

ICCE 2012: A Software Perspective

I'm a software guy who thinks in terms of market forces. As I listened to Dr. Samueli's talk at the 2012 ICCE conference I was frustrated.

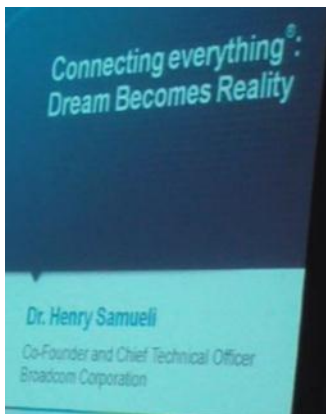
While I can quibble over particular errors of fact, after writing these comments my primary difference is a matter of opinion as to the forces that drive markets. This is a substantial disagreement but not a personal one.

While pairwise connectivity is indeed important, it is only a low level basis for the real application driver which is table pair relationships between arbitrary end points without regard to the path the bits take between them.

I'll admit there is also an emotional factor having lived through much of this history and thus being very aware of the details.

My goal here is to give others, especially those attending the ICCE, a sense of an industry in transition as the IEEE moves from a world of hardware and electrons to one in which a device is defined by how it is used and the software that is run.


My primary techno-economic point is that you get Moore's law benefits when you decouple markets, be it hardware/software or bits/applications (IP/TCP).



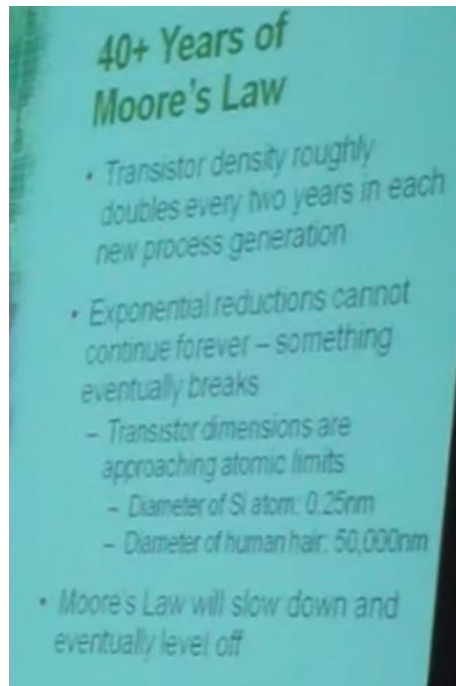
I agree with the spirit of connecting everything but the hardware aspects are only a minor part of the larger goal of relationships between people, devices and places. The chips cited only provide a single link in a far longer chain.

In the talk Dr. Samueli mentioned the importance of both the Electrical and Electronic engineers in the IEEE. It is a reminder that the organization itself hasn't quite come to terms with a software-redefined world.

While Marconi did play an important role in the [history of "wireless"](#) let's not forget others. As a kid I was fascinated by Nikola Tesla so I can't help but note that, eventually, the US Supreme Court awarded the radio patents to Tesla (alas, after he died).



Moving on to Shockley et al; the transistor was indeed important but it wasn't the start of electronics. We already had sophisticated electronic systems and digital signaling long before the transistor using tubes, relays and even player piano rolls (used by Hedy Lamarr for controlling frequency hopping in Torpedoes – an idea the Navy was too chauvinistic to accept).



The Intel 4004 was indeed a very interesting single chip CPU but let's not forget that it was prototyped in medium scale integration by Datapoint first. In a sense the software for the chip preceded the large scale integration.

This, of course, brings us to Moore's law – a topic in which I have a very strong interest. In

1997 I was asked to write about the limits of Moore's law and concluded that the law had little to do with physics and everything to do with markets.

As I wrote in [Beyond Limits](#), if you drill down you'll find all these neat learning curves are no longer smooth once you look at them close up. There is a different process at each point. Moore's law style hypergrowth can emerge when you decouple markets so that each can evolve on its own terms and you declare any success a victory. Thus a video board extends Moore's law to teraflops for home computers but those measures only apply to a certain class of problems. The technology may reach limits by one measure but can continue by other measures as long as we don't narrow our definition of success.


If, as in the talk, you focus narrowly on chip density you will indeed run into limits be they due to scale, heat or whatever. But if you take the same technology and decom-

pose a CPU into a standard computation portion and a portion that is tuned for graphics you get much more effective capability for large classes of applications with a graphical UI.

The increased cost of fabricating these chips leads to more and more concentration in the choices of chips. This makes it even more important to understand the interplay of hardware and software lest we be at the mercy of hardware designers doing us favors we can't afford and find harder to program around.


We already see this in cell phones and Bluetooth that have embedded smarts in the lowest levels of hardware and behind certification walls.

Samueli then goes onto to tell about how the ple][(mentioned Jobs but not Wozniak) and the IBM PC transformed the world without any mention of the role of software. In fact the Apple was failing into the market until [VisiCalc](#) re-defined it as a business machine.



This is backwards. IBM did give the PC added credibility but IBM was playing catch-up to a business market that had already started to deploy personal computers. Note that IBM wasn't just doing mainframes – it was already doing small specialized computers such as an APL machine.

He's far from alone in confusing the ARPANET with the Internet but there are fundamental differences between the




ARPANET which is a classic network and the Internet. The differences are at the heart of the idea of a network in which there are limiting policies and the Internet which is about opportunity¹. The ARPANET is a classic network that takes care of functions like reliable delivery in the network.

The Internet demonstrates a fundamentally different concept in which there need not be a network at all – just a best efforts means of exchanging bits. Following the de-

coupling principle that drives hypergrowth the sharp split between the exchange of bits and what we do with the bits is key to the power of the Internet as a transitional force. And that's just with today's prototype let alone what is possible in the future.

A quibble in that he says that the IMP was created by Honeywell. The DDP-516 was a hardened process control computer on the market. It was programmed to act as a router. "Act as" is fundamental to a software-defined world. It's a very different notion than a hardware manufacturer telling us the purpose of a chip.

The more important point is the claim that "Internet communications speeds evolved much more slowly than exponential Moore's Law model. Again this is a result of policy and not technology. The bits traversed the old creaky telecommunications infrastructure much like the Acela train slows down as it uses ancient tracks. As he says in the talk, we were using leased telephone lines.



Important Note: In <http://rmf.vc/Capacity> I explain that using the term "speed" is simply misleading. We don't say a 100Amp circuit is faster than a 10Amp circuit. It simply carries more electrons per second. The number of bits per second is a capacity measure. I'll continue to use the word "speed" because we're used to it is no more about how fast electrons move than a light year is a measure of time.

Back in May 1973 Bob Metcalfe was running 3ⁱⁱMbps on a coax. Let's step back a few decades and think about the capacity of television signals. We were already going megabits in the 1930's without wires.

The reason that upgrades were, and still are, expensive is because the network is coupled to the applications – the so-called smart network. As soon as we remove the coupling the speed shoots up beyond classic Moore's law cost/benefit rates.

Describing a communications infrastructure as an end-to-end link gets the Internet meaning of end-to-end completely backwards. Since this is fundamental to understanding the Internet it is vital to understand the difference between a traditional communications network and the Internet.

A telecommunications provider provides a valuable service like reliable delivery with guaranteed performance. This is what I call “womb-to-tomb” in the sense that everything along the path is managed by the provider.

The Internet is a very different concept in which the applications are created outside of a network. The meaning is only at the end points, hence end-to-end in the sense of “without the middle”. The major insight is that you can have reliable deliver (as with TCP) without a reliable transport. You only require “best efforts” which means that the path provider isn’t purposely preventing you from exchanging bits and is not favoring some traffic over others.

This creates a problem for today’s concept of telecommunications which is funded as a service (as I explain in [Thinking Outside The Pipe](#) and other essays). For the purposes of hardware vs. software the value comes from what we do with the network and the hardware. We see this repeated in network policies which limit what we do with the transport. Notice I’m trying to avoid even the word network since we no longer need a network provider to assure all the bits get delivered through a particular path.

One example is that a provider can take a copper wire and guarantee it operates at 56Kbps but can’t guarantee more so limits the user to what is guaranteed. With the Internet



approach we can let the user discover what works without setting an upper bound.

The speed of modems is a dramatic example of the impact of policy. I appreciate that he uses an Anderson-Jacobson ADC 260 acoustic coupler for the illustration; the very same model I got in 1966 and which is sitting near me as I type this.

The reason that the speed topped out at 56Kbps is that that is the carrier-defined speed for their “voice-band” provisioned network. It is clearly not the speed due to physics because the same wire that runs at 56Kbps for modems runs at megabits for 1980 vintage DSL. Imagine the speed if carriers continued to invest in taking advantage of the copper infrastructure

which reaches just about every home in the country!

Again, it is policy that is limiting us, not physics. Thanks to software we were able




to repurpose a voice network for data but only to an extent.

Even though I do disagree with many points I appreciate the enthusiasm for building chips that enable us to exchange bits at high speed without building in knowledge of the particular applications (putting aside QoS issues). This allowed us to use the bits as a plentiful “raw material” much like moveable type liberated the alphabet. Low cost is indeed important.

Too bad Samueli uses the term [Broadband](#) which confuses technology with a business modelⁱⁱⁱ. Almost all of the bits on a coaxial broadband cable as provisioned today are reserved for use by the service provider and not available for other applications. This is also true for fiber as we see with Verizon’s FiOS.

As I explain in my CES [debrief](#) the content business and the transport business are about to diverge and we’ll need to face up to a native business model for the raw copper, fiber and radios we use to exchange bits. This is going to have a major impact on Broadcom’s business but isn’t mentioned at all in the talk. In a sense it’s like the RISC (Reduced Instruction Set Computer) vs. the CISC (complex) as in Arm vs. x86.



The resolution change here is

another good example of policy. As we move towards a network of bits we no longer require a new infrastructure for content-centric changes such as new video formats. We no longer need different chips for each resolution – just more performance to handle more bits.

As an aside, in 1995 when decoupling home networks from the telecommunications business model (and thus enabling them to go at the gigabit speeds we are used to today) I chose phone wires as a “just works” medium because it was available but it was only an option and it didn’t really matter how fast it went.


Later when working with Tut we found that another company was also pursuing using the phone wire. The difference is that they tried to make it work as a home entertainment medium. Alas,

Broadcom paid hundred million or more for the technology which, as far as I know, never

had a serious market. Perhaps I’m jealous. Later I chose not to invest in Sand Video which Broadcom bought. Maybe one day I’ll find something to sell to Broadcom.



While it's interesting to see the evolution of the set top box it's an idea that has little future as the smarts move into boxes like Google's, Apple's, Microsoft's and others with content coming over IP rather than through a complex analog cable.



I compare analog television (including digital over analog cables) to sending email by fax and

then OCRing it back. I actually implemented this as a demo. You can make it sort of work but why bother? You don't need five tuners when you can support an essentially unlimited number of streams over IP.

We then get to "accessing the Internet" using broadband modems. This is a good example of repurposing an existing facility – we need to be careful to understand we have connectivity using the broadband technology but despite the broadband business model. This too is personal as these bits were connected to home networks. To emphasize the drift in meaning we use "broadband" for DSL, FiOS, LTE etc.


The important point is that these technologies repurposed existing infrastructure. We exchange bits despite the prior purpose just like dialup modems repurposed the phone network. And, as I've said many times, when you repurpose a video distribution system you find it's good for video, especially cute kittens on YouTube.



While he started to talk about the home in terms of video distribution he rapidly got to saying it's just a networked home which is closer to my take on it though I argue that the distinction between the Internet out there and the network inside is, again, a policy distinction. Too bad many people treat the distinction as an important boundary.

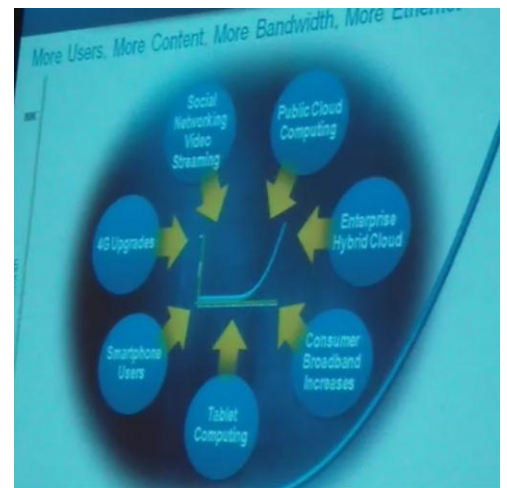
He then talks about the global infrastructure and switching in the transport going to a service provider network. True we are no longer switching connections. Less obvious is that the very idea of a network as a provided service is at odds with the Internet's edge-to-edge concept. See [Thinking Outside The Pipe](#) and other essays.

The inability to get the benefit of the value created becomes a problem if we expect carriers to invest hundreds of billions of dollars while all the value is created outside of a network. In the 1990's Lucent found no one was buying expensive PBXes. This shift in value is soon going to happen with networking gear.



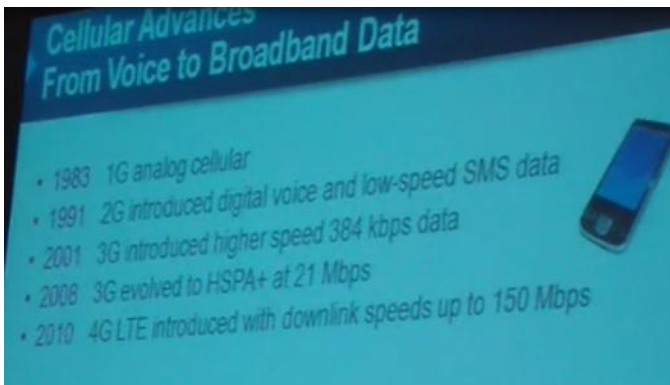
When he speaks about the improvement in Ethernet capabilities over the last 20 years he does recognize it happens because we have direct access to the physical infrastructure. (Should be 40 years but who's counting?)

This is a very important lesson – when we aren't depending upon intermediaries (as in decoupling or edge-to-edge (to avoid confusion over end-to-end)) we get dramatic price/performance improvements. To make this more clear, in the 1990's we had structured wiring which required coax and fiber be run throughout the house. But the market went with copper, taking it to speeds that weren't expected (and many thought impossible!) because it was better to build on what was available.



Outside the home where the carriers benefit from spending many billions of dollars on infrastructure they could pay for high priced fiber – at least until the market figures out that the value is not in the network but outside in the value created using the network.

While we can expect traffic to increase dramatically it isn't clear that all that traffic will be going through a small number of points nor that it will continue as if the Internet is about nothing but more and more video. Still, it is nice to know we still have a lot of latent capacity in all the unused fiber and copper and radio space around us. Again, all that capacity is available because of policy (and cute kittens) and not because of technology.



Just as the capacity of modems was limited by policy it is policy that limits the wireless capacity by segregating wireless bits from wired bits. Today's cellular protocols are burdened by billing at layer zero.

As an aside it's interesting how Broadcom and Qualcomm seem to be competing twins separated at birth (two UCLA grads) which are going to increasingly face each other in



the market.

I essentially agree that the Ethernet (or what we call Ethernet) is becoming the backbone networking choice in cars though there are many obstacles. The real challenge comes when we try to connect the car to the rest of the world and we're back in 19th century telecom.

He concludes by saying that connectivity everything has become a reality and I agree but I think the software that has driven the market with hardware playing a vital role in assisting.

Post Script – Open vs. Closed

After writing this essay I read Walter Isaacson's book on Steve Jobs and the issue of open/closed architecture is front and center. It's another take on these same issues. To the extent that value is intrinsic in a piece of hardware the closed, or whole system approach, has a major advantage.

But, in the long term, such approaches make it more difficult for society to reinvent. We need both approaches but let's not let the advantages of specialization in the short term evolution of systems blind us to the longer term necessity for reinvention.

ⁱ For more on opportunity see <http://rmf.vc/PurposeVsDiscovery>.

ⁱⁱ He is wont to remind me that his class project ran at 2.994Mbps.

ⁱⁱⁱ For more on this see <http://rmf.vc/BroadbandInternet> and other writings.