PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Moore's law, lithography, and how optics drive the semiconductor industry

G. Dan Hutcheson

G. Dan Hutcheson, "Moore's law, lithography, and how optics drive the semiconductor industry," Proc. SPIE 10583, Extreme Ultraviolet (EUV) Lithography IX, 1058303 (19 March 2018); doi: 10.1117/12.2308299

SPIE.

Event: SPIE Advanced Lithography, 2018, San Jose, California, United States

Moore's Law, Lithography, and how Optics Drive the Semiconductor Industry

G. Dan Hutcheson

VLSIresearch inc, 2290 North First Street, Suite 202, San Jose, CA 95131

ABSTRACT

When the subject of Moore's Law arises, the important role that lithography plays and how advances in optics have made it all possible is seldom brought up in the world outside of lithography itself. When lithography is mentioned up in the value chain, it's often a critique of how advances are coming too slow and getting far too expensive. Yet advances in lithography are at the core of how Moore's Law is viable. This presentation lays out how technology and the economics of optics in manufacturing interleave to drive the immense value that semiconductors have brought to the world by making it smarter. Continuing these advances will be critical as electronics make the move from smart to cognitive.

Keywords: Moore's Law, lithography economics, scaling, semiconductor wafer fab economics, semiconductor business models, Koomey's Law, Dennard scaling, Metcalfe's Law, Emergent Behavior

1. INTRODUCTION

Ever wonder why the many apocalyptic predictions of Moore's Law faltering due to soaring costs have never come true? Predictions of Moore's Wall have come and gone over many decades. In the era of semiconductors, Moore's Wall predictions have appeared more often than ones predicting the end of the world. The reason why Moore's Law is so resilient lies in how advances in lithography processes interact with Moore's Law.

That's putting it mildly, because lithography equipment advances the most critical driver of Moore's Law. Or as Gordon Moore once told the author when asked, 'Isn't advances in equipment the most important driver of Moore's Law?' ... "Equipment advances have been EVERYTHING when it comes to Moore's Law."¹

Now it's clear to anyone in semiconductor lithography that improvements in resolution and overlay are the mechanical reasons why transistors get smaller and density doubles per node — especially when optical lithography is viewed as a system of tools that encompass tools for exposure, resist, deposition, etch, and more. Demonstrating it can be easily done with grade-school math.

Yet, how is it possible that the business still thrives given the history of lithography, where the price of a single tool has soared from 10 cents to today's \$100M+ EUV behemoths?

Year	Node (nm)	Lithography Technology
1957	254000	Camel's Hair Brush
1958	127000	Silk Screen Printer
1959	76200	Contact Printer – Emulsion plates
1964	16000	Contact Printer – Chrome plates
1972	8000	Proximity Aligner
1974	5000	Projection Aligner
1982	2000	g-line Stepper
1990	800	i-line Stepper
1997	250	248nm Scanner
2003	90	193nm Scanner
2009	32	193nm Immersion Scanner
2019	7	EUV Scanner

Copyright © VLSI Research Inc. All rights reserved.

Figure 1. History of semiconductor nodes and the lithography tools used in production. In 1957, Camel's Hair Brushes, costing 10 cents were used by operators to paint wax masks on the surface of the resist. Today's EUV tools cost over \$100M dollars — an increase of nine orders of magnitude, which is a compound biannual growth rate of 100%.

And it's not just the equipment that's become more expensive: The cost of the entire fab has risen with it. How is this possible?

The answer is that there are fundamental economic engines underlying this paradox that make the rising cost of lithography affordable given that there are technical advances in both optics and the processes that allow the imaged patterns on a wafer's surface to be realized in the films below.

2. MOORE'S LAW AND ITS ALLIES

A fact that should never be overlooked is that there was much more to value creation in our industry than just "Moore's Law." Nevertheless, there is a tendency for people to bundle it or unbundle it depending on their purposes at hand. It is also denigrated in a similar fashion to promote new products and thoughts.

Problem is that one believes Moore's Law is dead, then ... they must also believe AI and IoT are dead: And that there will be few gains to be made in tech, because without Moore's Law, there is no foundation for advancement and growth. It's common to see experts in their field present Post-Moore's Law (PML) visions in which their favorite technology is the natural successor. While saying Moore's Law is over is certainly great click-bait, the logic in making the connection seldom holds. For example, arguing that Moore's Law is dead because MPU density-per-chip has slowed, when GPU densities have not. Then apply this to Artificial Intelligence, for instance: The first thing experts will say when you ask the why now question about AI, since it failed multiple times in the past five decades is, "The difference is that now we finally have computers powerful enough and cheap enough to make AI practical." The compute power for this has come from Moore's Law bringing more transistor density to GPUs. The gains in performance would have never been as much as they would have been limited to architectural gains, such as the shift from MPUs to GPUs. To pull this off, GPUs have needed

greater transistor densities and lower power densities (which is not Moore's Law, but it does come hand-in-hand with it) to do the massively parallel compute operations needed. And these chips certainly wouldn't be any cheaper. Therefore, there are few gains to be made in the PML era. This means all the business opportunities fade away with it. What about IoT? A lot of the future value in IoT comes from the ability to gather and pipe massive data streams to clouds for AI systems to distill the data into intelligible and actionable information. Just look at last year's Next Big Thing: smart speakers. Products like Amazon's Echo upended the world of music again as an IoT product on the edge that's linked to AI in the cloud. In fact, significant advances in almost all of tech, including green-tech and med-tech, will slow with Moore's Law.

So, what are 'Moore's Law and its Allies' and how do they create value in the semiconductor industry? There are five first-order mechanisms that create value in our industry. They are:

- Moore's Law: it drives Economic Performance
- Koomey's Law: it drives System Performance
- Dennard's Law: it drives Device Performance
- Metcalfe's Law: it brings more users, to drive demand
- Emergent Behavior: It creates the order needed for innovation to proceed across broad fronts.

Most people, when talking or writing about Moore's Law, are not purists about its specifics. Instead they are bundling as many as five of these separate mechanisms together under the name of Moore's Law. For the purist, this law can only be the areal doubling of transistors every two years, with no more than a 30% increase in the cost of the area,² so cost-per-transistor continues to decline. Even though the author is a purist, he must admit that he too often falls into this trap for good reason. The good reason is that they are all governed by what he calls Moore's Clock. While it is reasonable to believe that while the purist's version of Moore's Law may slow, it is more realistic to believe that the continuation of similar value benefits via the bundled version is more realistic. Thus, the realist's version is bounded by Koomey's Law, Moore's Law classic, Dennard's Law, Emergent Behavior, and Metcalfe's Law (it could also be said that bundling all 5 into Moore's Law is likely done because it's easier to verbalize the five as simply "Moore's Law").

The classic interpretation of Moore's Law is that it is defined by areal density growth over time without a significant rise in the cost of the area, which is derived directly from his 1965 article in Electronics Magazine.³

This mechanism makes price-per-transistor go down, which has made electronics more affordable to the masses over time, as it broke out of military applications, to industry, and ultimately consumers. More simply, Moore's Law is about more for less: more value for less cost. It drives semiconductor growth because it works to the law of supply and demand.

Figure 2. While the exact cadence has not been strictly held over the 50 plus-years of its existence, varying between 1 and 3 years, Moore's Law had driven growth by providing more for less: more value for less cost.

The actual cadence for a halving has moved all over the map, with extended periods where it averaged 1, 3, 1.5, and 2 years. The early period of a 1-year cadence came in the 60s, which is what Moore had written in his 1965 Electronics article. In the 70s and 80s, it actually slipped to 3 years as the technology became more difficult. During this period, the industry would make three major lithography transitions, staying in optical — in spite of Moore's promise that x-ray and e-beam would be waiting in the wings³ (they're still waiting more than 50 years later). Things sped up in the 90s and 2000s, due to evolutionary changes in lithography, etch, and deposition tools, combined with industry roadmapping that better aligned R&D resources, thus reducing waste. The important lesson in causality here is that Moore's Law is the result of innovation, not the cause of innovation — as it's often sold (more on this later).

Koomey's Law is more recent and took the classic interpretation of Moore's Law one step forward to shift the value equation from cost to computational power-per-Watt. It shows a fairly linear transition from tubes in the 1940s to extreme levels of IC scaling today. It's never gained much traction, as it could be critiqued for cherry-picking data to achieve the linearity. Yet, everyone intuitively knows it's true from both engineering and market perspectives, while outliers should be seen as the result of either poor designs or different design intents.

Figure 3. Kooomey's Law states that computer power efficiency grows at a 54% CAGR. Thus power-per-computation is declining at a 35% annual rate. This is due to gains in: System Architecture, Moore's + Dennard's Laws, Lithography Scaling, and Materials Scaling.

Kooomey's Law states that computer power efficiency in terms of computations-per-kWh grows at just over 50% each year.⁴ It's important because its value mechanism adds to the feature benefits needed to get people to continue to buy new computational capability. While Kooomey says nothing about cost, which is a supply-side function, this 'law' says a lot about creating demand-value (i.e. how much people are willing to pay).

In contrast, Moore's Law covers the supply part of the value equation, stating in 1975 that component density doubles every two years for roughly the same areal cost.³ Later, he amended this to a cost increase that was no more than 30%, which he observed had been seen consistently through the 90s, which he stated in his luncheon keynote address to the NTRS working groups in 1992 in Dallas, Texas.⁵

What's so important about Moore's Law is that 1) cramming more components onto a single piece of silicon creates more opportunities for better electronics — as noted in his 1965 Electronics article — and 2) it covers the most important side of the supply-side function: Cost. Basic economics tells you that lowering cost lets a supplier ride down the demand curve, which increases sales volume. Hence, Kooomey's Law wilts without Moore's Law, which is probably why the first is rarely cited.

Dennard is generally credited with being the first to codify the power and performance gains from moving transistor elements closer together via a set of electrical engineering rules set forth at the 1974 ISSC conference.⁶ It covers a second part to the supply part of the value equation, stating that cramming those transistors closer together results in proportional power and/or performance gains. It's more formally known as Dennard scaling, which technically ended when the industry took dimensions below 100nm, as IBM's Bernie Meyerson was the first to note early in 2003. However, it continues in spirit due to the semiconductor industry's unique innovation model that is based on emergent behavior. The 'in spirit' comes from the fact that when the 100nm barrier to Dennard scaling was reached, the industry simply switched from lithographic scaling to materials-enabled scaling to achieve the same feature benefits for end-users. What's so important about Dennard's Law is that it brings the specific feature benefits, absent from Moore's Law, of higher performance and

lower power-consumption so important to all smart-era of electronics. Also, Koomey's Law wilts even further without Dennard's Law, because aside from system architecture gains, there would be no power-performance gains.

Metcalf's Law⁷ is on the demand side of the value equation drawing on the network effect of economics for the creation of demand via emergent behavior. Metcalfe's Law states that the value of a network is proportional to the square of the number of connected users of the system. As more users join, a network becomes more attractive for other users to join. Facebook is the extreme example of this effect. Others are smartphones, PCs, telephones, and even cities. Metcalfe's Law brings an emergent behavior mechanism beyond the essential need for computing: Here demand growth occurs as a communications network grows. What's so important about this is it means demand growth is only limited by the population. A population of people first (IoP) and an unlimited population of things (IoT).

Figure 4. Metcalfe's Law states that the value of a network is proportional to the square of the number of connected users of the system.

Emergent behavior, based on complexity theory, is when order emerges from chaos, due to a few simple rules and the introduction of a chaotic attractor that provides the seed. For example, birds flocking or fish schooling with the simple rule of closely following the animal nearest you. Emergent order is the engine behind all value creation through innovation by the semiconductor industry. Calling Moore's Law a self-fulfilling prophecy does nothing to explain how it works, while dismissing something that is magical to behold.

Semiconductor innovation is far more about emergent behavior than it is about technology and collaboration. The latter are just the fruits of what could be called 'emergent innovation.' Our industry's innovation is not classical: the legendary kind where a single individual like a Thomas Edison or a Steve Jobs develops a product from an idea. Semiconductor innovation only occurs when thousands coordinate their efforts to one simple rule: follow the technology nearest you to be ready by the next node or be irrelevant by the next node. This rule is more formally known as Moore's Law, as I will describe.

While purists like myself may see Moore's Law as narrowly defined around density, the world tends to have a broader view. It's less precise, but far more impactful. This broader view tends to bundle all the formal laws into one and call it all Moore's Law and make it synonymous with the concept of steady progress through innovation. The most logical reason for this bundling is that Moore was the first to provide a clock cycle for the cadence. A clock cycle is critically important because emergent behavior is a dynamic process.

Why it all gets bundled into “Moore’s Law”

All the other technology laws are not possible without
Moore’s Law

- **Koomey’s Law** needs more transistors to improve architectures
- **Dennard’s Law** needs to get devices closer together
- **Metcalf’s Law** needs lower prices to attract more users
- **Emergent Behavior** needs a simple rule to create order from chaos

All Need a Clock

Gordon Moore in 1975 Source: Intel

VLSIresearch ... intelligence to make better decisions faster

16 Aug 2017 11

Copyright © 2017 VLSI Research Inc. All rights reserved. Distribution rights contained in T&Cs...

Figure 5. The other technology laws are often bundled into “Moore’s Law” as a broader concept because the others are not possible without Moore’s Law: Koomey’s Law needs more transistors to improve architectures. Dennard’s Law needs to get devices closer together. Metcalfe’s Law needs lower prices to attract more users. Emergent Behavior needs a simple rule to create order from chaos. And all need the clock that Moore’s Law provides.

But the other indications of and reasons for bundling broad-scale semiconductor innovation into the term “Moore’s Law” are that it is far better known outside the semiconductor industry than any of the other laws. It has the benefit of the network effect applied to knowledge transfer: The more times something is repeated, the more people know and believe that it’s a fact rises (unfortunately right or wrong, but in this case, I argue that it is close to right). Every time it is mentioned, there is a far lower probability that the others will be mentioned with it. Plus, it’s very difficult to say them all in one breath and you fall into the explanation fallacy if you do. If you mention them, you’ll have to explain at least one and this takes the conversation off track, deadening any point being made.

Moore's Law Extended through Bundling

- The bundling effect in search result numbers:
 - Moore's Law 1,930,000
 - Metcalfe's Law 67,900
 - Dennard scaling 31,100
 - Koomey's Law 27,000
- Who's going to say the future is here because of Moore's, Metcalfe's, Dennard's and Koomey's Laws?
- Thus 'Moore's Law' is more than scaling

'Moore's Law Dominates Popular Perceptions of Electronics
(Search Results, K, Nov 2017)

VLSIresearch ... intelligence to make better decisions faster

10 November 2017

Copyright © 2017 VLSI Research Inc. All rights reserved. Distribution rights contained in T&Cs.

Figure 6. Moore's Law extended through bundling

Another reason that there's a bundling of laws into Moore's is that there is so much more in his 1965 paper that gets overlooked. In fact, while seldom given credit for it he did make the point that more performance and lower power density important value adds over and above simple scaling.

There are other variant laws often mentioned as well. Moore's 2nd Law — the idea that fab costs are escalating out of control and will kill the industry — was wrongly attributed to him in a Business Week article quoting Bob Helms. Moore later said it was Arthur Rock who noted it in an Intel board meeting and it became known internally as Rock's Law. Separately, the author had noted it in a newsletter in 1990 and Siemens exec Hartwig Ruell later bandied it about as Hutcheson's Law.⁸ Call it, Moore's 2nd Law if you will, the trend has been quoted ever since as a reason for the ending of Moore's 1st Law.

No matter what popular opinion thinks is true, it's clear rising fab costs have not limited innovation. Just look at the copy of VLSI's 1990 chart from Dr. Ruell's slide, and you can see it was a pretty accurate forecast, with TSMC announcing it expected to spend \$20B for its next giga-fab targeted at 3nm. What people missed back then, was that far from it being a limiting factor, the author had seen it as a force for infrastructural change — believing that rising fab costs would force the industry to "share fabs" — as seen at the time with TSMC as the emergent model for this 'sharing.' TSMC was only three years old at the time and Morris Chang has said the term "pure-play foundry did not exist yet."⁹ All of which points to the importance of economics in predicting tectonic structural shifts.

Figure 7. Moore's Second Law hasn't mattered

More importantly, even though fab costs have steadily risen, they've been offset by increasing areal density so that price-per-transistor has steadily decreased. This was shown in figure 1, where there was no cherry picking of data, as it was a price-per-transistor average calculated by dividing all the transistors made in a year by the revenues for that year.

Beyond that, it should be noted that there are many more variants of Moore's Law that need mentioning, including 'More Moore' vs. 'More than Moore?' To understand it, all you have to do is replace the word 'Moore' with 'Transistors' and say it again. More importantly, Paolo Gargini has also told the author that it was obvious that scaling would go vertical once lateral scaling had run out. Obvious, because if you looked at any great city like New York, where the skyscraper came to be, lateral scaling was ultimately replaced with vertical scaling. Those denigrating it to assert their technology is better will argue that Moore's Law is: Tied to the PC, Tied to MPUs. Something too expensive. Not relevant to AI or Mobile. Something of Intel's that should be denigrated by competitors.

The problem with our internal squabbling is that to the outside world Moore's Law is the semiconductor industry's innovation engine. A magical thing that has brought magical products into their lives, a la Steve Jobs. It's become a brand of innovation that can be taken to the bank: Governments fund it. Financial Institutions bet on it.

So, we only hurt ourselves when we denigrate Moore's Law. Those who says it's dead never define anything meaty to replace it with the same power and rely on attacking small parts of it for their own gain. Worse, attacking it implies that that our innovation engine is slowing or stopped. Yet if that is true, why is that so many continue to invest in packing more transistors into a square mm. The answer is that the need for more compute power never abates.

So one must ask: Why do we need a clock? Its true importance is to align individual actors towards a common goal that creates the emergent innovation. Ironically when Carver Mead coined the term, he saw it as a way to spur innovation.¹⁰ As he put it, "This is a law [of] the way that humans are. In order for anything to evolve like our semiconductor technology has evolved, it takes an enormous amount of creative effort by a large number of smart people. They have to believe that effort is going to result in a successful thing or they won't put the effort in. That belief that it's possible to do this thing is what causes the thing to happen. The Moore's Law thing is really about people's belief in the future and their willingness to put energy into causing that thing to come about. It's a marvelous statement about humanity." Or as the author has argued: Moore's Law is Emergent Behavior leading to A Virtuous Cycle: which is defined as a complex chain of events that reinforces itself.

The 4 Laws that Create Virtuous Chip Cycles

— Each inspire action via our innovation model

Figure 8. How the four semiconductor laws create virtuous cycles through emergent behavior.

3. THE ROLE OF LITHOGRAPHY

So, if Moore's Law extended is still driving the industry and rising fab costs have not stopped the industry, where does lithography fit in and what about the steadily increasing cost lithography tools? As you might have guessed, what happens in the lithography bay, doesn't stay in the lithography bay. It affects the entire business model of the semiconductor industry by making Moore's Law possible. If we return to the position made in the introduction:

Now it's clear to anyone in semiconductor lithography that improvements in resolution and overlay are the mechanical reasons why transistors get smaller and density doubles per node — especially when optical lithography is viewed as a system of tools that encompass tools for exposure, resist, deposition, etch, and more. Demonstrating it can be easily done with grade-school math.

It should also be clear that the job of making transistors smaller can only be done practically with lithography. This includes the optics needed to image small dimensions, the photoresist, and all the processing tools from deposition systems used to put down films for hard masks to the etchers needed to etch patterns out of the underlying layers. Moreover, there could be no innovative progress without the brilliant engineers and scientists who painstakingly seek to shave nanometers and angstroms off patterns — while lowering defect density. Think of it like a race car driver who shaves seconds off lap times by shaving milliseconds of each critical section of the track, all the while, optimizing fuel consumption and tire wear to minimize pit stops. With the exception that what's done in lithography is far more amazing ... if unsung. That's the technical side of the equation on a human scale.

On the economic side, we have rising walls of cost, which much be overcome to spur growth. How this works is largely misunderstood due to the misleading shift to view manufacturing economics through the lens of CoO, which has an extremely shallow depth of focus. Shallow, because SEMATECH chose to focus on wafer cost and leave out yield. Hence, CoO is dominantly driven by tool cost and throughput, which are small intermediate steps to returning value in the form of chip sales. Wafer cost is not value. Revenue-per-wafer is value returned for what is spent in making the wafer. Revenue-per-wafer can be increased by 1) lowering defect density so yield is higher and 2) increasing the intrinsic value-per-chip by a: adding transistors, b: improving design, c: finding new applications that can be designed to. Of these, fab tools have a direct effect on 1 and 2a, while having an indirect effect of 2b and c by opening new markets with more transistors and/or

lower power consumption. The modern smartphone is a great example, as it was made attractive by cramming a supercomputer into multi-core chips, adding wireless mobility, while being able to run all day on a battery (most of the time).

A better way to approach this is to use a Cost-per-Good-Die model detailed in “Profitable Solutions to Lithography.”¹¹ This method starts by calculating cost-per-wafer, including all factors plus overhead. Then calculate differing die sizes based on resolution. Then calculate yielded good-die-per-wafer using defect density, die size, and mask counts. Cost-per-Good-Die is then calculated by dividing the number of good die into the cost of the wafer. When these calculations are summed across all process steps, the value returned can be calculated as the difference between the price of the die and the cost to make it. This method is best for evaluating differing types and generations of tools. But it is too laborious and costly to use on an industry-wide basis.

A more cost-effective way to address the problem at hand is to forgo yields and calculate the cost-per-pixel of the differing lithography generations. This has been particularly relevant since the 1990s, when steppers and pellicles became good enough to not add defects. Hence need to calculate yield was minimal at best.

Cost-per-pixel is particularly relevant to the economics of Moore’s Law because transistor density is the result of transistor size, which can be easily defined as a constant times feature-size squared. A pixel for the purposes here is defined as the feature size squared. Then all you have to do is divide the number of pixels that can be fit on a wafer into the cost of lithography and one gets cost-per-pixel. The author first took this approach in the mid-2000s to make the case that lithography tool costs were not out of hand in a proprietary presentation.¹² This was followed-up with an industry-wide presentation the next year.¹³ It has stood the test of time, as since then ASML and others have used this method to make similar claims.

Figure 9. Cost-per-Pixel trends versus Price-per-Transistor as presented in 2006 (Y axis is a Logarithmic scale).

But the real import of this approach, was that it showed how the Litho Cost-per-Pixel had steadily declined with Moore’s Law. This should come as no surprise as Cost-per-Pixel must be a building block of Cost-per-Transistor. Moreover, it must be less-than Cost-per-Transistor, which must in turn be less than Price-per-Transistor in order for positive value to

be achieved. This has to be the case at each and every node. Otherwise, the semiconductor business model breaks down, ending Moore's Law. All three levels are locked together in golden handcuffs.

There was a dark side to that study that's not apparent in the chart: Between 1990 and 2005, the Cost-per-Pixel had declined 33% per year or 58% per node. Price-per-Transistor had declined 30% per year or 54% per node. So while lithography costs were not yet a problem in 2005, they were gaining on Price-per-Transistor. We acknowledged that at the time, stating that a 22nm tool set costing ~140M would need ~250 WPH to not become critical.

So here we are, 13 years later, and does the relationship still hold? Back then, leading edge was 193nm dry. There were multiple NGLs, for which a large contingent of experts thought EUV was impossible — it was still a science project at best. 157nm was in the midst of failing. And immersion was yet to come to the rescue, as the industry circled its wagons in preparation for the end of Moore's Law. Then like an old western, the cavalry came to the rescue: the lithography tool companies beat 22nm, achieving ~250 WPH, while staying reasonably far from the \$100M projected at the time.

Though it is true that EUV tools will cross that \$100M and throughputs won't come close to 250 WPH anytime soon. Using current conservative numbers for what is expected when EUV is producing product wafers in 2019, this method does not paint a pretty picture at the Cost-per-Pixel level. Yet companies are still planning to insert the technology. Moreover, these plans are very real, considering ASML's backlog of EUV tool orders. Models never win a fight with reality, so the question is: What's wrong with the model? Or is the current reality wishful thinking distorted by confirmation bias?

My belief is that it's the model that's wrong because it's missing 3 critical factors: 1) yield is not considered and it's already known that this is a positive for EUV because pattern fidelity and overlay offer superior binning yields, hence revenue is higher. 2) Cycle times are faster due to fewer multi-pattern mask layers. This speeds time to market, which increases revenue over a product's life, as prices are higher earlier in a market cycle. 3) The model inherently assumes that all mask levels are EUV as no average weighting is factored in the fact that Mix-and-Match means only a small portion of the levels carry EUV cost loads. So, the total cost of lithography for all levels is lower, but the entire die gets the full benefit of the EUV layers in its end value.

The 2019 spike smooths out substantially when item 3 alone is accounted for. Moreover, even in the pre-EUV era, lithography costs have been falling faster than price-per-transistor, which explains improved profitability seen in recent years for semiconductor companies.

Figure 10. Cost-per-Pixel trends versus Price-per-Transistor as presented in 2006 (Y axis is a Logarithmic scale).

4. CONCLUSIONS

At this writing in 2018 — it is safe to say that Moore’s Law is alive and well in terms of increasing areal transistor density and decreasing price-per-transistor. More importantly, with respect to the purpose of this paper, the rise in lithography tool prices is being offset by mask-weighted price-per-pixel and thus will continue its contribution to pulling Moore’s Law down its long-term cost curve.

While not addressed here, the revenue effects of yield and cycle time are important factors to consider as further drivers of Moore’s Law. System level economics are another: as the larger share that lower areal cost silicon holds in the finished system means the incremental cost of the most expensive lithography is quite low. This is easily offset by the additional value new versions of the most advanced chips at any time bring to the system. This is an important economic driver in systems company’s renewed interest in designing their own leading-edge chips.

REFERENCES

- [1] From private conversations with the author.
- [2] Moore, G.E., “Progress in Digital Electronics,” Technical Digest of the Int’l Electron Devices Meeting, IEEE Press, 1975
- [3] Moore, G.E., “Cramming more components into integrated circuits,” Electronics magazine, April 1965. A copy of the original draft he submitted can be found at <https://www.chiphistory.org/20-moore-s-law-original-draft-1965>
- [4] Koomey, J., et al, "Implications of Historical Trends in the Electrical Efficiency of Computing", IEEE Annals of the History of Computing, March 29, 2010
- [5] An additional constraint that Gordon himself put on it at the first NTRS roadmap meeting.
- [6] Dennard, R.H., et al, "Design of ion-implanted MOSFET's with very small physical dimensions" IEEE Journal of Solid State Circuits, October 1974

- [7] Attributed to George Gilder by Carl Shapiro and Hal R. Varian in "Information Rules," Harvard Business Press, 1999
- [8] Ruell, H., "Major Trends Shaping the Microelectronics Industry," SEMI ISS Europe, May 18-20, 1992
- [9] Chang, M., Hennessy, J.L., "Morris Chang: An Emphasis on Excellence," Stanford eCorner, May 13, 2014
- [10] Mead, C., Courtland, R., "Q&A: Carver Mead," spectrum.ieee.org, Apr. 15, 2015
- [11] Hutcheson, G.D., Profitable Solutions to Lithography, SEMICON/West Technical Program, 1982
- [12] Hutcheson, G.D., "Is Lithography Affordable?" VLSI Research, 2006
- [13] Hutcheson, G.D., "Litho Tool Costs Out of Hand?" Chip Making Markets, VLSI Research, 2000