the art of boring™

EP 121 | What the end of Moore's Law means for the semiconductor industry

Rob Campbell (00:00):

Coming up on this podcast, I'm joined by one of our equity analysts, Mike Vogel, to do a deep dive into trends and technology impacting the semiconductor industry. Chips have been in the news a lot lately. Biden signed the CHIPS Act back in the summer, we've seen further restrictions on the ability to sell high-end equipment to China, and there's been a ton of investment earmarked for new fabrication plants. In many ways, leading edge technology in these tiny microchips have become a real source of tension globally. Mike and I talk about the impact of those geopolitical considerations toward the end of the conversation. But the real focus here is on the technology itself.

And when I say that we dove deep, Mike takes us right down to the atomic level to understand the chemistry involved in making these ever smaller chips. And then, with that understanding, we zoom out: how will the chemical boundaries that we're fast approaching impact global productivity going forward? What does it mean for various businesses in the semiconductor value chain? And then finally, whether it might be a good chance that we'll continue to be underwhelmed by future iPhone releases. Hope you enjoy.

Disclaimer (01:22):

This podcast is for informational purposes only. Information relating to investment approaches or individual investments should not be construed as advice or endorsement. Any views expressed in this podcast are based upon the information available at the time and are subject to change.

Rob Campbell (01:39):

Mike Vogel is here, and I'm excited for this one. We're here to learn about some of the work, Mike, that you've done into Moore's Law. So, welcome!

Mike Vogel (01:47):

Great. Thanks, Rob. Happy to be back on the podcast.

Rob Campbell (01:50):

Excellent. Now as I understand it, Moore's Law...it's not a law in the same way that we talk about Newton's Laws of Motion or the laws of thermodynamics, it's more of a prediction that's driven really jaw-dropping advancements in technology in the world of computing over the last 50 or 60 years. And so as we dive in here, Mike, can you just share why you chose to do a deep dive into this topic?

Mike Vogel (02:12):

Thanks, Rob. You're absolutely right. So, electronics are everywhere now—they're not just in computers, they're in cars, they're in industrial equipments, they're in our thermostats. And these advancements really have defined technology in modern life. One of the reasons I chose to look at this is there are hard limits to what we refer to as a law; there's real chemical boundaries here. And everyone knew ever since the transistor was invented that eventually we'd hit these limits. It's always just been pretty far off into the horizon. But now, we're looking maybe in the next five years there might be some tough limits as to how well electronics and the conductors can improve. And so, we had the opportunity—we can go into a little more about the top-down scenario discussion we had—where we took a little break from the day-to-day work of looking at companies to try to have some discussions about trends that we see over the next five or 10 years.

Rob Campbell (03:01):

Great. And so maybe just a little bit more on why exactly we do that, or how we do it—are you given sort of freedom as to what you want to explore or topics?

Mike Vogel (03:10):

That's right. I think there's kind of two big issues here that have to be balanced. And if we go to... I know in previous podcasts different people have talked about how we're constantly trying to optimize people's time. So, I think [there are] two big issues we try to find the right balance to. One is we definitely want an independent voice when it comes to research, and that would suggest people should work very autonomously, at least in the first half of investment research. We are a global team and we do want to share ideas as well. How do we balance the two? How do we balance people's time between real independent, autonomous work versus sharing the best ideas across the team? The second thing we try to balance is we do do very careful company-specific diligence. And there's always new products companies are releasing, there's always new announcements, there's always market changes; there's always ideas that we could consider investing in that we're not invested in yet.

So how do we balance time between going through all these company details, with taking a step back and just looking at larger trends? These are two things we're constantly trying to optimize and constantly trying to balance when it comes to people's time.

I think we have a few events per year. This was one of them where instructions were just given to the whole team: take a step back, take any high-level topic you want—it could be about something we're currently invested in, something we're not invested in yet—and just explore it and share it with the team. And so this was one of these things where again, we deliberately moved away from autonomous work and company-specific work to more sharing and a little bit more high-level abstract type research.

Rob Campbell (04:33):

Great. What is Moore's Law?

Mike Vogel (04:35):

If you go through the whole history of digital electronics, the key device is this idea called the transistor. That was invented shortly after World War II. And what people realized soon after was you could make multiple transistors on the same starting block of silicon. Maybe the first one was made out of germanium. But we think about making multiple transistors connected to each other on a single starting block of silicon. What people found was we kept getting more and more and more, and we were making each transistor smaller and smaller and smaller. And what that allowed was much higher performance computing, much more memory. And you hear these stories like you have more computing power in an Apple watch today than the entire planet had in the '50s and '60s. And that's true. And the idea is we kept shrinking this core building block, the transistor, smaller and smaller and smaller. So that's the idea of Moore's Law.

Your next question might be, okay, so why is it ending? I'll let you get to your next question if there's something else. [Laughing]

Rob Campbell (05:27):

I was going to ask two things, maybe just first, why silicon? What's special about that? Are there other materials that could do the same job?

Mike Vogel (05:34):

That's right. So I'll take a step back. The word semiconductor is actually not a very good word. There's real conductors (we think of the copper wiring in our house), and there's real insulators like plastics or different coatings that are on the outside of the wire that prevent us from getting shocked.

What a semiconductor really is, is an insulator that fails when we want it to fail. So if you were to take pure silicon and you were to make it very cold, you'd realize it's just an insulator.

What's nice about it is the way the electrons line up at different energy levels, it's controllable that by room temperature and with different impurities, we can make it a conductor when we want it to be a conductor, but for most of the time it's an insulator. There's other materials people use for very high-powered applications like in electric cars or in transmissions. Like grid transmission, people would use silicon carbide, they could use gallium arsenide. But for electronics that we have around the house, silicon's very cheap, it's very well understood, and it has the right power levels. Everything fits for what we think of as "traditional electronics."

Rob Campbell (06:36):

The next question I think you saw coming was just, well, how precise are we now? How small are we talking? Can you provide us with some context as to why we might be meeting those chemical limits today?

Mike Vogel (06:48):

So one thing before I get into it, industry observers will hear these references to transitions today being made on a three nanometre process or a two nanometre process. There's a very unusual marketing, or it's just kind of an industry practice in the semiconductor industry, what they call a three-nanometre transistor—it's off by about a factor of 15, and that's just for legacy reasons. So you're going to hear me talk about much larger numbers. And what I point out is these industry references? It's well understood that it's just for legacy reasons that they call it these smaller numbers.

So, what sets this limit? I said very pure silicon is inherently an insulator. The idea when we make a transistor is you start with very pure silicon and then you add a very faint amount of some other material. Usually it's phosphorus or boron. And when I say a very faint amount, it's about one in 10,000 atoms at most. At most.

It's usually much less than that. It's probably more one in 100,000, one in a million. A transistor needs three of these regions with these kinds of back-to-back junctions. If you do the math on that, at the fewest number you could have, if the doping level is one per 10,000, would be 3000 atoms. That's a peculiar limit that assumes you can make this thing absolutely perfectly one foreign atom in the middle, and just 9,999 atoms around it.

But we can't make it that way. We make it with these industrial processes that have all different random diffusion around it. The atoms are wiggling all around, so we have to be a few orders of magnitude off from there. If you back this up, and just as a rough rule of thumb, a million silicon atoms as far as their spaced away and apart from each other in a crystal would be a cube of about 30 nanometres. And a nanometre is 1 billionth of a metre.

Now, there's a long series of cascading assumptions I've made having to do with the doping levels, having to do with random diffusion, and just kind of statistics to make this all work. The shape, these things aren't perfect cubes. There's different shapes they can make these out of to sort of save real estate, if you will. That's kind of how you end up with these numbers around in the ballpark of 30 nanometres.

More importantly than this math, there's an industry consortium. It's called the IEEE. They do the standard setting for all the electronics, and every year they get the whole semiconductor industry together.

So these are the semiconductor designers, the tool makers, the chip buyers, academics, and they put out a manufacturing roadmap of what everyone expects is going to happen for the next few years. And if you look at that roadmap this year, they talk about transistors were the mid points, the gates are about 48 nanometres apart. The pace of that shrink is definitely slowing. And in about five years' time, it flattens out. Now, there's other engineering tricks that are going on. There's ways to build transistors on top of each other, but generally speaking, this whole pace is really slowing down. And this is the entire industry getting together and saying, we're kind of reaching the end.

Rob Campbell (09:26):

I'm blown away by just how small we're talking here. The wavelength of visible light tends to range between 300 and 750 nanometres. So we are at a fraction of that. And I just wonder, because I find it so interesting, can you talk about just the technology that we've developed to be able to etch the silicon at this level of precision?

Mike Vogel (09:46):

You nailed it. So visible light, it's about 300 to 700 nanometres. These are made with laser sources that are in pretty deep into the ultraviolet. It's not quite X-rays, but it's way out there. There's all sorts of physical alignment tricks to keep these lined up. And yes, the main advancements really was this whole, particularly on the optical side, getting these laser sources and all these controls handling light, again, far into the ultraviolet spectrum.

Just to back up, the way these semiconductors are made—for older listeners who have ever been in a photography dark room, the idea of having what we call the "film negative," it goes into a machine and you shine a light through it and there's a chemical reaction where the light hits and where it doesn't—in the semiconductor industry, the semiconductor uses that. It creates sort of a stencil and there's different coatings that are put on top of it. And then you do this over and over again and you build a series of patterns and coatings.

But yeah, I mean the precision here is incredible. If we think about the time of year now for our Canadian listeners, when we think about how precise or not precise our door frames and our window frames are, these are off by trillions and trillions of atoms. And here we're counting (again) thousands of atoms. So I mean, this is many, many, many orders of magnitude off from kind of our day-to-day experience.

It's been an incredible success, but it's kind of ending is what's changing now.

Rob Campbell (11:01):

I guess from there, I'm struck by two things. One is, so it's ending, and I think there's a narrative that just Moore's Law and the way that it's developed over the years almost metronomically has been such a driver of productivity and advancement. Even the way that we're recording the podcast today [laughs] is a result of the advancements in Moore's Law. So that's the one side. But then the other side is just, even if it does level off, just the sheer difficulty and jaw-dropping science that goes into actually doing this today, it's still crazy hard to do this, is it not?

Mike Vogel (11:33):

Let me answer the second question first and then I'll get to the first one. It is, and I'd say there's certainly steps of this that there's really only one company in the world that can serve this. And I'm thinking about the lithography stuff that ASML does. That's a true monopoly business. It is very sophisticated. There's only a few locations in the world that are really pushing this limit. But what we wonder is, if you look through the history of the industry, there was this constant leapfrogging. One company would build a new factory with all this equipment, spend billions on it, and then within months that was obsolete because another company came around and built up this new system and they were the leaders now.

So there was this constant leapfrogging and when times were rough, companies would go bankrupt if they were even briefly leapfrogged. What I think is more likely to happen over the next five years is there is this steady state where there's not going to be dozens of competitors, but at least this leapfrogging will go away and there'll be some more consistency about the useful lives of manufacturing equipment that companies buy, and neither the benefits of being the winning leapfrogger nor the pain of getting leapfrogged.

Mike Vogel (12:34):

Now the earlier question…sorry [laughs] can you remind me what the first question was?

Rob Campbell (12:38):

The earlier question was more about my iPhone. I mean, we tend to replace our phones every two to three years, and it was a big gap between the iPhone I and the iPhone X or whatever I've got today. It's a huge deal. But I just wonder what the slowdown in Moore's Law or if we are reaching the end, does it matter as much if I keep my phone for the next four or five years? I'm just not going to see the improvements that we've seen historically. And so just going back to productivity and thinking the economy at large, what does the end of Moore's Law imply there?

Mike Vogel (13:05):

That's what we were hoping this whole research was going to be about. I know the chemistry sounds complicated. I think the chemistry's pretty straightforward. And if the industry participants are putting out roadmaps that kind of say, “yeah, it's slowing down,” I think it's a very trustworthy source. What I try to do is I try to say, “well, what's the impact for global productivity?” And this is a very hard question, and I'll be honest, I shared this with the whole group—I really came up with nothing. [Laughing]

I'm grateful that I work with a team that accepted that as an answer. There's no research we do that we're just trying to create attention or generate trading activity or something like that. What I mean I came up with nothing was, so first there's been this, I guess what some economists call the “productivity paradox” of the last few decades, more like the last decade.

There was a lot of corporate spending on IT, and it just hasn't shown up in the productivity numbers. If you look at these good macro sources like the OECD, you'll find a lot of material about this. This is a little bit of a mystery in the first place. I can give a few of my own ideas of why Moore's Law won't be too disruptive to the global economy. Again, this is something I couldn't find an answer to. I can't put a lot of numbers and precision around this. I think there's a few layers to it. I think first there's a lot of productivity improvements that aren't in what we consider core technology areas. So, if you think about something like big-box retail like Costco, I think most families would recognize that that improved shopping efficiency. That was certainly a productive enterprise.

Do we think of Costco as a tech company? Certainly not. It's just buying in bulk. I think one layer down, there's a lot of opportunities for technological gains outside of what we consider the traditional information technology sector.

So, if you look at something like Tesla—okay, some great computation in there. But things like the batteries and the motors, it's not really semiconductors. There's other power semiconductors in the supply chain. But again, there's other material science that's in there.

Mike Vogel (14:51):

Healthcare, diagnostics, not necessarily having anything to do with Moore's Law. So I think when we look across the world of technology and improvement—and by the way, even within IT, there's a lot of other non-semiconductor stuff. There's wireless coverage. And that's cement and steel antennas all over the place. There's software delivered through cloud services where rather than have all the clients out there doing their own maintenance, just the provider does the maintenance for everyone. That's a real efficiency gain. So, I think there's just many factors that ended up really, I'd say, watering down a little bit the impact of Moore's Law on global GDP. And at least what I shared with the group was this probably is a little bit more an industry-specific consequence rather than a global growth factor.

Rob Campbell (15:33):

Can I go back to just more micro implications? (Although some of these companies are pretty big.) Recognizing that it may take us a few years for Moore's Law to really sunset, I mentioned Apple earlier. What does that mean for Apple?

Mike Vogel (15:47):

So, a few things before I answer the Apple [question]. Moore's Law, let's say it does slow down. There's always new engineering tricks. There's always a few additional things to do. You can build transistors on top of each other. It's not necessarily as thrifty as making them smaller, but you can get more out of it. There's slightly different geometries you can make with it. And then possibly, much longer term, non-silicon materials. Quantum computing. There are potential breakthroughs in, I guess, what we consider digital electronics today. But getting to Apple, yeah, this is tough. I mean, you think about inside an iPhone there's limited real estate for semiconductors, and you have only so much computing power there. And I know there's these jokes going around about how the last iPhone was a little bit underwhelming. And I think in five years it's going to be a lot more underwhelming, unless there's something they can do outside of the processing and the memory.

If there's a whole new tactile system or a whole new camera system, there might be something there. But yeah, the real horsepower under the hood, that real estate is tight. You can't just add twice as much and prices might go up. The cost per area of semiconductors actually has been increasing as the processing became more and more intense. Now, it was compensated for with kind of cheaper per computing stuff. But I think the future iPhones are going to be as underwhelming as they've been, sorry to say it. And maybe even more expensive. Well, plus of some other examples. We have Amazon in the portfolios. Their computing infrastructure that they lease out called AWS, this has been a huge success, great business, great future ahead of it. They've certainly benefitted from constantly having faster, cheaper semiconductors that they would buy and lease out to people. Cheaper memory, cheaper processors.

That might slow down for them. That kind of benefit on the procurement side might slow down with them in five years. I'd caution this isn't the only factor that's affecting these companies five years out. There's plenty of reasons and plenty of different scenarios we can discuss where things can go right and things can go wrong for these businesses within five years and beyond five years. I think this is going to hit those kind of companies—again, you'll start to see the slow down as we approach five years and beyond. Then just the cadence of improvement will moderate.

Rob Campbell (17:45):

You talked earlier about TSMC and ASML specifically, but what about some other companies that are in that semiconductor supply chain or equipment manufacturers?

Mike Vogel (17:54):

I had a really good discussion with Alex on this. So, my first reaction when I was researching this is I said, "Well, this is great for the equipment companies. If you want more computing, if computing demand is strong and you don't get more computing per area, you just need more area." And all these machine companies, they ultimately sell processing time on an area per time basis. You need a lot more of these tools.

And this was an interesting discussion I had with him. He said, "You think computing demand is strong? Let's check that assumption. Is computing demand really strong or do people just want the next updated thing?" So much to say is you can think about computing demand in terms of some sort of combination of bits and speed and storage or something like that. Is that really how demand is created? Or is just when there's a better chip out there, people want the better chip? Does demand drop in half if things really are stagnant for a while?

That was an interesting debate we haven't quite figured out yet. We do have other exposure around there. Just so the listeners know, we don't have a particularly high weight in technology. I mean, we look at all industries. We look at all opportunities equally on the basis of competitive advantage and discount to what we calculate as an intrinsic value of the business. And even within IT, we probably have more software and services today than we do semiconductors.

But yeah, we have KLA in the U.S. midcap equity portfolio. That was one of the companies I was discussing with Alex. You have Samsung in the emerging markets [equity] portfolio. We have a few. [I] just want to make [it clear] if there's any new listeners, we're by no means concentrated in the IT sector just for the sake of any particular growth tilt or unnecessary enthusiasm about the sector.

Rob Campbell (19:22):

Got it. I've got a bigger question, which is semiconductors have been at the heart of geopolitics lately just given their strategic importance in everything that we do. We've heard a lot about China versus the U.S. playing out in Taiwan. What might the end of Moore's Law mean for that looking out over the next five or 10 years?

Mike Vogel (19:40):

The whole development of the semiconductor industry in Taiwan, this was an incredible success. The idea that designers from across the world could come there for the most sophisticated design libraries and manufacturing techniques. I mean, this was really an impressive feat that took decades to grow. But first guess, I think these are two separate issues. If there really is some sort of unfortunate change from the piece we have today, and duplicate capacity needs to be built across various continents, I think that's just its own effect.

Mike Vogel (20:00):

Now, maybe that's a fair question. If things slow down—back to what I said earlier—it's hard to be a sustained leader when the technology is stagnant five years out. Maybe it's a little easier to catch up, but I'm sure TSMC has very good customer relations. It's risky to build a whole new fab without existing clients to get people to come in. So yeah, at first order, I think they're just two separate issues. But taking this to its logical extreme, it'll be easier to catch up to TSMC when there's no consistent technology advances.

Rob Campbell (20:42):

If you're not leapfrogging, it's easier to catch up.

Mike Vogel (20:44):

Exactly.

Rob Campbell (20:45):

Interesting. But still really hard to do. And for those who are interested in the topic, there's a great podcast, the Acquired podcast that did an episode on TSMC, and just fascinating to hear the history of how that all developed in Taiwan, the company specifically. Mike, any last words?

Mike Vogel (21:00):

I hope we get the chance to talk about some of the other ideas we had from this top-down scenarios meeting.

Rob Campbell (21:06):

Well, thanks for coming on, Mike. A fascinating topic. And I think you're right—I suspect we'll dig into a few other topics that came up at that top-down scenarios forum over the next couple months.

Mike Vogel (21:16):

Good.