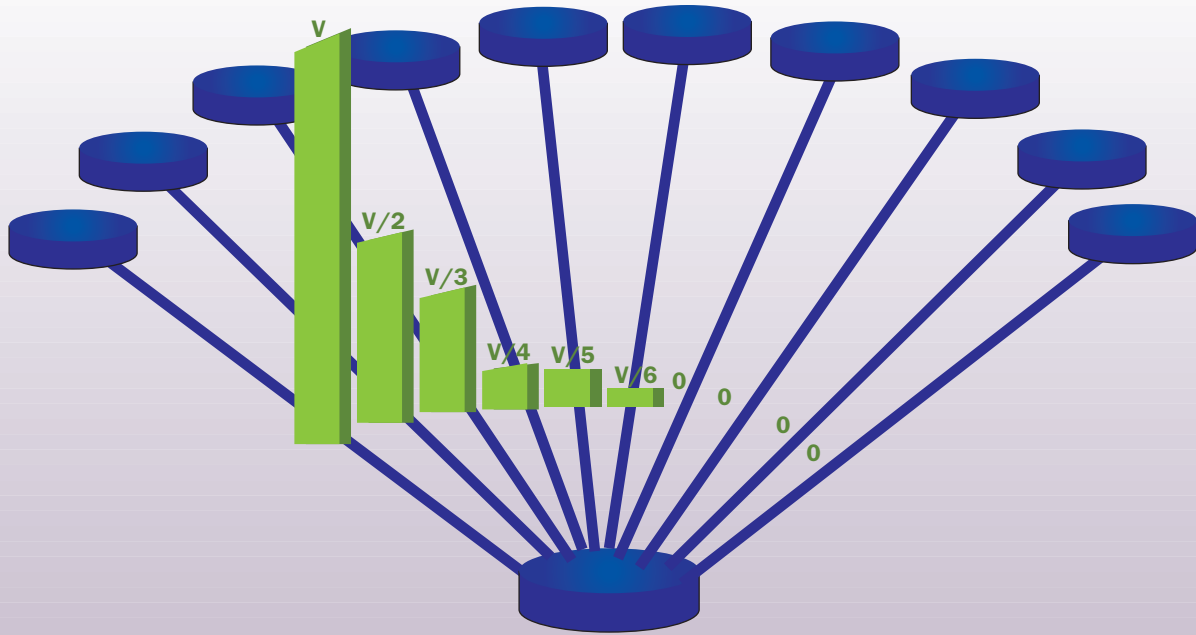




FIGURE 4 Plenty Of Value, But A Bottleneck In Consumption





A variety of strategies exists to accelerate the rate of network value growth

increase value to customers and therefore revenue as a network service provider. There are a few canonical strategies for service providers to use in such an environment:

■ **Increase the limits of consumption**—Even though an individual may only have time to physically watch, say, four hours of video each day, a personal video recorder can “consume” (i.e., record) more than that. This is not just “wasteful” consumption, but provides value by enabling more options for the consumer. In general, automation scales better and therefore tends to have fewer limits of consumption. But even individuals can increase their limit of consumption, e.g., by substituting one form of entertainment for another or multitasking.

■ **In a multi-constraint environment, optimize against all constraints**—In a consumer television network, customers may only watch four hours a day, but there may be more value—both in terms of customer flexibility and willingness to pay—if those four hours are on-demand services. In fact, there may be multiple limits of consumption. For example, someone may have an attention span constraint of TV watching of 4 hours per day, a constraint for typical broadcast fare of 2 hours a day, and a budget constraint for entertainment of \$200 per month. This person may be at one limit, e.g., already watching two hours per day of mainstream content, and having more options such as indie films down the long tail may allow him to reach four hours per day. Then, he may not be anywhere near the financial limit, e.g., I may only be spending \$70 per month on entertainment. Shifting the four hours a day from “free TV” to pay-per-view content keeps him within both limits, but enables the consumption of more value within the financial limit.

■ **Eliminate access and edge bottlenecks**—A limit of consumption might be due to bottlenecks at the edge of the network. Moving, for example, from dial-up access services to broadband DSL to fiber, or from GPRS (General Packet Radio Services) to EDGE (Enhanced Data rates for Global Evolution) to HSPA (High Speed Packet Access), enhances the ability to consume services such as streaming video.

■ **Eliminate core bottlenecks**—Although in converged telecommunications and video delivery networks engineered for scalability and peak loads it is rarer to find bottlenecks today, there are historical cases. A classic example is the first Victoria’s Secret fashion show to be broadcast over the Web in 1999. The Web servers crashed, limiting consumption to zero for thousands of consumers. Moving to a content delivery network the following year eliminated core bottlenecks.

■ **Increase the discoverable content available for consumption**—For services to be consumed, they must first be discovered. The use of search, program guides, collaborative filtering, recommendations, metadata tagging and similar tech-

niques makes more of the content that is generated available to be consumed. The appreciation in market value of some companies nominally focused on search highlights the value created by such a function.

■ **Create a Long Tail**—Chris Anderson [7] has asked “Does the Long Tail grow the pie or simply slice it differently?” In other words, does moving from a concentrated network to an expanded network increase the total value, or merely rearrange it? His research has indicated that the answer depends on many factors. However, since a finite, truncated value distribution clearly creates only linear value, creating a long tail probably can’t hurt (see “The Gem In The Long Tail”).

Conclusion

Assessing the value of real networks, as they grow or consolidate, is more complex than outlined here. One can look at the total value as being not just what is provided to users or subscribers, but also adjusted for the cost of operations and capital expenditures. Then, there may be value creation synergies through consolidation, marketing synergies or scale economies.

However, the point behind this article is not to

The Gem In The Long Tail

One of the arguments for network expansion and long tail effects is the possibility that, somewhere deep down in the long tail, one will find a gem: the perfect trippin’ track on an on-line music downloading service, or a book with a deeply insightful assessment of the human condition at an on-line book retailer. Therefore, the argument goes, substantial value is created by adding a track or book to the long tail.

While this may in fact happen, it does not invalidate the arguments here. The point of this analysis is that, when the shape of the value distribution has a certain curve, it is convergent, and then the total connectivity value of the network is proportional to n . Having something of surprisingly high value appear way down in the long tail—stated differently—means that there were plenty of things of low or zero value between the “hits” and there. As Chris Anderson [7] points out, “whatever you are looking for, there’s more stuff you aren’t looking for the farther you go down the Tail.”

The sum representing the partial series then tries to play “catch up” to an optimally sorted distribution. Even though they have the same terms, the partial sums of $1 + 1/2 + 1/4 + 0 + 0 + 0 + 0 + \dots$ get bigger faster than do the sums of $1 + 0 + 0 + 0 + 1/4 + 0 + 1/2 + \dots$ —but in any event, they have the same limit□

provide an exact quantitative analytical process for determining the value of a network, but rather to outline arguments as to why some networks may some of the time, if not all of the time, experience growth in value that is only linearly proportional to the growth in nodes.

There may be networks whose value grows as n^2 , $n \log(n)$, or even 2^n , as long as underlying assumptions regarding connection value distribution and nodal consumption hold regardless of how large n grows. For many networks, though—if not always at their early stages of growth then at least beyond some critical point—we can state a new “law:” The connectivity value of a network is linearly proportional to its size when either the distribution of connectivity values to each node is convergent or the ability to extract value from the network at each node is limited.

In the final analysis, there is no universal “law.” Rather, it is important to understand that rules of thumb are of limited value, and that one must recognize the implications that different distributions of value and limits of consumption have on total network value □

References

1. Don Tapscott and Anthony D. Williams, *Wikinomics: How Mass Collaboration Changes Everything*, Penguin

Books, 2006.

2. David P. Reed, “The Law of the Pack,” *Harvard Business Review*, February 2001.

3. George Gilder, “Telecosm: How Infinite Bandwidth Will Revolutionize our World,” *The Free Press*, 2000.

4. Robert Metcalfe, “Metcalfe’s Law: A network becomes more valuable as it reaches more users,” *Infoworld*, Oct. 2, 1995.

5. Andrew Odlyzko, and Benjamin Tilly, “A refutation of Metcalfe’s Law and a better estimate for the value of networks and network interconnections,” available at: <http://www.dtc.umn.edu/~odlyzko/doc/metcalfe.pdf>.

6. Bob Briscoe, Andrew Odlyzko and Benjamin Tilly, “Metcalfe’s Law is Wrong,” *IEEE Spectrum*, July 2006, pps. 26–31.

7. Chris Anderson, *The Long Tail: Why the Future of Business is Selling Less of More*, Hyperion, 2006.

8. Pankaj Ghemawat, “Why The World Isn’t Flat,” *Foreign Policy*, March/April 2007.



In the final analysis, there is no universal “law”

Companies Mentioned In This Article

AT&T (www.att.com)

Fox Interactive Media (www.fox.com)

Google (www.google.com)

MySpace (www.myspace.com)

YouTube (www.youtube.com)