***The Chip Insider®***

***April 15, 2016 – Strategy and Tactics:*** *The new semiconductor market structure and what it means competitively.* ***WildPhotons****: Death and Taxes (it's tax day in America). . .* *By G Dan Hutcheson*

**Focal Points:**

The new semiconductor market structure and what it means competitively

- The new driver is the emergence of heterogeneous integration
- Structures effect on the fall and rise of giants
- Back to the future with horizontal market structures
    - A historical walk through semiconductor market structures
        - How companies sold chips then, today, and tomorrow
    - Today's vertical structure evolved
        - Looking for volume
        - Intel was the first vertical
        - Emergence of the Fabless/Foundry/IP Structure
- Electronics OEM Vertical Reemergence
- **Foundry, Subcon, IP, EDA Ecosystem - Platform Horizontals**
- The future of integration …
    - It's heterogeneous
- How the new structure changes who wins and who loses.

**The new semiconductor market structure, the return of horizontals, and the strategic opportunities:** The structure of the semiconductor industry is changing again and with that will come new strategic opportunities for chip companies. The new driver is the emergence of heterogeneous integration via advanced packaging. Every time this has happened in the past, it has resulted in the fall of existing giants that did not see it and the rise of new giants in their place. To better understand this, let me walk through how it has evolved strategically over time.

The earliest market structure for semiconductors was horizontal. Back then there were three tiers of the semiconductor supply chain: 1) Electronics companies, who controlled architecture, 2) Distributors who sold to them, and 3) Chip companies, who sold functional building blocks in packages. Electronics companies approached buying like going to a hardware store. There were more than 5000 of them in the U.S. alone, so distributors ran the semiconductor store and chip companies lined up to get their chips in the distributors' catalogs. The key issue was that all these functional blocks fit together, so distributors' sales efforts favored chip companies that

offered a wide spectrum of building blocks. Back then all chip companies designed their own chips. There was an electrical diagram, but the real work started at the physical layer, with the hand cutting of mask templates from rubylith with a knife. They also owned their own fabs, which was key to making sure everything worked together. Fairchild, GE, Motorola, Philips, RCA, and Texas Instruments were the dominant companies of their day. The one vertical at the time came in the form of electronics companies. They were called captives because they designed their own chips and fabricated them on their own fabs. IBM and AT&T are the classic examples, but almost all auto, computer, telecom companies of significant size owned their own fabs. This was easy to do, because fabs were cheap compared to the giants of their day.

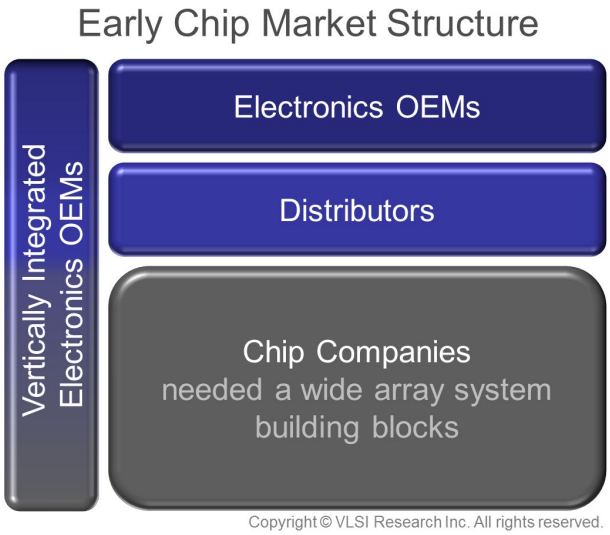


Then a vertical structure began to evolve. Like all market disruptions, it began with small start-ups looking for a blue ocean to get around the distributors, which were controlled by Fairchild, GE, Motorola, Philips, RCA, Texas Instruments and others. Start-ups didn't have the capital to compete directly, but they needed volume to survive. Intel was the first vertical chip company. It chose its market carefully, using the "Goldilocks Strategy," developed by Gordon Moore. First it focused on memory, which would quickly evolve into the first high-volume commodity. Then came MPU, MCU, DSP, DRAM, NAND, etc.  Intel was built focusing on verticals. And so were the Japanese giants that followed them.

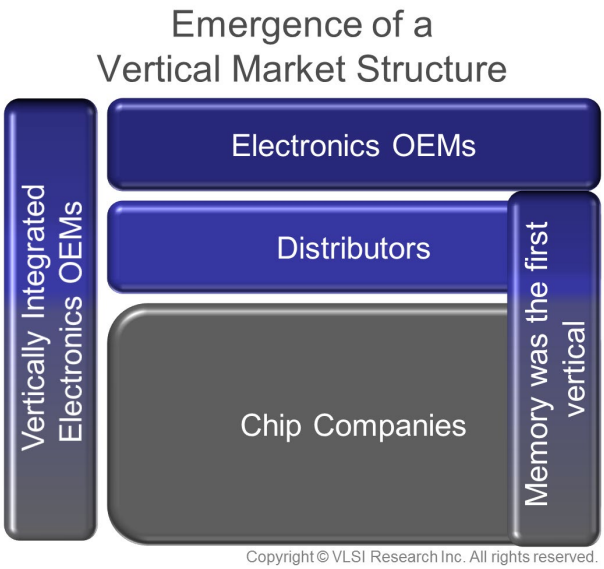


And then it got complicated with the Emergence of the Fabless/Foundry/IP Structure. This new structure was again driven by scale and software — in this case, the rising cost of wafer fab and the emergence of EDA with hierarchical design techniques enabled the fabless emergence. An explosion of small verticals emerged with each company looking to avoid owning a fab, which could not be afforded with a single vertical. Foundries became the new horizontal. They were looking to fill a fab w/o a distribution channel. As the SoC era came, the emergence of IP verticals became essential to the fabless emergence, as well. They provided soft functional blocks and that combined with EDA, allowed them to glue it all together. The result was a significant competitive threat as new market leaders, such as Broadcom, Qualcomm, and TSMC emerged.

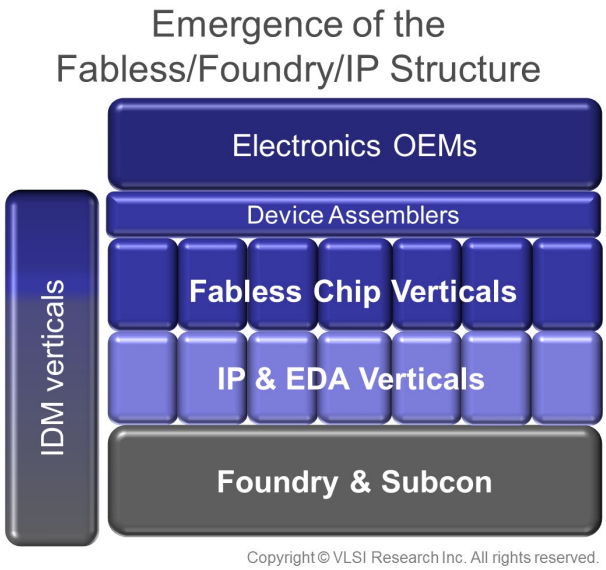


This would also enable a reemergence of OEM verticals. This was especially true as device assemblers made it possible for OEMs to be only designers and marketeers. First, the big

electronics companies came back. Cisco and Nokia were early entrants, as were ASUS and Logitech. Then came Apple, which made it clear. What they wanted was control over IP integration. They wanted to lock-in differentiation, as well as build secure platforms. That made the market very crowded, which was a consolidation catalyst.

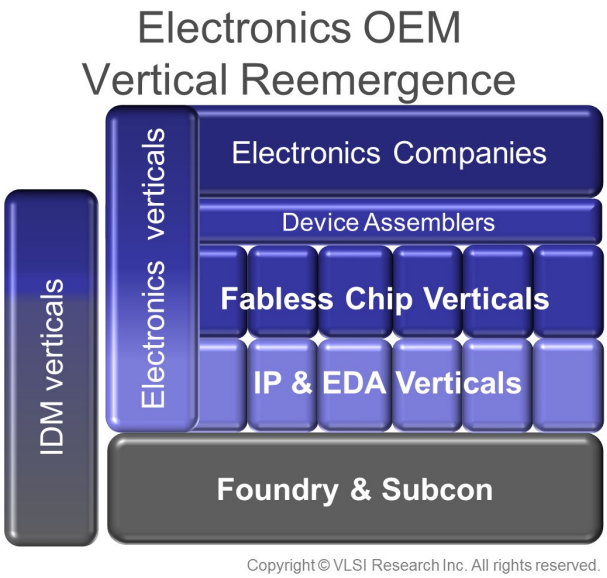


Market horizontals began to return, as all of this was happening. It was first driven by the 130nm node, which was the first to need DFM for tapeouts to be successful. Design rules began to explode.

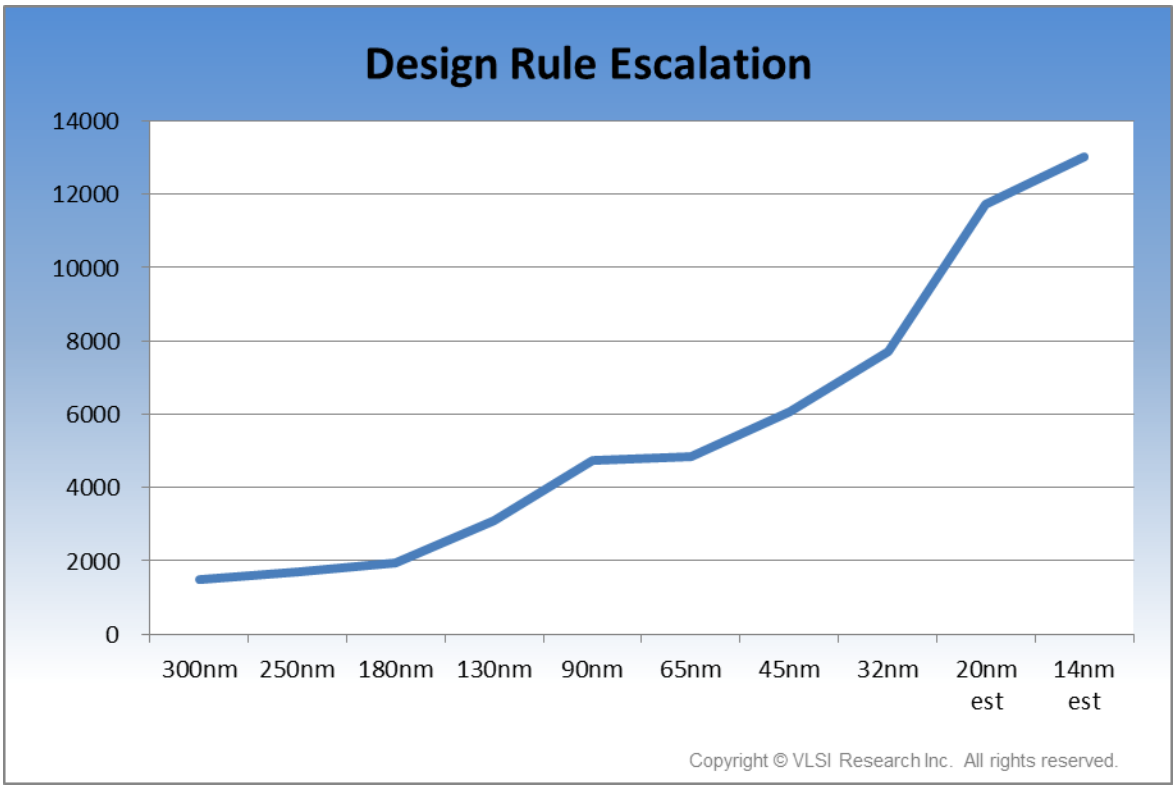


As a result, ecosystems became important. Though it would be a while before we called them ecosystems, what was important was the recognition that that a foundry was at the center of each. Technology platforms were even more important, as IP had to be fully qualified to run on each foundry's process. The growing size of SoC density forced this horizontal. Much like the early industry, breadth became an important factor. But this time it was IP breadth … where more was better. Today the breadth of a platform is a critical decision factor for any fabless design commitment.

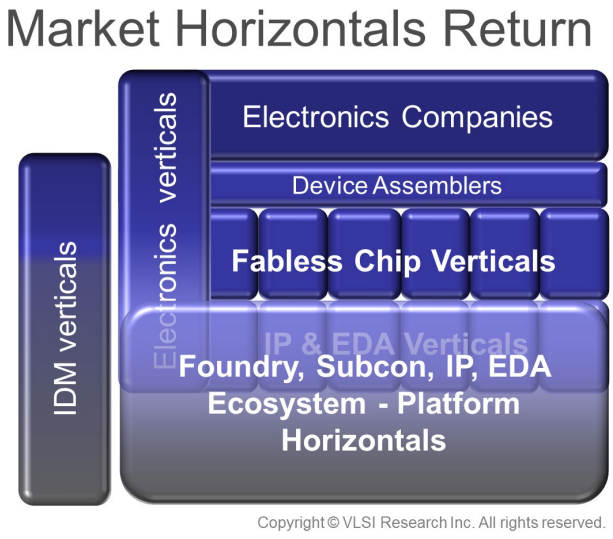

The fact is that we are well on the way to 100B transistor designs. Moore's Law is far from dead, as density is enabling new market approaches. This creates new scale issues in the form of NRE cost amortization. Currently the number of Designs is a third of what it was two decades ago, while design spending, at roughly $34B, is six times what it was then. So markets have to be bigger. So new tactics like Redundant Logic to increase yields, for example, are emerging. Here, instead of cutting out blocks of memory to have a sellable chip, you're cutting out core in memory. I believe Redundant Logic will become as common as it is in memory. Why, because as Paul Otellini put it, "If you understand anything about Moore's Law … You understand that eventually… Everything gets integrated."

This is where Heterogeneous integration comes in to change things. Integration must branch out in different directions: for differentiation and to control cost. Planar integration is only effective when added structures don't add process complexity. It's already branched out horizontally with 2.5D packaging and is going vertical where costs allow.

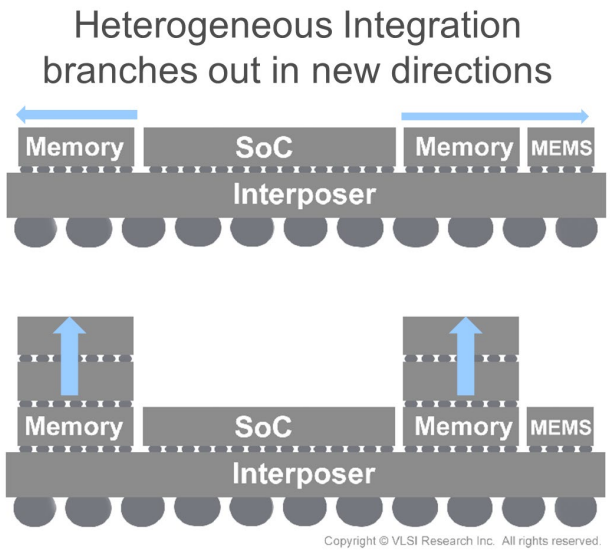


The future of integration is heterogeneous. A new future of competitiveness will likely hinge on having the best heterogeneous platform. It will be a new horizontal that you can already see TSMC putting together. Commanding horizontals across IP and CHIP platforms will matter most. Competitive advantage will come from platforms that access the best sensors and memory. Especially in a power stingy … performance driven world. Memory that has high bandwidth and low power drain will be prized. That puts Samsung, and potentially Intel, in the best position.

## The future of integration …

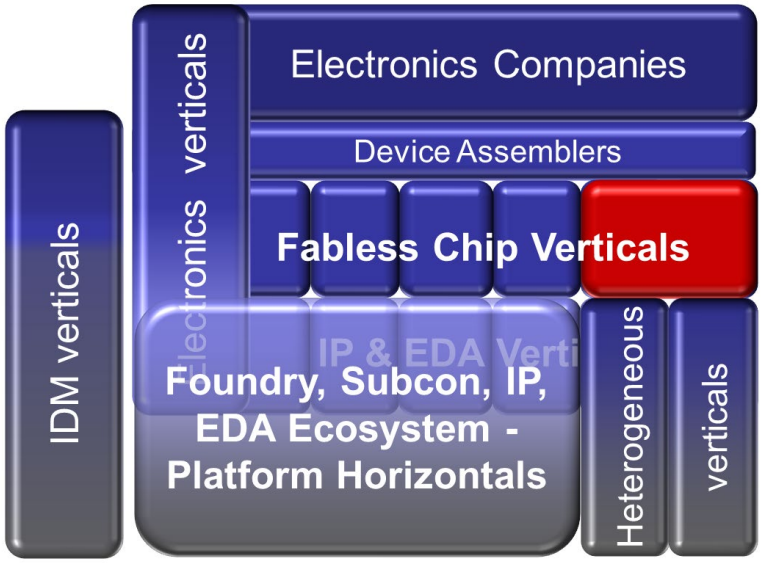




## WildPhotons: light life lessons

The difference between death and taxes is that
death doesn't get worse every time Congress meets
Will Rodgers

Old barn in snowstorm near Crater Lake, Reference number: NCal_0712k_274

follow Dan Hutcheson's photos on twitter @wildphotons
see them all at flickr.com/photos/wildphotons/

***July 1, 2016 – Strategy and Tactics:*** *The Cook's Tour: The IBM Research Semiconductor Group* ***WildPhotons****: The present is the new future. . .*

**Focal Points:**

The Cook's Tour:

- The IBM Research Semiconductor Group
  - They spun-off its Microelectronics
    - But kept research
  - What IBM is thinking, doing, and the role they expect to play in the future
- Does this high horsepower engine have the drivetrain to move systems down the road?
  - Is this just another toothless fab-lite?
  - Or is IBM inventing a new research model?
    - IBM is often so far ahead that it looks behind
- A tour of the Albany, NY fab
- Why IBM Semiconductor Technology Research represents a new R&D Model
  - The differences between it, the fab-lite research model, and consortia models
- Why systems companies need a deep understanding of devices and interconnect
  - Reverse engineering IBM's vertical research model
  - Apple, Google, and Facebook contrasts
    - How they are evolving their research models
  - IBM's ROI on research

**The Cook's Tour:** The IBM Research Semiconductor Group is one of only three logic research centers left in the world that does credible broad-spectrum device research. This one's a jewel with a heritage. Older than Moore's Law, IBM's researchers have led the semiconductor industry for decades. So IBM's decision to keep semiconductor research, as it spun-off its Microelectronics division last year, created a question of conflicting track records. It was IBM's versus the dim record of other IDMs going fabless and keeping research. Fab-lite R&D efforts to stay on the leading edge have typically built a process to nowhere. Keeping semiconductor R&D

internal while shifting to a foundry for manufacturing may placate shareholders, but it fails when it is an engine without a drive train.

At the same time, IBM is often so far ahead that it looks behind. Sustaining IBM's heritage of leading the industry through its technical insights is critical, not only to their partners, but also to the world. It has always been a great corporate citizen, publishing more insightful research than anyone over the decades. IBM Research was lighting the roadmap's path decades before there was a roadmap. So losing it would be a great loss.

So I was excited to get the invitation to visit. The goal was to find out what IBM is thinking, doing, and the role they expect to play in the future. That meant that this would be no normal Cook's Tour that relied mostly on critiquing what I saw. Instead, this one is more about what I heard.

There is no question that IBM researchers can be innovative. The critical questions are: does this high horsepower engine have the drivetrain to move systems down the road? Is this just another toothless fab-lite … where the R&D capability will quickly fade away? Or … as it has often done in the past … is IBM inventing a new research model?

The first time I witnessed IBM change the research model was back in 1992 when they introduced true partnering with competitors. IBM put together a $1B effort between themselves, Siemens, and Toshiba to develop next generation DRAM technology. It was the first major JDA to successfully bring top-line category, leading-edge products to market: 64Mb and 256Mb DRAMs. Moreover, they successfully overcame major cultural difficulties in getting American, Japanese, and German researchers to work together productively — something few thought possible at the time. The $1B spent collectively saved around $2B in R&D.

This was taken to a higher level, just four days after the 10$^{th}$ anniversary of the IBM, Siemens, and Toshiba partnership — when IBM was again at the center of partnering, with the announcement of the Albany Research Hub. John Kelly, who was director of research at IBM then, realized his vision that government and academia could play a significant role with industry in furthering the scientific research essential to the future of semiconductors while delivering significant economic benefits to the parties involved. In 2009, he would receive the semiconductor industry's highest honor, the Robert N. Noyce Award, for his accomplishments.

But technology never sleeps as it intertwines with economics to make advancements in semiconductors evermore expensive. The value is always there, but it does not come cheap. IBM had beaten the bullet longer than anyone due to its raw innovativeness and the fact that it was not an IDM … It was an Integrated Systems Manufacturer (ISM). The growing problem IBM would face in the 2010 decade was that, while the research was achievable, the scale needed to manufacture semiconductors was not.

Wafer fab ROI is extremely dependent not only on its initial cost, but also on the ability to fill it with large volumes of wafers passing through every month. And that's not just because

depreciation eats up margins as most think. Fewer wafers mean fewer cycles of learning, which is as big an issue because it creates a steadily growing yield handicap against competitive fabs running higher volumes of wafers. One of the reasons why fabless companies grew so big was that the emergence of the foundry business model solved the IDM's problem of needing to fill a fab. The high fabless-to-foundry ratio meant that a foundry didn't need to find markets large enough to fill its fabs. Somebody else was finding the markets for the foundries, while the lack of a need to fill a fab meant that fabless companies could be more creative in the markets they addressed – thereby creating steady demand for large volumes of wafers. It was a virtuous relationship. But it also drove the need for even greater scale to be successful.

That led to what must have been the hard decision to spin off its microelectronics manufacturing capability to its partner, GLOBALFOUNDRIES. With a company as forward thinking as IBM, its choice to keep a semiconductor research group leaves open the critical question asked above: Is IBM inventing a new research model?

First can they do real research without a real[1] fab? Fab-lite versions tend to be capital constrained, as the R&D is more for financial optics than products. What's different here is that IBM is still a critical research source for its chip making partners and the systems designers within the company.

Walking around the Albany, NY fab, I would say they definitely can do real research here. IBM researchers have access to the most cutting-edge tool sets in the world via their partnership with the State of New York and SUNY. IBM's own EUV lithography tool and track have been operational here since 2014 and they were finishing their latest upgrades, for example.


The only thing missing is the automation systems that only get in the way inside a real research lab[1]. But the tool sets they are using are in fact cutting edge, thanks to the fact that so many chip equipment companies are doing research in the Albany Research Hub as well. The tools themselves are not research playthings. They are the same fully-automated tools that can be found, or will soon be found, in leading-edge manufacturing fabs around the world. So the fab runs like a small manufacturing operation. The evidence is that there are few people around the tools.

Information turns are of critical importance to learning cycles. Here, they are capable of achieving 1 day-per-mask-layer and 8 weeks for a full-flow wafer lot. In other words, this fab runs research like a manufacturing line.

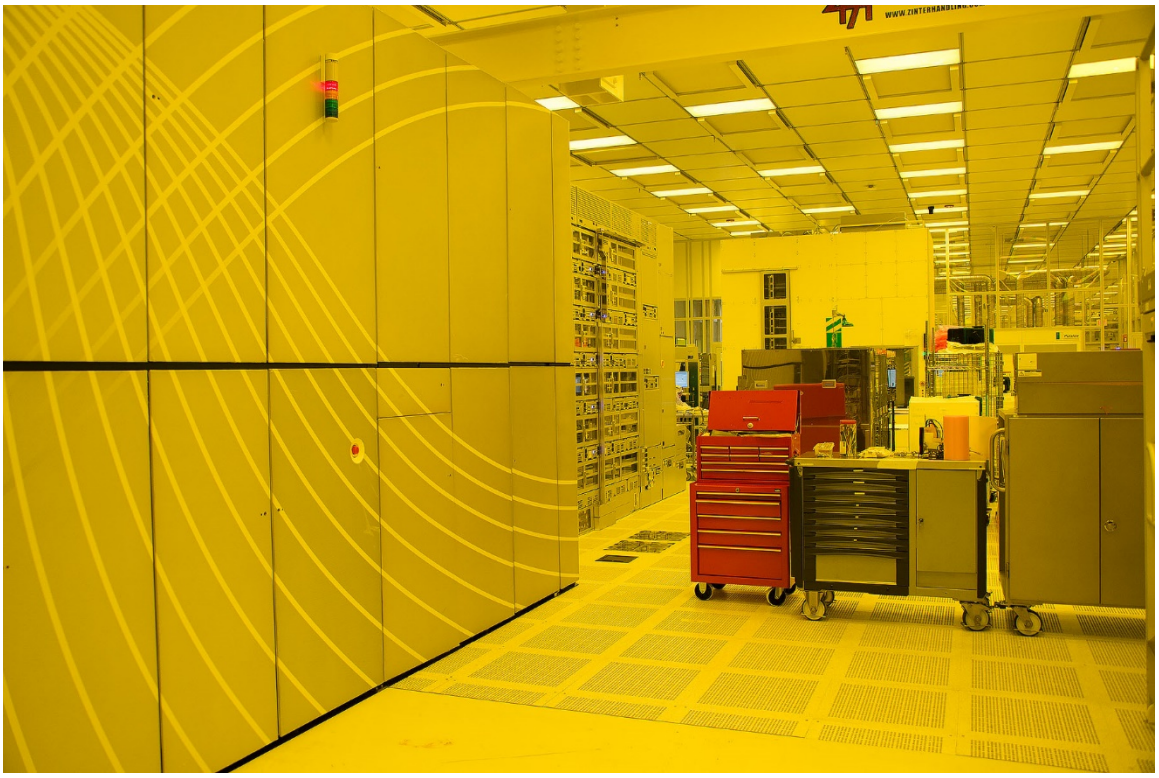

**Why IBM Semiconductor Technology Research represents a new R&D Model**

The differences between IBM Research's Semiconductor Group and the common fab-lite research model are pretty significant. One is that the group's focus is more on research, leaving development to their partners. Partners like GLOBALFOUNDRIES and Samsung are IBM's drivetrain with the primary advantage being that they have the scale needed to continue investing in building out foundry capability for new nodes. Also important is their foundry's need for IBM's research, which is seldom the case with the common fab-lite research model. This need is both economic and technical – saving hundreds-of-millions in duplicative research, while shortening time-to-market in countless ways.

Another difference is IBM's demonstrated ability to partner well. IDMs converting to fab-lite models come from a background of being extremely secretive, which carries cultural baggage that is counterproductive to partnering. Moreover, IBM's partners already incorporate IBM's research as a critical part of their pipeline. This has not been the case with the fab-lite model, because their foundries had their own separate pipelines. Because IBM's already a part of their foundries' pipeline, they can look farther ahead without duplicating efforts. That's been the whole point of IBM's partnering efforts over the last few decades.

| Key Success Factors for Stand-Alone Logic Device Research | IBM Model | Fab-lite R&D Model | Consortia Models |
|---|---|---|---|
| Use a Foundry for manufacturing | Yes | Yes | No |
| Stop investing in leading edge capacity | Yes | Yes | NA |
| Maintain leading-edge R&D facility | Mostly R | r&D | R&d |
| Invest in leading-edge tool sets | Yes | No | Limited |
| Use OPM$^2$ for access to tool sets | Yes | No | Yes |
| Established channel to transfer results | Yes | No | Limited |
| Ability to partner well | Yes | Maybe | Yes |
| Full flow process transfer to >1 Foundry | Yes | No | No |
| Have a system-level need for device knowledge | Yes | No | No |
| Ability & need to look 50 years out | Yes | No | No |

Are the research results transferable? Definitely: there is really little change here other than ownership. GLOBALFOUNDRIES' fab is not far away and GF has researchers and engineers at Albany, NY Research Fab as well. Moreover, IBM has developed partnering into an art form over the last three decades. Fab-lite's results are typically a duplicative waste; since the foundry will do its own research and create its own PDKs. Consortia have an ability to transfer. But they are more project oriented with the effort driven by their customers. Customers typically hand off these projects because either they need independent validation of what they already know or they want to outsource work that will not result in a differentiable advantage.

In contrast, IBM Research Semiconductor Group's contribution to a foundry's finished process will be significant, critical, and is differentiating. For example, when their foundry partners needed a finFET process, IBM had been working on the technology since the beginning.

There's also a significant difference between it and the consortia models. This is an important distinction, because the value of consortia dropped significantly when consolidation created fewer chipmakers than consortia. One of the big value propositions of consortia was to do pre-competitive research for lower cost by eliminating the need for duplication. This value falls with the number of companies in the pool. Moreover, consortia are also engines without drivetrains, limiting their value. Plus, their ability to attract leading-edge tool sets varies over time, limiting their quality.

**So what's in it for IBM?**

So it's clear IBM Research's Semiconductor Group is essential to IBM's future and its partners. It's also clear that IBM can continue to do cutting-edge semiconductor research. That leaves the business case question to be answered. Many in the semiconductor industry would argue that there is no business case today. But the strategic moves of other companies suggest that IBM is ahead in the race.

Think about how several years ago the world's largest fabless producers, such as Nvidia and Qualcomm, began to spend significant resources to reach down into the supply chain to understand how new equipment, materials, and production techniques coming with future nodes would affect designs in development.

Then Apple bought a fabless processor company and soon after started to do the same. More recently Google and Facebook have followed in their steps, first building chip design capability and later hiring process engineers before they came out with their own chips. In short, they have been visibly reverse engineering IBM's vertical research model over the last ten years, which starts at the top of the supply chain and reaches down.

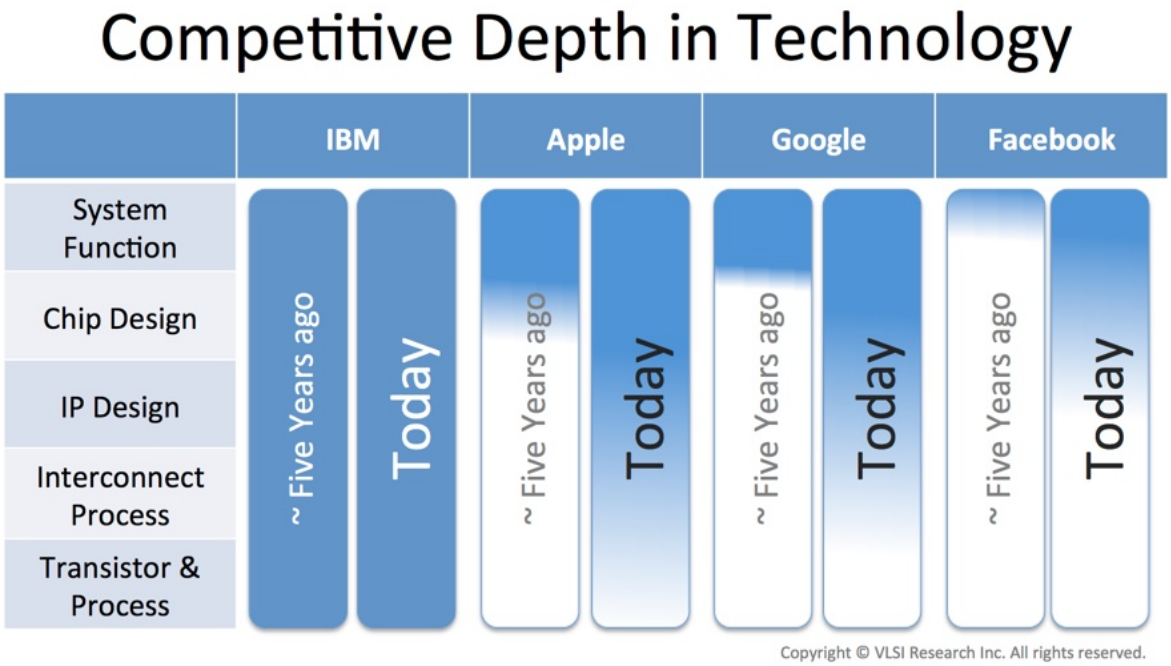


Given these strategic movements, there must be something to having an active semiconductor research capability. The question is why?

Why reverse engineer IBM's vertical research model? The reason is that system architecture advancements are constrained by the process technology advances. It is transistor and interconnect technology as well as variability in manufacturing that ultimately define the performance and functionality of the finished product. If you're fabless, like IBM now is, you also need to have the technical understanding to know which foundry is ahead, as this will determine the competitiveness of the finished product.

But it's more than just performance and functionality. It's also provides assurance that a company will successfully tapeout its first chips on a new node. Early designs on a new node are the most expensive and have the highest probability of a tapeout failure. For a company like Apple, they risk missing a product announcement cycle for which the revenue and profit loss would far outweigh the upfront cost of early assurance. Top-Tier-Tech (T-cubed$^3$) companies attempting to design chips at the leading edge quickly find this out once they start making their own chips.

Don't make the mistake of seeing this as a single design either, because it's never about designing one chip. Success is defined by a consistent ability to successfully design the next chip, the next chip after that, and so on. Otherwise, your products and your company will lose their relevancy. That means you have to look further ahead if you are a systems company. The risks for systems companies today are even greater as radically new computing architectures appear on the horizon.

IBM has over a hundred-year history of innovation. They've had to reinvent themselves multiple times. To do this successfully you have to know where the world is going long before it gets there. You must have vision with a horizon that's far out. For IBM it reaches out thirty to fifty years. On IBM's horizon are exotic technologies like neuromorphic computing, cognitive hardware, and quantum computing. Knowing those are the mountains they will climb allows everyone to align to making it possible. You can't do this from a typical foundry's vision that stretches out two to three nodes, or four to six years.

The practice of looking far ahead has been the case at IBM all along. At IBM research sites around the world, people are figuring out what future system requirements will be. This is tightly intertwined with what the technical requirements will be. They call it DCTO or Design-Technology-Co-Optimization, which is akin to DFM, but starts much earlier and covers far more ground. The task is to take these requirements and translate them into the technical elements they'll need and vice-versa. At each layer, people are looking to the future and then co-optimizing what they find up and down the chain from process to device and interconnect, to chip, to sub-system, to full system and back. Then the Semiconductor Group works on figuring out the value proposition for each of the various elements. Finally, they finish with proof of concepts for manufacturing.

The breakthrough of the 2020s will be getting past atomic dimension limits. There are plenty of possible ways to break through these limits: gate-all-around, the vertical transport FET, nanowires, nanosheets, photonics, and/or 3D multi-chip stacking, just to mention a handful. When it comes to scaling, only a few people commit to 2 to 3 nodes out and most are always in a next-node frame of mind. IBM's horizon on scaling's limit is 2032. So how do they do this?

IBM looks at the future from multiple vectors that end up as a system. They start with the application and then figure out what the system will look like – figuring out the system requirements of Watson for example. From there, they work their way back in time and down into the component details. The reason is that in order to tackle the next big application, you need more performance. The quest for performance starts in the arcane details at the atomic level of transistors and interconnect. Their quest is to take atoms and molecules to technology readiness.

For example, the next big app is big data analytics which, will be needed to process the volumes that will come from IoT edge devices. In order to have the big data analytics, you need a system. In order to have a system, you need an architecture, which for IBM today is POWER. In order to realize that architecture, you need a processor. You also need to understand memory and interface chip requirements which will glue it all together. In order to have a processor, you need a transistor and interconnect. In order to have a transistor and interconnect, you need to understand the materials, tools, and processes that will make them possible. These affect performance parameters like contact resistance. Moreover, it's not just a single transistor in development. Systems engineers want them to figure out how to deal with multiple threshold-voltage gate stacks – each with different work functions. Fail at any one of these and it all falls apart.

So the ROI for IBM is having the technology manufacturing ready for the systems, whatever form they may take, that will be introduced in the 2020s and 2030s. The ROI is ultimately about business continuity … not about the next quarter or the next year. It's closer to be able to double-down on a 100 years of innovation to make it 200 years.

[1] A 'real fab' is one that manufactures semiconductors in volume. In contrast, a 'real lab' is one that can deliver R&D that is either pilot or full-volume production worthy.

[2] OPM: Other Partners Money

[3] While an Apple can be easily labeled as an ODM, it's harder to put this label on a Google, Facebook or Amazon – which are really Internet companies. One thing is clear is that they are on the top tier of technology supply chain, hence my creation of the term T-cubed.

## WildPhotons: light life lessons

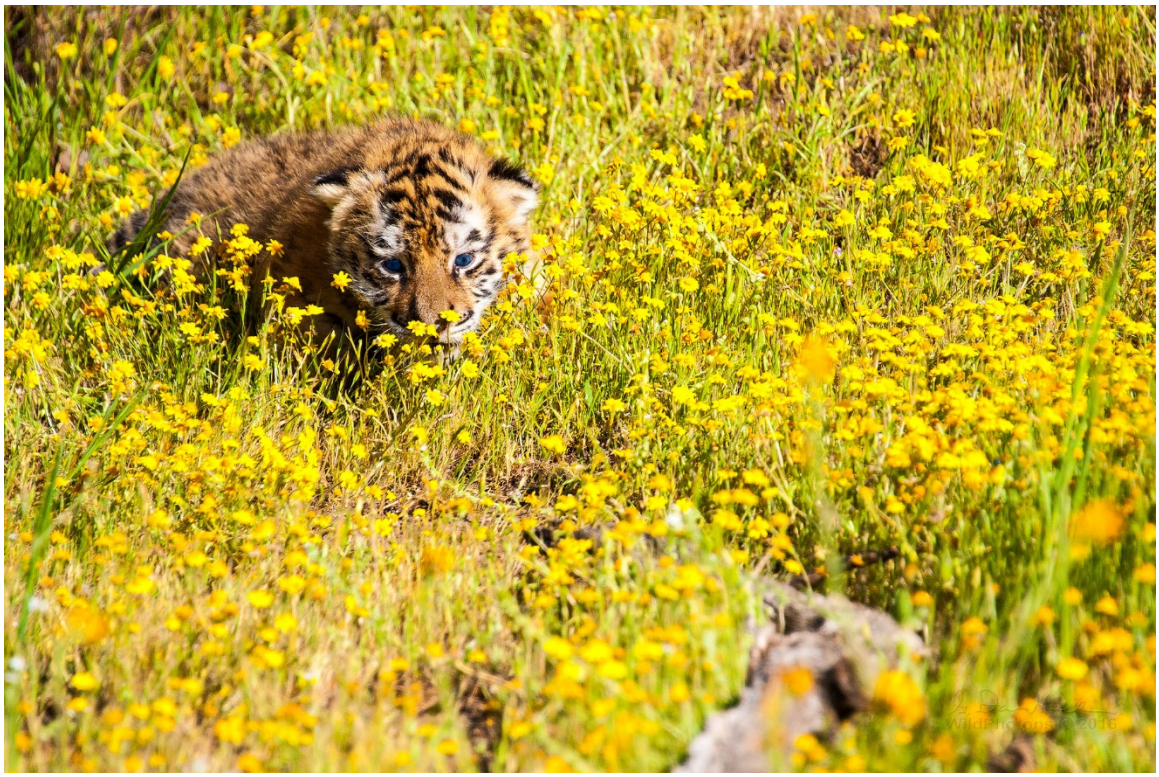

The present is the new future: That where you sit, you create everything that's going to come

Sarah Jones

A 4 week old tiger cub stops to smell the wildflowers, Panthera tigris, CA. Tigers are largest of the cat species. Males can reach a body length of 11 feet and weigh as much as 670 lb.
Reference number: TDtg_1304_394
follow Dan Hutcheson's photos on twitter @wildphotons
see them all at flickr.com/photos/wildphotons/

## ***August 4, 2017 – Strategy and Tactics:*** *Intel's Strategy: Just One Word … "Data"* ***WildPhotons:*** *Leadership …*

**Focal Points:**

*Intel's Strategy: Just One Word … "Data"*

- Brian Krzanich's strategy for Intel
- Extracting value from Data Rivers
    - Battle Map of IoT versus the Cloud
    - Four Value Platforms of tech emerging
    - Data Rivers and Streams
- Cognitive Clouds with IoT on the edge
- Everything runs on Silicon
    - Why Moore's Law is needed for tactical execution
    - Making network
    - The monetization horizon of silicon

*Chip History Legends*

- Robert N. Noyce, Intel's first CEO
- Thanks to **Applied Materials** for funding a renovation of The Chip History Center

*"History never repeats itself but it often rhymes."*
*—as Mark Twain is reputed to have said*

**Intel's Strategy:** If you ask Brian Krzanich what Intel's strategy is he'll answer with one word, "data." I've met a lot of CEOs who can distill their strategy down to a single sentence. I've never met one who could do it with a single word. On the surface it may sound simplistic … like the advice given to an Alfa Romeo driving college boy in a movie. But it is actually very deep … and yet distilled like a single malt. Here's why:

A wise man once told me that to get rich, all you have to do is get next to a river of money and skim off a little as it passes by. That's what credit card companies do, what investment banks do, and virtually all financial institutions do. Now keep that image in your mind and imagine this river of money morphing into a river of data — with those rivers converting to money every day.

Data is the currency of the digital age. Rivers of data crisscross the wired and wireless virtual world to pop into the real world with real value virtually everywhere. Tapping into these rivers is the business model that's driving technology today. In place of financial institutions, it is Amazon, Google, and Facebook that dominate the space today. So where is Intel in this space?

It is data enablement that glues together Intel's strategy going forward. Many see Intel's strategy beyond PCs and datacenter as haphazard. Drones, robots, virtual reality, smart city, supply chain logistics, and self-driving cars look to be all over the map from a traditional electronics market perspective. Acquisitions of companies in AI, VR, vision, 3D reconstruction, semiconductor, and more may seem to be an out-of-focus spectrum of technologies. But it all comes into focus when you look through a lens focal point that is data.

Since Brian Krzanich became CEO, Intel's major moves have been towards having the capability to enable its customers to extract value from Data Rivers. There is a flood from these growing between the clouds and IoT sensor networks emerging in verticals on the edge. Some are huge, amazon-like rivers, with many tributaries. Others are fresh little spring creeks. What's common to all is deriving value from data flow.

This is a much more complex strategic map of the future than what I drew five years ago, where there was a cloud dominated by Intel pitted against an army of things on the edge, dominated by ARM.

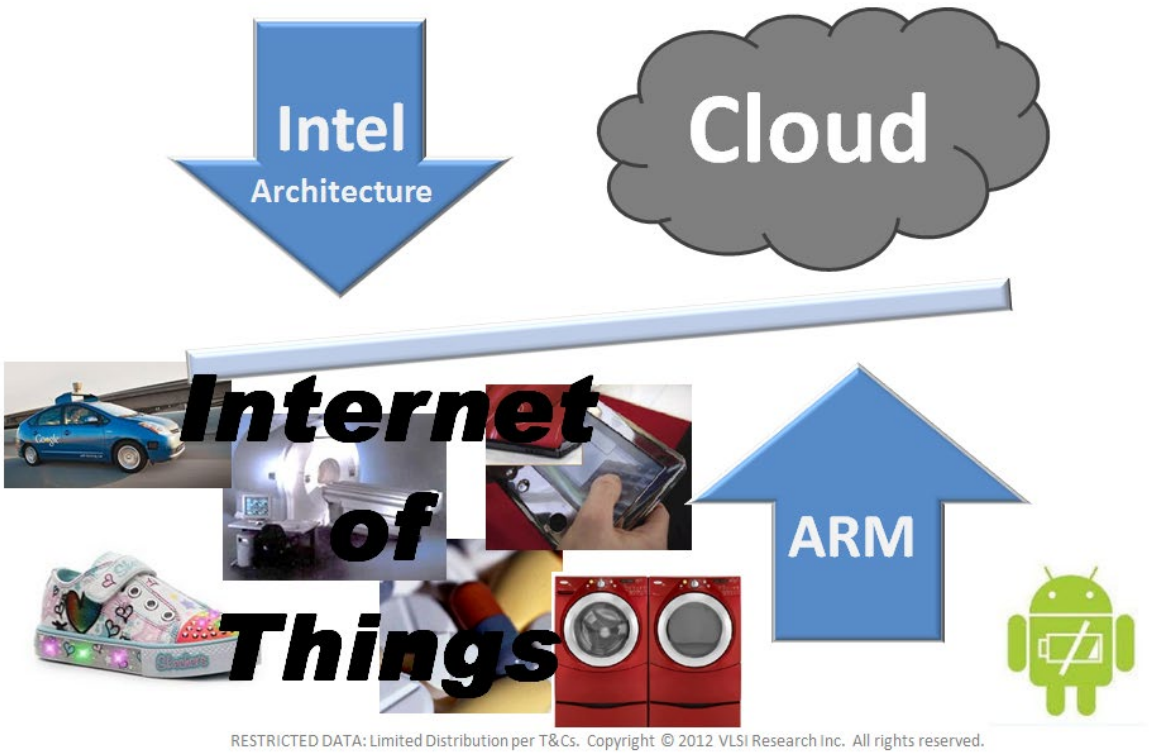


Back then, there was little visibility between the cloud and IoT different than what was well established at the time. Two years ago I described how the technology value platforms were emerging. While it was predictive of how things would evolve, the structure was still not strong enough:

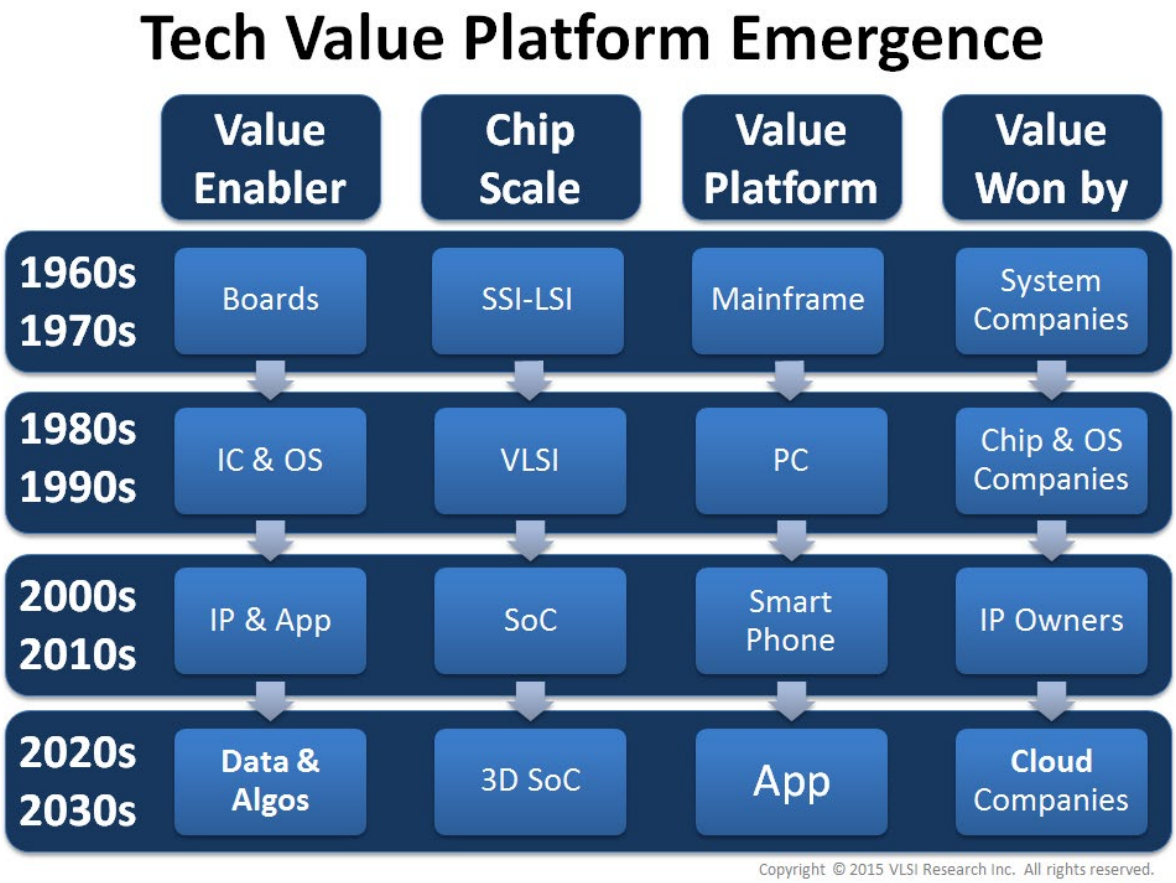


Today, a structure is evolving around the four value platforms of the next twenty years. This structure binds everything together with rivers of data. These Data Rivers and their tributaries flow to connect skies full of cognitive clouds with sensored-and-connected IoT objects around the edge in market verticals. There are also streams of data within and between the market verticals.

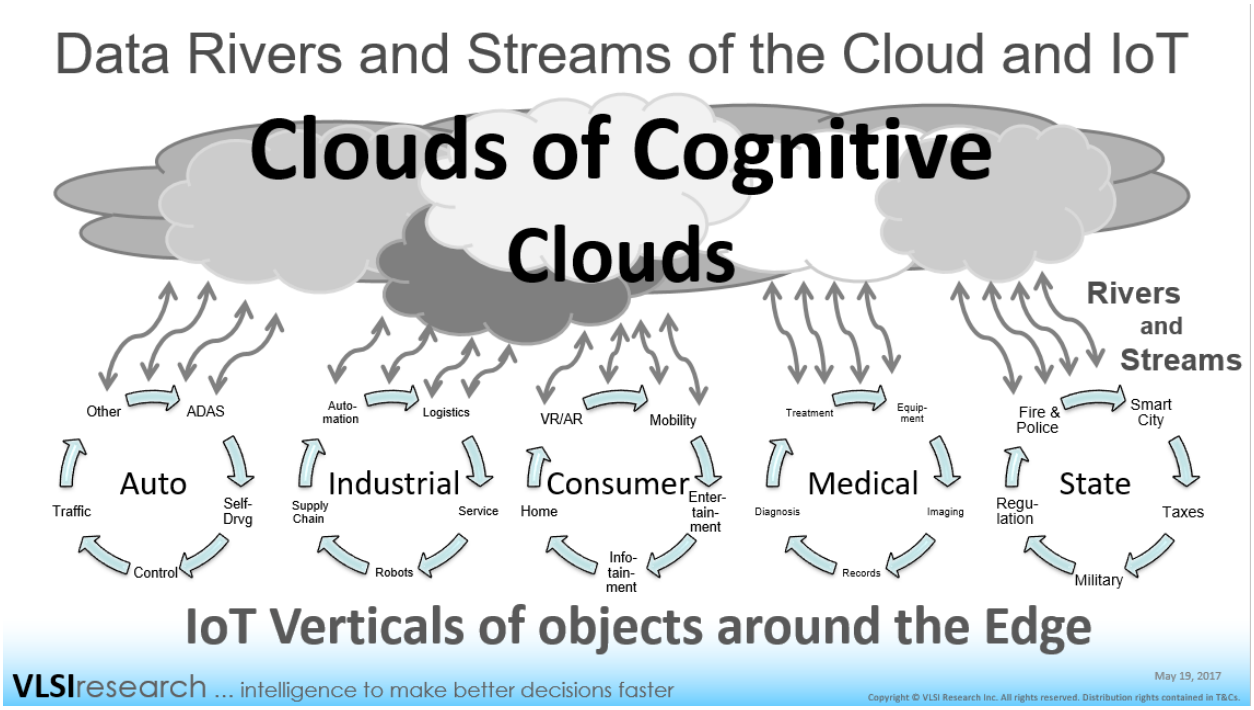


The immense complexity of this structure is what makes seeing it so difficult to imagine, because most only see clouds or edge devices when attempting to conceptualize it. We have a bias built by a history of thinking that the world is centered around single products, like PCs and Smartphones. So we tend to see the emerging products as just a new thing in a field of many different things, while waiting for one to rise above the rest.

But all of these things are like the castles one sees along the strategic rivers of the old country. Back then, rivers were where the strategic value lay. Kings used the castles to extract the value. Today it's the data that's important. Like water is the source of life… In this new world: Data is the source of value.

So how can Intel capture the value of this data? It's pretty simple. The building blocks of all the new castles and forts are semiconductors: Everything runs on silicon. But the silicon cannot stand alone as individual chips anymore — just like the castles of old could not stand alone and needed a kingdom to connect them for strength. IoT objects on the edge form the castles on the tributaries that flow data to rivers, which need to be connected to the cloud and cognitive analytics to form the strength of new data kingdoms.

Some are the cloud empires that already exist, like Amazon, Google, or Facebook — all names that come readily to mind. Then there's Alibaba, IBM, Huawei, Microsoft, Tencent, to name some more.  What's common to all is that their business models run on silicon. On top of this, a chip company must bring standards to life that fertilize new markets, while having the scale needed to implement a wide variety of applications and workloads into silicon that work seamlessly together. These are core strengths of Intel that can be brought to bear on executing a data centric strategy.

Tactical execution of this strategy requires having the breadth to span what is one of the broadest range of applications anyone could imagine. It also requires bringing a useable framework to the many companies who have plans for an IoT strategy but can't execute because they don't know how to put it into a tactical framework that is executable by their organization.


So why does Intel need Moore's Law? The monetization horizon of all semiconductor silicon is hardware. Even when the value is in software, it must be derived in hardware. Hardware is the end point of value creation with silicon because that is what people pay for. In general, people have historically done something with the hardware they bought. In the PC era, that was typically some computational task based on data that was really hard to come by and was manually entered. Cell phones were first used to communicate, but as they became smart, they began to take on compute workloads as well, because that's what smart things do.

Hardware is still important. What's changed is that instead of people doing something with the hardware, the hardware is doing something for them. That's the difference between smart and cognitive. For example, asking Alexi or Siri to play a song is a smart task. Cognitive would be "tell me if Julie comes back after curfew."[1] That change is a revolution in what needs to be done with silicon. Moreover Intel's ability to bring together many disparate point solutions in a manner similar what was done with the PC is a key competitive advantage.

But unlike the PC, it will be across multiple verticals that are often served by large corporations who want a total solution. The giant cloud companies of today are just the start of what is a

revolution in how companies serve their customers. They have the data. So now it's a matter of what to do with it. Brian intends to provide the silicon to enable their success.

1) In the first instance, all that has to done is identify the title as a song you want to play, search a data base and then stream the data to a speaker. In the second, it has to understand that Julie is your teenage daughter, have learned what time curfew is for her from previous discussions, and then know the best way to alert you if Julie comes home late. For example, you could be out or you could be home asleep and want different alerts for each case.

**Chip History Legends:** Robert N. Noyce, **Intel's** first CEO. When considering Intel today, I always find it useful to review the history of the company. A great place to start is with Bob Noyce, who really saw the organizational value of getting away from an Eastern hierarchical structure to establish a meritocracy … something he was most proud of. Here are some links to learn more at The Chip History Center:

Legends, Circa 1990 - Remembering Robert Noyce
Interview with Leslie Berlin on her Biography of Bob Noyce - Segment 1 Segment 2
Robert N. Noyce: Early Career
Robert N. Noyce: The Intel Years

Many thanks to **Applied Materials**, who recently funded a renovation of The Chip History Center website, including transferring all files and code from old format standards to current ones. When we started this project a decade ago, everyone thought once it's in digital, it would always be in digital. Sadly, that's not true as file format standards and what browsers will display constantly change due to security issues, which arise as old open standards get hacked. So without Applied Materials generosity and the generosity of many others, we would not be able to continue this effort to preserve our industry's history.

## WildPhotons: light life lessons


follow Dan Hutcheson's photos at [flickr.com/photos/wildphotons/](flickr.com/photos/wildphotons/)

***May 26, 2017 – Strategy and Tactics:*** *Samsung's Foundry Forum* ***WildPhotons:*** *Grabel's Law …*

**Focal Points:**

*Samsung's Foundry Forum:*

- Samsung Semiconductor's new business structure
    - What it means from perspectives of:
        - Competitive counter-marketing
        - The role Samsung expects it to play
        - Shareholder value
- Structure
- Production, Packaging, and Test
- Samsung Foundry's differentiation strategy
    - Addressing Fabless 2.0
    - What it means for process technology
- Technology Roadmap
    - EUV & EDA

*Maxims: Brand is more than a promise*

- Brand is a reputation for consistently delivering…
- The great brands

- The boneyard of failed brands

*This Week, 10 Years Ago: Industry Legends*

- MAXIMS applied:
    - TSMC and its proven path to success
    - Samsung's redefinition of the vertical integration model

*"History never repeats itself but it often rhymes."*
*—as Mark Twain is reputed to have said*

**Samsung's Foundry Forum:** featured an announcement laying out Samsung Semiconductor's new business structure. The company is carving its foundry group out of the System LSI business. They named **ES (Eun-Seong) Jung** to be GM. He jokingly says the ES initials stand for Engineering Sample, though not joking that his goal is to get Engineering Samples to customers fast and first-time-right. Backing him up are **Jong Shik Yoon** — heading up foundry technology, **Charlie Bae** — running strategic marketing, and **Siyoung Choi** — head of manufacturing.

This directly addresses competitors positioning it in the corner of not being a 'pure-play foundry.' Now they are 'pure-play foundry' in name at the least, which won't stop competitors counter-marketing efforts. However, I have found that the concept of pure-play foundries having customer value has lost the punch it had in the nineties. These days, when it comes to customer impact, 'pure-play foundry' garners the attention of a small web page banner ad. Customers I talk to are pretty satisfied with Samsung's IP control. Especially given the design complexity of today's chips.

More importantly, spinning its foundry out means Samsung sees the business as large enough and sustainable enough to stand out on its own. It also gives visibility to a group with the potential to grow faster than the company as a whole, which is of great shareholder value.

One interesting aspect is that they intentionally sandwiched in-between the Memory and S. LSI group. This demonstrates the growing product ties between memory and logic. Another important recognition of these ties are that the Memory and Foundry groups will share Samsung's R&D and Test & Packaging Centers. Advanced packaging has become highly visible customer value center, which gets most of the marketing glory in the back-end these days. But pay attention to test.

**Test**, while far less visible marketing-wise, is no less critical. As the number of fabless companies has fallen from the many to the few, brand value has risen substantially. Brand used to not matter in the days where crappy fabless products could hide in the trash bins behind the big-box stores return's desk. The remaining giants are too large to hide. Plus, the sheer size of production from them makes the cost of field failure painfully high (Just ask Samsung about the Galaxy Note 7 … ah … Forget that … Reminding them might bring a MA rating for impudence,

language, and possibly violence). Point is that trends in SoC and Advanced Packaging are driving test off the chip and back into the tester, which is an opportunity for test companies.

**Samsung Foundry's differentiation strategy** is to bundle the company's strengths in technology for logic and memory, test and packaging into a total turnkey solution. The reason why this is important is what **Qualcomm's Roawen Chen** describes as Fabless 2.0. It's a mirror to Ajit Manocha's Foundry 2.0, while almost rhyming poetically with his point is that the 2 year clock cycle of Moore's Law is too slow: It has to be done in a year.

So 1.0 goes to 2.0, while 2 goes to 1. This means you have to go from tape-out to mature yields in less than a year, which is a task that can only be done with the massive scale that a Samsung can bring. The key foundry equation for Roawen is that 1+1>>2. Of course this violates Grabel's Law. But even equaling 2 is better than the more common merger formula of 1/1, which can be doubled down on with $1^2$, and some have even worked at 1-to-the-$n^{th}$, all of which just results in just 1. Math humor aside, he does have a key point in that the serial foundry-fabless 1.0 model doesn't work in a world where you have to hit 60-80% die yields in a quarter or 2 (where 1 is optimum).

Fabless 2.0 has to be like high-bandwidth-memory with each group's timeline stacked on top of each other and plenty of TSV's[1] connecting them all together. That has to happen for everything to come together by Christmas. It's turbo-charged partnering: more performance with less power consumption.

**What this means for process technology**, I have described before, which is node-splitting. Something like LELE, except its DRDR (Develop-Ramp-Develop-Ramp, which can be pronounced as Doctor-Doctor). This is actually the same strategy Micron used in its early days to beat the Japanese.

In Samsung's case it means that 8LP will be squeezed between 10LP and 7LP. Then 6LP between 7LP and 5LP. Before that will come the more normal serial process improvements. So 14LP will have four 3$^{rd}$-letter versions: E, P, C, and U — while 10LP breaks into three: E, P, and U — these will all be finFET processes.

Samsung was clearly proud that they smoked all the foundries to 10nm, with the chips having shipped all year to appear in the wild inside the Galaxy 8. More interesting is that they boldly issued challenges: intending to be first with EUV at 7LP in 2018 and first with GAA[2] at 4LP in 2020. Also announced their next generation FD-SOI technology: 18FDS, which will come in 2019. In the meantime, they will add RF and eMRAM to it. Now I know this number soup is pretty confusing, so here's their roadmap.

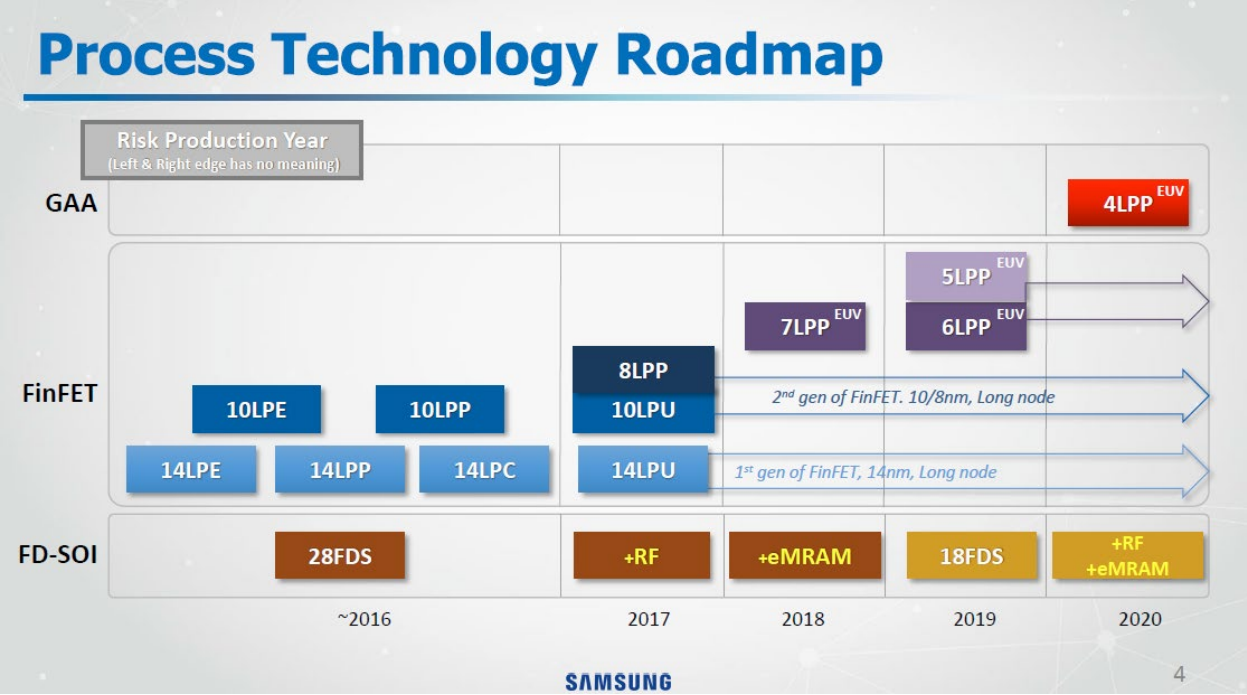


Wait-a-minute … did you just read that Samsung has a roadmap??? Shock of shocks, they did release a roadmap. Normally, Samsung motto has appeared to be, 'don't say it until you can do it,' because failure is not an option for them. I don't think they have adopted carefree marketing motto of 'say it, try to do it … then don't be a fool about it: update the damn roadmap.' I just doubt that their culture could change that fast even if they wanted to. Instead, I take this to mean that they've made significant progress with EUV and GAA. Progress enough to confidently say they see a path to 1.5nm! So my prediction is that we'll see GAA 4LP go to 3, 2, and 1.5.

**EUV** is another space when Samsung has taken clear public lead since last year. The big news this day was that Samsung had summited the magical 250W laser-pulsed Sn plasma source mountain in the prior week, achieved 70% availability, and see 140 WPH with 90% availability in the near future. I should add that there was plenty of supplemental oxygen from **TEL** and **ASML**. While the others can buy the same oxygen, it still took a lot of dogged determination to pull it off. Nevertheless, exposure is just one thing. They still need masks, which they are confident the infrastructure is in place. Confident because what's not available in the commercial world they have work-arounds and their own tools, including their actinic defect review tool that was presented earlier this year at SPIE. I would not be surprised if they don't have their own actinic inspection tool, as well, but that has not been announced.

**EUV EDA Tools**: **Synopsys**'s **John Koeter** gave a sharp presentation on how they've already tackled the design tool challenges for single-patterning EUV insertion at 7nm for every step of the flow. They've also dealt with EDA needs for coverage space extraction, via-stapling, contact centering, dealing with non-Gaussian parametric on-chip variation, and advanced wave-form

propagation, to name a few. ARM, Cadence, and Mentor also presented their capabilities as well.

**IoT Security**: Turning to the more elegant side of innovation, one of the real diamonds was SAMPUF. It stands for SAMsung Physically Unclonable Function. What's so elegant is that it works by using the parametric variation signature of a string of gates. The area penalty is infinitesimally small (250 square microns). Plus they can be put anywhere, such as security gates for each core or memory block. I was impressed, seeing it at an elegant level equal to 'Breakfast at Tiffany's.

P.S. Please let me know if you enjoyed reading this, as I'm experimenting with a new writing style. The idea is to use more metaphors and similes to communicate the more qualitative views I came away with, while having fun writing it. Also, don't think I am favoring Samsung as it is an experiment. If you like it, I will write this way about you too.

1 TSV: Through Structure Via
2 GAA: Gate All Around

## Maxims: **Brand is more than a promise**

Brand is more than a promise. Brand is a reputation for consistently delivering a result, not just a product or a solution. Cesar Millan, the celebrated dog trainer, started off charging $10 per session. As he became well-known, his fees soared to $10,000 and even $100,000 for a 'private consultation.' So what changed? His English did get smoother and he did become better at marketing basics. But none of this explains the thousand to ten-thousand increase in the price he could charge. The bigger difference came from increased brand value. This increase came from delivering systematically good results that were initially spread via word of mouth.

Just take a look at a few of the great brands in our industry: Samsung, Intel, or TSMC on the semiconductor side. Applied Materials, ASML, KLA-Tencor, TEL, Teradyne, Advantest, Besi, K&S, or VAT to name a bunch on the supplier side. They, and all the names I couldn't name for the sake of brevity, all became pack leaders because due to one thing in common: they deliver great results. This means delivering consistent value for what the customer's pay, a track record for an ability to develop value, and a believable roadmap to the future.

If you dig around the boneyard of failed brands in our industry, you'll find they failed on one or more these points. That's not to say you can't fail-to-deliver once, because the brands above have produced some unruly dogs at times. But when they did, the difference between greatness and failure is how they dealt with their failures. They fixed the problems and made the customer whole. It may have been a rough road to the end, but they did deliver the results.

***Maxim: "an expression of a general principle."*** *– Webster's Dictionary*

Go to Amazon to get a copy of the latest: **Maxims of Tech** –
Rules of Engagement for a Fast Changing Environment
or how to thrive in what is the extreme sport of business

## This Time, 10 Years Ago:

**MAXIMS applied: TSMC and its proven path to success —** What made TSMC different from all foundries before it was that **Morris Chang** astutely believed that the foundry business was a service business, not a chip making business …

**Samsung's redefinition of the vertical integration model —** Attempts at vertical integration in the semiconductor business have typically failed. Customers don't like their suppliers competing with them. So no one attempted them since the Silicon Valley four failed at it in the late seventies. Yet … Samsung has figured out sales tactics to tier competitive instincts horizontally so they can compete on one level and partner on another… Samsung is levering off the fact that systems companies that use a distribution model should not see their chips as competitive… Samsung also strategically levered off structural changes that occurred in the electronics market… As a result, the electronics market developed walls that isolated areas like computing, cellular, and consumer… Samsung astutely figured they could compete in one area without offending customers for their chips in other areas. So when Nokia abandoned them as a strategic partner for building their cell phones to PICOSing Samsung, they walked away and went into the cell phone business. Why would a Dell care, or anyone else buying their memories? It was a risk on Samsung's part, but it was a well-considered risk where the probability of failure was low.

2 is not equal to 3, not even for large values of 2

Grabel's Law

Out of place Dromedary Camel, Camelus dromedarius, with handler in White Sands National Monument, NM. Why it Works: Rarity of finding a camel in New Mexico's White Sands desert. Reference number: NM_1109_333



WildPhotons: light life lessons

follow Dan Hutcheson's photos at [flickr.com/photos/wildphotons/](flickr.com/photos/wildphotons/)

**VLSI**research

**Ph: 408.453.8844 | 2290 North First Street Suite 202 San Jose, CA 95131**

**… intelligence to make better decisions faster**

**VLSIresearch.com** **ChipHistory.org** **weSRCH.com**