


January, 2018

# **Menalto Advisors Thought Report:**

## **An Overview of the Semiconductor Industry**



### **Contents Include:**

- Highlights of the semiconductor industry, including growth, issues, recent public company acquisitions, investments, and other topics
- A review of major semiconductor segments including processors, memory, FPGAs, analog, MEMS, wireless, wireline, and optical
- Summaries of recent M&A transactions and activity
- Brief profiles of more than 100 private semiconductor companies

---

**Mark Grossman**  
Managing Director

## Table of Contents

About Menalto Advisors .....	4
Author's Bio .....	5
Other Managing Director Bios .....	6
Previous Thought Reports .....	8
<b>Introduction .....</b>	<b>10</b>
Why a Report on the Semiconductor Industry? .....	10
Report Overview .....	12
A Note on the Private Semiconductor Companies .....	14
We Look Forward to Hearing from You .....	14
<b>Chapter 1: Semiconductor Industry Highlights .....</b>	<b>15</b>
Introduction .....	15
General Semiconductor Market Highlights .....	15
Many Recent Acquisitions of Public Chip Companies .....	22
End of Moore's Law? .....	27
China and Semiconductors .....	30
Fewer VC Investments .....	35
Some OEMs Making Chip Acquisitions .....	43
Semiconductor M&A .....	44
<b>Chapter 2: Semiconductor Technology/Segment Overview .....</b>	<b>45</b>
Introduction .....	45
Semiconductor Technology .....	45
Many Semiconductor Segments .....	56
Processors (and GPUs, DSPs, Microcontrollers, etc.) .....	57

Memory .....	79
FPGAs .....	93
Analog .....	98
Compound Semiconductors .....	109
MEMS Components .....	115
Optoelectronics .....	124
Wireless Communications Chips .....	136
Wireline Communications Chips .....	155
Video/Audio Chips .....	169
Discretes and Passives .....	177
Semiconductor Capital Equipment .....	180
EDA Software .....	185
 <b>Chapter 3: Select Semiconductor M&amp;A Transactions .....</b>	<b>188</b>
Introduction .....	188
Recent Semiconductor M&A Transactions .....	188
Public Semiconductor M&A Premiums .....	205
 <b>Appendix: Select Private Semiconductor Companies .....</b>	<b>209</b>
Introduction .....	209
Brief Summaries of Select Private Companies .....	209

## About Menalto Advisors

### Preeminent Technology M&A Boutique Investment Bank

Menalto Advisors LLC (“Menalto Advisors”) is a new and different type of technology M&A advisory firm located in Silicon Valley. We are a close-knit, diverse group of talented professionals who strongly believe that early partnership and proactive guidance based on each of our client’s unique circumstances are the best way to help them attain their strategic objectives and exceed their M&A goals. We help our clients discover, strengthen, and communicate their capabilities and the potential of their business to optimize each M&A outcome.

Prior to founding Menalto Advisors, our investment banking team completed over 160 transactions around the globe and has deep experience in M&A, engineering, finance, law, management consulting, and entrepreneurship. We understand the dedication, experience, pragmatism, foresight, tenacity, creativity, and exceptional execution it takes to consistently achieve transactional success. Importantly, our partnership approach to M&A works because our team consists of bright, fun, motivated individuals who are passionate about technology and entrepreneurship.

Please visit our website at [www.menaltoadvisors.com](http://www.menaltoadvisors.com) to learn more.



## Author's Bio

### Mark Grossman, Managing Director

Originally from New York City (Brooklyn), Mark has primarily lived and worked in Silicon Valley and Boston. He has more than 20 years of experience in technology and finance, including mergers & acquisitions, technology industry equity research, and engineering. When he isn't traveling the world advising clients, Mark enjoys spending time at home with his son and wife, and playing guitar.



Prior to Menalto Advisors, Mark was a Managing Director at Pagemill Partners and began his investment banking career at SVB Alliant. As an investment banker, he has successfully completed more than 45 transactions (domestic M&A, international M&A, and strategic investments) across a broad range of technology sectors and geographies. Mark has also written a variety of thought pieces on emerging technologies and M&A topics.

Before becoming an investment banker, Mark was a Managing Director and technology industry research analyst at several Wall Street firms. Mark performed in-depth analysis of companies across the tech industry and wrote numerous reports on industry trends and technologies. He received a number of industry accolades, including the Wall Street Journal "Best on the Street" award, and was a guest analyst on a variety of TV financial news programs. Mark began his financial career as an Associate at Goldman Sachs.

Prior to graduate school, Mark worked at General Electric where he managed a variety of complex electronics projects (including both software and hardware technologies). Mark earned a Bachelor of Science in Electrical Engineering, with a Computer Science minor, from Polytechnic University (now NYU Tandon School of Engineering), and both a Master's degree (Electrical Engineering Department) and an MBA (Sloan School of Management) from the Massachusetts Institute of Technology, where he completed a two-year dual-degree program.

Mark's contact information:


M: +1.650.521.1599/T: +1.650.453.5846

[mgrossman@menaltoadvisors.com](mailto:mgrossman@menaltoadvisors.com)

## Other Managing Director Bios

### Nels Nelsen, Managing Director and Co-CEO

Nels was born, raised, and educated in the Bay Area and has spent over 25 years in Silicon Valley advising and working with technology companies as a Certified Public Accountant, lawyer, and investment banker. Away from the office, Nels can be found cheering on his three children and his favorite Bay Area sports teams, at the beach on Kauai, and occasionally on the golf course.



Prior to Menalto Advisors, Nels was a Managing Director at Pagemill Partners, an M&A advisory firm, which he and his partners sold to Duff & Phelps. Nels began his investment banking career at Alliant Partners, which was sold to Silicon Valley Bank. As an investment banker, Nels has closed nearly 100 transactions in a broad range of technology sectors. Nels has represented clients all over the globe, and nearly half of his completed transactions have been cross-border.

As a lawyer, Nels was a Partner at Gray Cary Ware & Freidenrich (now DLA Piper), one of the largest law firms in California, and he was listed in "The Best Lawyers of America." Prior to practicing law, Nels was a Senior Accountant at Price Waterhouse.

Nels earned a Bachelor of Science in Commerce in Accounting from Santa Clara University and a J.D. from Santa Clara University School of Law. He is a member of the State Bar of California (inactive) and a Certified Public Accountant in California (inactive).

Nels' contact information:

M: +1.650.996.6727/T: +1.650.453.5859

[nels@menaltoadvisors.com](mailto:nels@menaltoadvisors.com)



## Charles Welch, Managing Director and Co-CEO

Originally from Michigan, Charles was educated and has lived in both Northern and Southern California and surprisingly has loyalty to sports teams in both parts of the state. Charles is the proud father of two daughters, is a self-proclaimed watch nut, and has a black belt in Aikido.

Charles has extensive experience in mergers and acquisitions, business and corporate development, finance, capital raising, investor relations, human resources, and law. He is a serial entrepreneur. He has been a principal at three different businesses, all of which he sold. Over a 14 plus year investment banking career, Charles has successfully completed numerous M&A, strategic capital raise, and licensing transactions in a variety of sectors and geographies.



Prior to Menalto Advisors, Charles was a co-owner and Managing Director at Pagemill Partners, an M&A advisory firm, where his closed transactions included negotiating the sale of Pagemill Partners to Duff & Phelps. Prior to Pagemill Partners, Charles co-headed several different practice groups at SVB Alliant, the then investment banking arm of Silicon Valley Bank.

Before joining SVB Alliant as an investment banker, Charles was the Vice President, Finance, Corporate Development & Strategy at NetScaler, an Internet infrastructure company that was eventually acquired by Citrix. Prior to NetScaler, Charles held the position of Vice President, Corporate Development at Solectron, an EMS company, and he was the Vice President, Business Development and General Counsel at SMART Modular Technologies, a memory module manufacturer that was acquired by Solectron. Charles was also a corporate securities attorney at Wilson, Sonsini, Goodrich & Rosati.

Charles has a Bachelor of Science in Civil Engineering from Virginia Tech, a Masters in Real Estate Development from the University of Southern California, and a J.D. from the University of California at Berkeley (Boalt Hall School of Law).

Charles' contact information:

M: +1.408.242.1040/T: +1.650.453.5835

charles@menaltoadvisors.com

## Previous Thought Reports

### Highlights of Previous Reports

In case you missed it, below are summaries of four previous Thought Reports. If you would like a copy of one or more of these reports, please contact one of the investment bankers at Menalto Advisors.

- **Emerging Automotive Technologies (ADAS, Self-Driving Cars, V2X Communications, In-Vehicle Connectivity, Electric Vehicles, and More)** – This 180+ page report was issued in July 2017. The report provides an overview of emerging automotive technologies including autonomous vehicles, automotive sensors, V2V communications, in-vehicle sensors, electric vehicles and batteries, user interface technologies, automotive cyber security, and even flying cars. It includes descriptions of more than 100 relevant acquisitions during the past few years, as well as summaries of more than 100 private technology companies addressing the auto tech market.
- **The High-Profile Enterprise Software Acquirers (an Overview of M&A by Oracle, Microsoft, IBM, Salesforce, SAP, Cisco, VMware, HP Enterprise, Adobe, and CA)** – This 140+ page report was issued in April 2017. The report covers M&A activity by the “high-profile” enterprise software acquirers. It includes summaries of more than 260 acquisitions (since 2013) made by the major enterprise software companies, including Oracle, Microsoft, IBM, Salesforce, SAP, Cisco, and others. The report also discusses some of the major M&A highlights, trends, and technology themes across these companies.





- The “High-Profile” Tech Acquirers (an Overview of M&A by Apple, Google, Microsoft, Facebook, Intel, Samsung, and Amazon)** – This 120+ page report was issued in January 2017. The report covers M&A activity by the “high-profile” technology acquirers. It includes summaries of more than 325 acquisitions (since 2013) made by Apple, Google, Facebook, Microsoft, Intel, Samsung, and Amazon. These are the companies that we are most often asked about by private tech company executives and VCs. The report also discusses some of the major M&A highlights, trends, and technology themes across these companies.



- Artificial Intelligence and Machine Learning** – This 130+ page report was issued in November 2016. The first chapter provides a comprehensive overview of the artificial intelligence (AI) and machine learning landscape, including some basic terminology, a history of AI, AI applications, recent trends, strategic and VC investments, concerns, and other related topics. The second chapter contains a detailed primer on AI and machine learning technologies. The third chapter provides summaries of more than 70 AI-related M&A transactions, and the fourth chapter highlights more than 70 private AI companies across a variety of sectors and applications.



## Introduction

*“The microprocessor is a miracle.”*

- Bill Gates, Microsoft Cofounder and Chairman

### Why a Report on the Semiconductor Industry?

Two of our previous Thought Reports were Artificial Intelligence/Machine Learning and Emerging Automotive Technologies. We selected those topics because they were viewed as “hot” segments of the technology industry, each with a growing number of start-ups and significant increases in venture capital investment. In contrast, the semiconductor industry is now generally viewed as mature, with only modest projected growth, and with many VCs no longer interested in the sector. Yet, it remains one of the most critical segments of the technology industry.

The semiconductor industry is a sizable market (over \$400 billion), but more significantly it has been, and remains, the driving force behind the electronics industry. Most of the enormous improvements in performance and cost in computers, smartphones, tablets, servers, game systems, and other electronics have been due to semiconductor technology. As one example, the original IBM PC (introduced in 1981) used 16Kb DRAM chips whereas many PCs now have 16Gb DRAM chips, a density increase of about one million times (at approximately the same price)! Microprocessor performance has similarly also increased by about a million times since then.

While most technologies tend to improve over time, it is difficult to find many examples of non-semiconductor based products that are a million times better than they had been just a few decades ago. As an analogy, if the price-performance of cars improved at a similar pace, a car model from 35 years ago would now cost only a couple of cents. Alternatively, the price of a car would be the same as 35 years ago, but cars would be capable of traveling 100 million miles per hour (around the world in less than one second) and would get about 30 million miles per gallon.

Most of the innovations in the electronics industry during the past 50 years have relied on advancements in semiconductors. Smartphones, for example, would not have been possible without the enormous level of chip integration that has occurred, not to mention the development of MEMS sensors, CMOS image sensors, semiconductor touch controllers, and other chip-related functions.

We previously noted that AI and automotive electronics are hot markets right now, but many of these technologies are possible only because of semiconductor advancements. Most of the core algorithms for AI and machine learning (e.g., neural networks) were developed several decades ago, but there wasn't enough processing power back then to make them useful. It is only recently that processor performance reached the point to make neural networks commercially advantageous. Similarly, many of the emerging automotive technologies (e.g., autonomous driving) would not be possible without the dramatic improvements in processors and memory.

As national and international economic growth is largely tied to productivity gains, and a large portion of global productivity enhancements has been generated from computer systems and electronics, and most of the improvements in electronics are attributable to semiconductor advancements, it could be argued that much of the economic prosperity over the past few decades has been due to the semiconductor industry.

However, despite the importance of semiconductor technology, it is now viewed by many as an outmoded industry. Many of the high-profile Silicon Valley venture capital firms that historically made dozens of chip investments, now only rarely invest in semiconductor start-ups. There have been very few semiconductor IPOs (in the U.S. and Europe) in recent years, and a substantial number of public semiconductor companies have recently been acquired, resulting in fewer public chip companies (which typically occurs in maturing industries).

The industry also faces a number of challenges. Many of the historical innovations in the semiconductor industry have come from emerging private companies (that were often later acquired). However, with fewer start-ups being funded there are concerns that there could be a slowdown in innovation across the industry.

There are also serious concerns that the historical reductions in process nodes (transistors sizes) that have driven semiconductor improvements for the past several decades (“Moore’s law”) may be coming to an end, as transistor sizes are approaching atomic levels. This could have a substantial impact across the entire technology industry.

### **Good Time for a Report on the Semiconductor Industry**

While the semiconductor industry is no longer viewed as a hot industry, it remains a critical one. In addition, while there are certainly fewer chip start-ups relative to 10 or 15 years ago, there are still hundreds of private semiconductor companies, many of which have developed highly innovative solutions and technologies (this report alone highlights more than 100 of them). Based on this, we felt the timing was right for a thought piece on the semiconductor industry.

## **Report Overview**

---

The rest of this report is divided into three major chapters and an appendix:

- **Chapter 1: Semiconductor Industry Highlights** – The first chapter provides a general overview of the semiconductor industry and covers a variety of topics including overall market size and growth rates, major industry segments, top chip suppliers, top OEM customers, recent public company acquisitions, concerns about the end of Moore’s law, strategic and VC investment in private semiconductor companies, semiconductor M&A, and other chip-related topics.

- **Chapter 2: Semiconductor Technology/Segment Overview** – The second chapter begins with a brief overview of semiconductor technology, including explanations of commonly used terms (e.g., transistors, CMOS, logic, analog/digital, integrated circuits) and some of the basic concepts of chip design and manufacturing. It then provides a more in-depth overview of numerous major semiconductor sectors (e.g., processors, memory, analog, FPGAs, MEMS, wireless, wireline, optoelectronics, etc.). Each section includes a technology overview, a discussion of the major players, and brief summaries of several private companies addressing the segment. In total, more than 100 private companies are included and these summaries highlight some of the key emerging technologies across the industry.
- **Chapter 3: Select Semiconductor M&A Transactions** – The third chapter provides an M&A overview including a table of more than 135 semiconductor acquisitions that were announced during the past two calendar years (2016, 2017). The chapter also includes highlights of some of the larger acquisitions that occurred in 2014 and 2015. The last section includes a table highlighting the M&A premiums paid for public company acquisitions during the past few years.
- **Appendix: Select Private Semiconductor Companies** – The appendix includes an alphabetical list and very brief summaries of the private semiconductor companies that were discussed in more detail throughout chapter 2.

## **A Note on the Private Semiconductor Companies**

---

As previously noted, chapter 2 includes brief summaries of more than 100 private semiconductor companies. This was not meant to be comprehensive as there are several hundred private chip companies and, due to space constraints, we could include only a subset of them. We tried to highlight a broad range of different types of companies (some high-profile, some lesser-known). Importantly, we included only public information about these companies, even if we previously met with management and have more detailed non-public information.

We also note that this report was finalized in early January, 2018. As such, some recent news items (e.g., press releases from CES 2018) may not be included in the company profiles.

## **We Look Forward to Hearing from You**

---

### **Comments Welcome**

We hope you find this report useful. We plan to update it from time to time as the industry develops. If you have any comments, catch any errors, or have recommendations for the next edition, please let us know!

### **Happy to Chat!**

Of course, if you know any companies in the semiconductor sector (or any sector within technology for that matter) that may be considering M&A now or in the future, please feel free to have them contact us! We look forward to hearing from you!



## Chapter 1: Semiconductor Industry Highlights

*“In the entire history of the human species, every tool we've invented has been to expand muscle power. All except one...the integrated circuit, the computer. That lets us use our brain power.”*

- David Gerrold, Science Fiction author

### Introduction

---

This chapter provides a general overview of the semiconductor industry. It includes highlights of the market and growth trends, a breakout of the industry by major segments, an overview of the largest chip companies, and a summary of the largest end customers. It contains a list of acquisitions of public semiconductor companies during the past few years and some of the reasons for this activity. The chapter also discusses several important topics such as the potential end of Moore's law and China's interest in the semiconductor industry. As we are often asked by semiconductor start-ups if there is anyone that still invests in chip companies, the chapter also highlights some of the venture capital firms and strategic investors that are still active in semiconductors.

### General Semiconductor Market Highlights


---

#### \$400+ Billion Market

The worldwide semiconductor market was \$338.9 billion in 2016, according to the Semiconductor Industry Association/World Semiconductor Trade Statistics (SIA/WSTS). The 2017 results have not yet been finalized, but it was a very strong year. Specifically, in November 2017, the SIA/WSTS projected that the 2017 semiconductor market would exceed \$400 billion for the first time ever, and reach \$408.7 billion. That represents year-over-year growth of nearly 21%, which is the highest growth rate since 2010. This was driven primarily by memory (Flash and DRAM) which grew about 60% Y/Y in 2017 (the rest of the industry was projected to grow approximately 9%).

## Historical Industry Growth

The graph below illustrates semiconductor industry market size since the mid-1980s. As indicated, the semiconductor market was only approximately \$26 billion in 1986.




Source: SIA/WSTS

From 1986 through 2017E, the semiconductor industry grew at a compound annual growth rate (CAGR) of approximately 9.2%, although the growth rate has generally slowed. For example, the CAGR from 1986 through 2004 was 12.3%, and the general rule of thumb in the 1980s and 1990s was that the industry would likely grow at an average of about 12% to 15% per year. However, from 2005 through 2016, the average growth rate was less than 4%, a substantial change from the past.

As previously noted, the industry experienced strong growth again in 2017 (up over 20%), but that was mostly driven by memory, as Flash and DRAM demand was strong while the industry was capacity constrained. DRAM in particular saw substantial pricing increases (nearly doubling year-over-year in some cases). The rest of the industry (non-memory) was projected to see 9% Y/Y growth, which is better than the recent past, but certainly not explosive. The WSTS is estimating 7% industry growth in 2018 (to \$437.3 billion) and most long-term forecasts project mid-single digit industry growth per year in subsequent years (typically 3% to 5% or so).

The following graph illustrates year-over-year industry growth for the past few decades.



Source: SIA/WSTS

The industry appears to have become a bit less volatile than in the past with fewer wild swings in growth rates. One factor is that the end markets for chips have become more diverse, leading to less dependence on a particular end product (e.g., during the 1990s, the industry was driven primarily by PC sales and PC upgrade cycles). Semiconductor growth has become much more correlated with general economic growth (GDP) than in the past.

Another factor is the transition to the fabless model, as most chip companies now outsource manufacturing to independent foundries. In the “old days” most semiconductor companies had their own fabs and when demand was strong each would invest in significant capacity expansion at the same time. By the time this capacity came online, however, end demand would often decline, resulting in significant overcapacity and price erosion. With a handful of foundries now accounting for a substantial portion of capacity, the industry appears to be in a slightly better position to manage capacity (although there hasn’t been a major economic downturn recently so it remains to be seen what happens when this occurs). One exception to this is the memory sector in which the major players generally own their own fabs, and this segment still experiences significant market swings.

## **Growth Drivers: More Devices and More Complex Semi Content**

Two of the major industry growth drivers have been the proliferation of complex electronic devices and the growing amount of semiconductor content per device. Consumers own a growing number of electronics products (multiple computers/notebooks, tablets, smartphones, home electronics, etc.) which has been a positive for the industry. In addition, electronic products are becoming more complex with a growing level of semiconductor content. Even historically simple home appliances (refrigerators, washing machines) are becoming more sophisticated, requiring more semiconductor functionality.

Overall semiconductor content in electronic products has grown from under 20% in the late 1990s to about 25%, and projections are this will increase towards 30% over the next few years. We noted in a previous report (Emerging Automotive Technologies) that electronics accounted for less than 5% of the content in cars a few decades ago and now represent more than 25%, with some projecting that the percentage will increase towards 50% as electric vehicles and autonomous driving systems become more commonplace.


## **However, PC and Smartphone Markets are Maturing**

Partially offsetting these growth drivers is that many of the largest end markets for semiconductors have begun to mature. For several decades, PCs were the major catalyst for chip sales, but PC sales (including desktops and notebooks) have generally declined during the past few years. Smartphones have been a significant growth driver for the industry since the iPhone was introduced in 2007, but growth has sharply slowed. According to an IDC press release (August 2017), for example, smartphone units are projected to grow at a CAGR of less than 4% from 2016 to 2021 (from about 1.47 billion units to just over 1.7 billion units). The low growth from these large end markets makes overall semiconductor industry growth challenging, even with a variety of emerging segments such as automotive, smart home electronics, and Internet of Things (IoT) devices.

Of course, as always, the industry is under pricing pressure which tends to limit growth. There is some speculation that industry consolidation (as discussed in more detail later in this chapter) could result in less pricing pressure in some semiconductor segments, as there are fewer players in a number of markets. For example, the DRAM market has largely consolidated into three players (Samsung, Micron, Hynix) and, as least in 2017, DRAM pricing was very robust, resulting in strong growth.

## Semiconductors by End Market

Semiconductors are used in an extremely broad range of products. For a number of years, personal computers and servers accounted for the majority of the market for semiconductor chips. More recently, communications (which includes smartphones, as well as other wireless and wireline communications equipment) has become the largest user of semiconductors, exceeding the




Source: SIA/WSTS

PC/server market. While semiconductor content in cars is expected to continue to significantly increase, automotive accounted for less than 12% of the semiconductor market in 2016. The chart above illustrates the breakout of chip sales by end market for 2016, according to a press release by the SIA/WSTS (the breakout for 2017 has not yet been published).

## Semiconductor Segments

The semiconductor market includes many different segments. At a high level, the industry is commonly divided into memory (mostly DRAM and Flash), analog (power management, amplifiers, data converters, and application-specific analog chips), microprocessors (including Intel/AMD processors for PCs and servers, as well as microcontrollers and digital signal processors), logic (a very broad category of chips that includes many different types of digital



Source: SIA/WSTS

and mixed-signal semiconductors for various end markets), optoelectronics (LEDs, laser diodes, image sensors, etc.), discretes (individual transistors and diodes), and sensors (MEMS sensors, silicon sensors). The above pie chart shows the estimated breakout of these segments for 2017 (again, the 2017 results have not yet been finalized as of early January 2018).

Each of these segments has different competitive aspects. For example, memory tends to be a commodity-like market with significant swings in pricing based on supply and demand. In contrast, analog chips are often proprietary with fewer pricing swings and less volatility. Most semiconductor companies focus on only a subset of the overall market.

The WSTS breaks out semiconductor shipments by geography, but it is based on shipment location rather than where the end product is sold (e.g., a chip shipped to Asia for incorporation into a smartphone that sells in Europe would generally be included in Asia not Europe). As most electronics are now manufactured in Asia (and especially China), not surprisingly, most sales are now to Asia. Specifically, the breakout for 2017E is: Asia (excluding Japan) – 61%, Americas – 21%, Europe – 9%, and Japan – 9%.



## Top Customers

A relatively small number of OEMs account for a substantial portion of semiconductor sales. Companies such as Samsung, Apple, Dell, and Lenovo sell a broad variety of products, each of which incorporate many semiconductor chips. According to a press release from Gartner, the table to the right lists the top ten purchasers of semiconductor chips in 2016 (2017 results not yet available). In total, these ten companies accounted for over 38% of industry sales (BBK Electronics is a Chinese OEM that sells the Oppo line of smartphones and displaced Cisco in the top ten in 2016).

Top Buyers of Semiconductors (\$M) - 2016			
Ranking	Company	Amount (\$M)	Share
1	Samsung	\$31,667	9.3%
2	Apple	\$29,989	8.8%
3	Dell	\$13,308	3.9%
4	Lenovo	\$12,847	3.8%
5	Huawei	\$9,886	2.9%
6	HP Inc.	\$8,481	2.5%
7	HP Enterprise	\$6,206	1.8%
8	Sony	\$6,071	1.8%
9	BBK Electronics	\$5,818	1.7%
10	LG Electronics	\$5,172	1.5%
Top 10		\$129,445	38.1%
Others		\$210,238	61.9%
Total		\$339,684	100.0%

Source: Gartner press release (Feb 2017)

## Top Semiconductor Suppliers

According to a press release by IC Insights, the top ten (non-foundry) semiconductor chip suppliers in 2016 were Intel (\$57.0 billion in revenue), Samsung (\$44.3 billion), Qualcomm (\$15.4 billion), Broadcom (\$15.2 billion), SK Hynix (\$14.9 billion), Micron (\$13.5 billion), Texas Instruments (\$12.5 billion), Toshiba (\$10.9 billion), NXP (\$9.5 billion), and MediaTek (\$8.8 billion). Intel accounted for about 15.6% of the total market and the top ten represented over 55% of the market. If foundries were included, TSMC would have been third with nearly \$30 billion in sales.

Although final 2017 numbers have not yet been reported by these companies, IC Insights has projected significant share gains by the memory chip companies as the DRAM and Flash memory markets sharply increased during the past year. IC Insights specifically projected that Samsung would surpass Intel for the first time when the final 2017 results are reported, and that both Micron and SK Hynix will increase share. IC Insights also projected that Nvidia would displace MediaTek in the top ten.

One interesting data point is that in 1985, when there were relatively few semiconductor companies, the top 10 accounted for approximately 60% of semiconductor revenue. This dropped during the 1990s (e.g., it was approximately 45% in 1996). However, due in part to industry consolidation, the percentage has increased in recent years (as noted above, it was approximately 55% in 2016).

## **Many Recent Acquisitions of Public Chip Companies**

---

### **Large Number of Public Semiconductor M&A Transactions**

One of the hot topics in the semiconductor industry has been the large number of acquisitions of public chip companies during the past few years. The semiconductor industry has always been very active with respect to M&A. This historically included a large number of acquisitions of private companies and a small number of public company acquisitions each year. However, since late 2013 there has been an unusually large number of acquisitions of public semiconductor companies.

The number of recent public semiconductor company acquisitions isn't completely "off the charts" as there have been several years in which numerous public chip companies were acquired. For example, in 2009 announced deals included acquisitions of ARC International, California Micro Devices, Hi/fn, Ikanos, Intellon, Semitool, SiRF, Tundra, and Zilog. However, these were mostly relatively small companies. The combination of the number of recent deals, the size of some of the targets (ARM, Altera, Cavium, Linear Technology, NXP, etc.), and the lack of recent semiconductor IPOs, has led to talk about industry consolidation, as there are certainly fewer public chip companies than there had been in 2013.

## Public Chip Company M&A Examples

Some examples of public chip companies that were acquired by other companies since the fourth quarter of 2013 are listed below (including who the acquirer was, the approximate consideration, and the announcement date). Note that a few of these have not yet closed. We did not include the potential \$100+ billion Broadcom/Qualcomm deal as this list includes only transactions that have been agreed to or closed.

- Sigma Designs (by Silicon Labs, \$282 million, December 2017)
- Cavium (by Marvell, \$6 billion, November 2017)
- Imagination (by Canyon Bridge, £550 million, September 2017)
- IXYS (by Littelfuse, \$750 million, August 2017)
- VIXS (by Pixelworks, \$20.2 million, May 2017)
- Exar (by MaxLinear, \$700 million, March 2017)
- Mobileye (by Intel, \$15.3 billion, March 2017)
- GigPeak (by IDT, \$250 million, February 2017)
- InvenSense (by TDK, \$1.3 billion, December 2016)
- E2V (by Teledyne, \$780 million, December 2016)
- AppliedMicro (by MACOM, \$770 million, November 2016)
- NXP (by Qualcomm, \$39.2 billion, October 2016)
- Intersil (by Renesas, \$3.2 billion, September 2016)
- Actions Semi (by Supernova, \$97.3 million September 2016)
- Linear Technology (by Analog Devices, \$14.9 billion, July 2016)
- ARM (by SoftBank, \$32.4 billion, July 2016)
- QLogic (by Cavium, \$1.4 billion, June 2016)
- Anadigics (by II-VI, \$61 million, January 2016)
- Atmel (by Microchip, \$3.6 billion, January 2016)
- Micronas (by TDK, \$217 million, December 2015)

- Inotera (by Micron, \$4 billion, December 2015)
- PMC-Sierra (by Microsemi, \$2.5 billion, November 2015)
- Fairchild (by ON Semiconductor, \$2.4 billion, November 2015)
- EZchip (by Mellanox, \$811 million, September 2015)
- Pericom (by Diodes, \$400 million, September 2015)
- Vimicro (by Vimicro China, \$454 million, September 2015)
- Ikanos (by Qualcomm, \$47 million, August 2015)
- Altera (by Intel, \$16.7 billion, June 2015)
- ISSI (by Uphill Investment, \$765 million, June 2015)
- Micrel (by Microchip, \$839 million, May 2015)
- Broadcom (by Avago, \$37 billion, May 2015)
- Audience (by Knowles, \$129 million, April 2015)
- OmniVision (by Hua Capital, \$1.9 billion, April 2015)
- Vitesse (by Microsemi, \$389 million, March 2015)
- Freescale (by NXP, \$11.8 billion, March 2015)
- Entropic (by MaxLinear, \$287 million, February 2015)
- Silicon Image (by Lattice Semiconductor, \$600 million, January 2015)
- Spansion (by Cypress, \$1.6 billion, December 2014)
- Oplink (by Koch Industries, \$445 million, November 2014)
- CSR (by Qualcomm, \$2.5 billion, October 2014)
- Peregrine Semi (by Murata, \$465 million, August 2014)
- International Rectifier (by Infineon, \$3 billion, August 2014)
- Capella (by Vishay, \$205 million, July 2014)
- Montage (by Shanghai Pudong Science, \$693 million, June 2014)
- Hittite Microwave (by Analog Devices, \$2.45 billion, June 2014)
- ISSC (by Microchip, \$328 million, May 2014)

- Wolfson (by Cirrus Logic, \$488 million, April 2014)
- Integrated Memory Logic (by Exar, \$223 million, April 2014)
- TriQuint (by RF Micro Devices, \$1.6 billion, February 2014)
- Supertex (by Microchip, \$394 million, February 2014)
- LSI (by Avago, \$6.6 billion, December 2013)
- Mindspeed (by MACOM, \$272 million, November 2013)

## Reasons for Consolidation

Although each acquisition was different, it is generally believed that some of the major reasons for the increase in public company acquisitions include:

- **Industry Growth Slowing** – There is often consolidation in industries when growth begins to slow. While 2017 was a good year for the industry, overall growth has been modest for the past decade, which has likely been a factor in the wave of acquisitions.
- **Significant Cost Synergies** – Many of the major chip companies have significant overlap in operations (managing the same foundries) and sales (addressing the same customers). There are also often engineering redundancies, and administrative costs of the target can be cut. In addition, higher volumes can generate better pricing from foundries and assembly/test providers. As such, these acquisitions can produce significant cost synergies. As one example, when Microsemi acquired PMC-Sierra, it announced annual cost synergies of at least \$100 million (nearly 20% of PMC-Sierra's operating expenses).
- **Better Positioned for Top OEMs** – Many of the major end markets for semiconductors are now dominated by a small handful of major OEMs. For example, the top three smartphone OEMs (Apple, Samsung, and Huawei) account for about 45% of smartphone units (and a larger portion of revenue) and the top three PC OEMs (HP, Lenovo, Dell) account for well over half of PC units. Having a broader range of solutions makes a chip supplier more important to these large OEMs, so there are benefits to having a broader chip portfolio.

- **Stock Appreciation** – Historically when a large semiconductor company acquired a smaller chip company, it wasn't clear how investors would view the deal and often the acquirer's stock would decline. However, in many of the previously noted deals, the stock of the acquirer significantly appreciated after the deal was announced. This has provided an incentive for companies to make acquisitions.

If we had to select one deal that appeared to be a catalyst for public-to-public chip acquisitions, it would be Avago's acquisition of LSI at the end of 2013 (right at the beginning of the surge of public acquisitions). Traditionally, chip companies were acquired because of the strategic fit with the buyer (e.g., both companies sold similar or related components). However, Avago and LSI were generally unrelated. That is, LSI primarily sold digital chips for storage and wireline communications, whereas Avago mostly sold analog components for wireless markets. The main driver seemed to be a general interest in "getting bigger." This acquisition helped propel Avago to become a much larger company (both in terms of revenue and market cap) and enabled it to later acquire Broadcom (and become Broadcom Ltd), one of the ten largest chip companies (and Broadcom has recently offered to acquire Qualcomm, one of the five largest chip companies in the world).

### **Few Semiconductor IPOs**

Further driving industry consolidation is that there have been relatively few semiconductor IPOs (in the U.S. and Europe) during the past few years. Ambarella had a successful IPO in late 2012, but since then there have been only a handful of new public chip companies. Some examples include Applied Optoelectronics (fiber optic components, IPO in September 2013), Resonant (SAW filter technology, May 2014), Adesto Technologies (low power memory chips, October 2015), Acacia Communications (optical communications chips, May 2016), Impinj (RFID components, July 2016), Quantenna (Wi-Fi chips, October 2016), X-Fab (specialty chip foundry, April 2017), and Aquantia (high speed communications PHYs, November 2017).



There are a handful of other private chip companies that reportedly are considering filing for an IPO in 2018. Overall, however, the number of new public chip companies is a tiny fraction of the number of public companies that have been acquired, resulting in fewer public chip companies relative to a few years ago. That said, there are still several dozen public semiconductor companies in the United States. In addition, there have been a variety of semiconductor IPOs in China.

## End of Moore's Law?

---

### Potential End to Transistor Shrinking

One of the biggest issues across not just the semiconductor industry, but the entire technology industry is the potential end of “Moore’s law.” Every couple of years or so chipmakers have been able to shrink the size of transistors such that the number of transistors in a given area doubles. The net result is that amount of functionality or performance roughly doubles while cost remains about the same. This trend of doubling transistor density every 18 to 24 months is often referred to as “Moore’s law” (after Gordon Moore, the Intel co-founder, who first noticed this trend), although it is more of an industry target than a law (more details are provided in chapter 2).

Over a period of decades, the compounding effect of these “transistor shrinks” results in enormous improvements in performance and reductions in cost. As previously mentioned, for example, DRAM density and processor performance is more than a million times greater than it had been in the early 1980s at about the same cost.

While there have been design improvements unrelated to Moore’s law, transistor shrinking (and its associated benefits) has enabled the vast majority of the enhancements across the electronics industry during the past several decades. Virtually all of the major breakthrough electronic products during the past fifty years have been enabled by the dramatic improvements in semiconductor technology. As noted in the Introduction chapter, one could argue that economic prosperity is tied to productivity, which has been driven by chip improvements and Moore’s law.

## Many Technical Hurdles Have Been Overcome

For the past few decades there have been many predictions about Moore's law coming to an end, as engineers have faced numerous significant technical hurdles in shrinking transistor size. In each case, the industry has been able to develop innovative solutions to address these problems, and Moore's law has generally continued. As a couple of examples:

- **Current Leakage/FinFETs** – As transistor sizes have shrunk, a major issue has been that even when a transistor is supposed to be “off,” electrical current leaks which results in significant increases in power consumption and heat. One of the ways the industry has addressed this is by developing “multi-gate” transistors (e.g., FinFETs) which help reduce current leakage.
- **Photolithography Optical “Tricks”** – A critical step in semiconductor manufacturing is photolithography in which light and photo sensitive materials are used to create features on a chip during production. Ideally, the wavelength of light used would be much smaller than the desired feature sizes. However, today's transistors are so small that this would require light in the X-ray range, which is not feasible to use. As such, engineers have come up with a variety of complex optical “tricks” such that, amazingly, a 193nm light source is still used even for today's 10nm process nodes.

## However, Soon Facing “Laws of Physics” Issues

While engineers have successfully resolved numerous major technical hurdles during the past couple of decades that some had viewed as insurmountable, it does appear that the industry will soon be facing “laws of physics” issues that could potentially put an end to these improvements. Specifically, the industry is currently at a leading-edge process node of 10nm, with expectations for 7nm ramping this year and a path towards 5nm in the not too distant future. However, the diameter of silicon atoms is on the order about 0.2nm and multiple atoms are required to do anything useful from a transistor perspective. It does not appear possible to shrink transistors to sizes that are close to atomic levels. Therefore, there are reasons to believe that Moore's law may soon be coming to an end.

### 3D Transistor Stacking a Potential Alternative

One potential technology for continuing Moore's law is 3D transistor stacking. That is, currently transistors are implemented on the surface of a wafer in two dimensions. However, it is technically possible to build layers of transistors, adding a third dimension. That is, rather than doubling transistor density by shrinking the size of the transistors, transistors could be kept the same size and the number of layers of transistors could be doubled (in the vertical direction). There have already been some examples of 3D stacking in Flash memory chips and Intel and Micron's XPoint memory uses 3D stacking. However, yields for some of these 3D solutions are still relatively low and it is not yet clear how easy it would be to continue to double the number of vertical layers over time. Even if 3D stacking is possible with memory, which has a very uniform structure, it is unclear how well this could be used for more complex chips such as microprocessors. At a minimum, it would require substantial changes to the manufacturing process, as well as a variety of new design tools.

### What if Moore's Law Does End?

We are a little surprised that the potential end of Moore's law has not received more attention in technology circles and the press. Perhaps it is a "boy who cried wolf" issue in that it had been talked about for decades, but the industry has historically overcome the technical hurdles. As previously noted, however, after another few shrinks it seems like the industry will start coming across some very fundamental issues that could dramatically slow down further shrinking, which could have significant consequences.

Engineers would still be able to make modest improvements to chips and products through better circuit designs and software, and costs could be modestly reduced over time through conventional yield and learning curve improvements. Fabs costs would likely decline a bit as foundries may not have to invest in rolling out new process nodes. There could also be a variety of "system in package" solutions with multiple die per package. However, the enormous, and almost automatic, improvements in chips and electronics during the past few decades could be a thing of the past.

The performance and capabilities of PCs, smartphones, and other electronics might improve only modestly over a decade (rather than by 100X). There could be fewer reasons for consumers and enterprises to upgrade their products over time (much longer upgrade cycles) and it would likely become much more difficult to develop “breakthrough” products. Semiconductor companies with fabs that relied on being at the leading-edge process nodes for a competitive advantage, might find it more difficult to differentiate from the competition. On the other hand, if shrinking transistors becomes extremely difficult, it could provide an even greater advantage for companies that can scale down (e.g., if only one company can get high yields at 3.5 nm, they would have a big advantage). In any case, an end to Moore’s law would certainly have significant ramifications across the entire technology industry and potentially even the world economy.

## China and Semiconductors

---

### Chinese Chip Companies

Another concern across the semiconductor industry is the impact of China, which has made semiconductors a strategic focus, in contrast to most other countries in which chips don’t appear to be a national priority. China-based SMIC (Semiconductor Manufacturing International Corporation) has become the fourth largest chip foundry and several others



are in the top twenty. HiSilicon (the semiconductor arm of Huawei) is believed to have revenue of about \$4 billion, while Spreadtrum (Unisplendour) has revenue on the order of \$2 billion. There are numerous other Chinese chip suppliers with a few hundred million of revenue, many of which are not well-known globally and sell primarily to Chinese OEMs. According to various reports and articles, there are hundreds of semiconductor start-ups in China, although many of these are hard to track.

While China has already built up a significant chip industry, it has expressed strong interest in becoming a much larger player in the market. In 2015, China announced that it planned to invest up to 1 trillion yuan (about \$160 billion at the time) over a ten-year period to develop its semiconductor industry. Since then, there have been a variety of other similar announcements, as well as anecdotal information about each of China's provinces looking to obtain and invest in semiconductor technology.

## China Semi Acquisitions

To accelerate the ramp-up of its semiconductor industry, China (through various state-sponsored entities or private equity firms) made several acquisitions in 2015 and early 2016. However, with concerns over China buying up critical technologies, the U.S. Government has taken a more aggressive stance on these acquisitions and has either delayed or rejected several deals through CFIUS (the Committee on Foreign Investment in the United States). There have also been reports that Europe may begin taking a more aggressive approach towards these types of deals.

## China Semiconductor Acquisitions

Some examples of China-related acquisitions and attempted acquisitions of U.S. or European semiconductor companies include:

- **OmniVision** – In April 2015, OmniVision (a leading supplier of CMOS image sensors) agreed to be acquired by a group of Chinese investors led by Hua Capital for \$29.75/share or about \$1.9 billion. The investors previously offered \$29/share. The transaction closed in early 2016.
- **ISSI** – In June 2015, Integrated Silicon Solutions, Inc. (ISSI) agreed to be acquired by China-based consortium Uphill Investment for about \$765 million. This occurred after a long bidding war with Cypress Semiconductor over several months (originally ISSI agreed to a \$639 million offer, but Cypress then made a higher bid, and there were multiple counterbids). ISSI provides memory chips (SRAM, DRAM, Flash) and typically targets niche segments that the largest suppliers don't address. The transaction closed in December, 2015.

- **Micron** – In July 2015, China's state-owned Tsinghua Unigroup offered to acquire Micron (DRAM and Flash memory) for \$23 billion. However, due to concerns over CFIUS review and other issues, Micron rejected the offer and a deal was never agreed to.
- **Lattice** – In November 2015, Lattice Semiconductor agreed to be acquired by Canyon Bridge Capital Partners for \$8.30 per share (about \$1.3 billion, net of debt). Canyon Bridge is technically based in Silicon Valley, but one of its primary investors is reportedly the Chinese government. Lattice is the number three supplier of FPGAs and programmable logic solutions (behind Xilinx and Altera), but had also acquired Silicon Image (HDMI and interface chips). Although this deal was first announced in late 2015, CFIUS approval was delayed several times and CFIUS ultimately rejected the deal in September 2017. As such, the transaction was never completed and Lattice remains an independent public chip company.
- **Anadigics** – In November 2015, GaAs Labs agreed to acquire Anadigics, but reportedly it was subsequently outbid by a Chinese acquirer. This led to a bidding war, which new entrant II-VI eventually won. According to press articles, concerns over CFIUS was one of the factors that led the Anadigics board to select II-VI.
- **Mattson** – In December 2015, Mattson Technology agreed to be acquired by Chinese P/E firm Beijing E-Town Dragon for about \$300 million. Mattson provides semiconductor processing equipment, including dry strip, etch, and rapid thermal processing tools. The transaction was completed in May, 2016.
- **Aixtron** – In May 2016, Germany-based semiconductor capital equipment maker Aixtron agreed to a \$723 million offer from China's Fujian Grand Chip Investment Fund. However, several months later the deal was blocked by the United States and Germany. As such, a transaction was not completed.



- **NXP's Standard Products/RF Power Units** – In June 2016, NXP agreed to sell its Standard Products unit (mostly discrete components) to China-owned investment firm Jianguang Asset Management (JAC Capital) and private-equity firm Wise Road Capital for about \$2.75 billion. The deal was completed in early 2017. NXP previously divested its RF Power unit to JAC Capital in 2015 for \$1.8 billion.
- **Imagination Technologies** – In September 2017, UK-based Imagination Technologies (primarily GPU and multimedia IP cores) announced that it was being acquired by Canyon Bridge for about \$675 million, although the deal excluded Imagination's MIPS unit which was separately sold to Tallwood Venture Partners. As previously noted, it is believed that Canyon Bridge's main investor is the Chinese Government. This deal closed in November 2017.

In addition to these, there have been several take-private acquisitions of China-based Nasdaq-listed public chip companies by entities in China. For example, Spreadtrum (RF components), RDA (RF components), Vimicro (multimedia processors), and Actions Semiconductor (mobile device chips) were all acquired by Chinese entities. Another example was Montage, a China-based chip company that focused on analog and mixed-signal components for set-top boxes and memory modules. Montage had an IPO on Nasdaq in 2013. However, in early 2014, its stock sharply dropped as a result of allegations from short-selling research firms that Montage's executives also controlled Montage's primary distributor (LQW) and was using LQW to artificially boost Montage's financial results. Montage denied the allegations, but subsequently failed to file its 10-K report. In June 2014, Shanghai Pudong Science and Technology Investment Co., a state-owned limited liability company, agreed to acquire Montage.

There have been other technology deals called off due to CFIUS. For example, China-based Unisplendour was planning to make a nearly \$3.8 billion investment in Western Digital, but the deal collapsed over U.S. government concerns. Western Digital was going to use the funds for its acquisition of SanDisk, but ended up buying SanDisk anyway.

Inseego agreed to sell its mobile hotspot unit to TCL, but eventually called the deal off due to “delays and uncertainty in securing CFIUS approval.” Some deals have been approved by shareholders but are still awaiting CFIUS approval. For example, Xcerra (which provides test equipment including equipment for testing semiconductors) announced in April 2017 that it agreed to be acquired by a Chinese investment firm. However, CFIUS approval has still not yet been provided (as of early January 2018). Given the uncertainties around CFIUS approval, many U.S. chip companies have become reluctant to sell to Chinese buyers (without substantial breakup fees if a deal isn’t approved by CFIUS) and Chinese acquirers have become less aggressive in trying to acquire U.S. chip companies.

### China Capacity Concerns

Another China-related concern is capacity. With China looking to make substantial investments in semiconductors and now generally unable to acquire large U.S. chip companies, and with each province in China looking to independently add semiconductor manufacturing capabilities, there are concerns that an enormous amount of capital may go into building new semiconductor fabs, without much concern over return on investment.

China has already announced several major new upcoming fabs (e.g., in early 2017, China’s state-owned chipmaker Tsinghua Unigroup announced plans to build an enormous \$30 billion memory-chip factory in Nanjing). These investments could potentially result in significant excess manufacturing supply across the chip industry for many years. While excess capacity is clearly a negative for foundry suppliers, there is an argument that it could be a positive for fabless chip companies (lower costs). Historically, however, overcapacity causes a variety of serious industry-wide problems and is generally a negative for most chip companies. Some believe China’s strategy may be to drive down pricing and make the industry less economically attractive for non-Chinese companies, enabling chip companies in China to gain significant share. In any case, the potential incremental chip manufacturing capacity in China over the next several years is a major concern within the industry.

## Fewer VC Investments

---

### Fabless Model Led to Strong Growth in Chip Start-ups

Another major issue across the industry is the decline of venture capital funding for semiconductor private companies, as many of the major venture capital firms (outside of China) now generally avoid semiconductor companies. As many of the innovations in the chip industry have come from start-ups, there are concerns this could reduce breakthroughs.

The semiconductor industry has always been entrepreneurial with numerous start-ups. However, up until the mid-1990s, chip companies generally had to have their own fabs. This required a meaningful amount of capital and expertise in both chip design and manufacturing. However, the emergence of 3<sup>rd</sup> party foundries in the late 1990s (such as TSMC, UMC, and others) allowed outsourcing of manufacturing and eliminated the need to own a fab, significantly reducing the barriers to entry for starting a new chip company. This resulted in a huge increase of “fabless” chip start-ups.

For many years, each year there were typically 50 to 100 new private chip companies and well in excess of \$1 billion invested in private semiconductor companies (e.g., according to data from PitchBook, well over \$1 billion was invested in chip start-ups in at least five of the eight years from 2000 through 2007). However, this declined sharply beginning in 2008 with investment dropping to about half of previous levels and much fewer new start-ups.

One could argue there were far too many private chip companies. At one point, for example, there were more than 50 private chip companies addressing the Wi-Fi market (Atheros went public, a few were acquired at high value, but the vast majority had unattractive exits or were shut down).

### Much Less VC Investment Since 2008

VC investment in semiconductor companies plummeted during the 2008/2009 recession and remained lackluster for many years. Some reasons for this include:

- **More Challenging to Compete** – A general view is that it has become more challenging for chip start-ups to compete. Many of the largest OEMs that account for a sizable portion of end markets are reluctant

to use chips from start-ups. Additionally, start-ups are often at a disadvantage to larger companies that offer broader ranges of ICs, can bundle chips together, and receive lower pricing from foundries.

- **Rising Costs** – The fabless model greatly reduced the cost to start a new chip company. However, this has reversed in recent years. As chips have become more complex (especially digital chips), they require larger and larger design and verification teams. Mask costs for leading-edge processes have significantly increased and EDA design tools are expensive. Chip companies now need sales and support around the world, and in many cases semiconductor companies have to develop complete reference designs of end products to demonstrate their solutions for customers. As a result, some chip start-ups raised twenty-five to fifty million dollars just to get a couple of chips into the market, which is a major negative for many investors.
- **Alternative Sector Investments** – As the costs for chip start-ups increased, VCs were drawn to other sectors, such as Internet and social media, which often require only a relatively modest investment to get up and running. These companies may also ultimately require significant capital for marketing and expansion, but the investments can be made over time as the business ramps up, rather than making a large bet upfront. Of course, these sectors have their own issues, but the general view for many VCs has been that the level of risk to potential reward is better in these sectors than in semiconductors.

### Fewer Start-ups, but Still a Significant Number

While the number of new semiconductor start-ups has declined from a decade ago, there are still a significant number of them. According to PitchBook, there have been over 110 new semiconductor private companies funded since the beginning of 2013 (about 22 per year). That is far below the level at the peak, but on the other hand reports about the demise of semiconductor start-ups have been exaggerated. These numbers also exclude many Chinese chip start-ups that generally aren't well-tracked (reportedly there are hundreds of them). This report alone includes profiles of more than 100 private semiconductor companies.

There are also signs of growing interest in semiconductors. According to a CNBC article in September 2017, CB Insights projected that private semiconductor companies were poised to raise approximately \$1.6 billion in 2017, up from \$1.3 billion in 2016 and \$820 million in 2015 (these numbers don't match PitchBook's, and are "apples" and "oranges" with some of the numbers noted above, but show an upward trend). Much of the increase was reportedly due to investments in AI-related chip companies.

In a number of cases, start-ups have focused on less capital-intensive segments of the market such as analog or sensors, which typically require much less investment. While general semiconductor investing is down from the boom years, there are certain "hot" chip segments (such as artificial intelligence/machine learning, automotive, and quantum computing) in which there have recently been numerous significant venture investments.

### Still Some Large Capital Raises

While many chip start-ups have taken a capital-efficient approach with smaller funding rounds, there are still a small number of private companies that have raised substantial amounts of capital during the past few years. As previously noted, a number of these are in "hot markets" such as automotive, AI, or quantum computing. Some examples include:

- **Aquantia** – Raised \$37 million in 2015 (total of more than \$100 million). It sells high-speed PHYs for networking. IPO in late 2017.
- **Autotalks** – Raised a \$40 million round in 2017. It sells chips for vehicle to vehicle (V2V) communications in cars.
- **Barefoot Networks** – Reportedly raised \$130 million in 2015/2016. It developed an advanced programmable Ethernet switch chip.
- **Butterfly** – Reportedly raised over \$100 million. It develops MEMS ultrasound chips and complete miniature ultrasound scanners.
- **Cavendish** – Raised a \$36 million round in 2015. It sells MEMS-based RF switches for smartphones and other markets.
- **Cerebras** – Reportedly (per Forbes) raised \$112 million in three rounds. It is a stealth-mode AI processor company.

- **Crossbar** – Raised a \$35 million Series D round in 2015 (total of \$85 million raised). It sells resistive RAM (ReRAM) memory chips.
- **Graphcore** – Raised a total of about \$80 million in 2017. It is focused on designing AI/machine learning processors.
- **LeddarTech** – Announced a \$101 million capital raise in 2017. It provides chips and software for automotive LiDAR.
- **Nantero** – Raised a total of \$51 million in 2015/2016. Nantero has developed carbon nanotube based memory chips.
- **Rigetti** – Raised a total of \$64 million in series A and B rounds. Rigetti is developing quantum computing processors.
- **ThinkForce** – Raised \$68 million in a Series A round in 2017. The company is developing an AI processor and AI algorithms.
- **Transphorm** – Raised \$70 million in 2015. It sells gallium nitride (GaN) chips for high power applications.
- **Valens** – Raised a \$60 million round in 2017. It sells communications chips including those for HDBaseT for automotive networking.
- **Vayyar** – Announced a \$45 million capital raise in December 2017. It uses wireless technology to capture object images without cameras.

### Some Positives for the “Survivors”

One positive for chip start-ups is that while it can be challenging to raise capital, there are often fewer competitors than there might have been in the past. As mentioned, at one point in time there were 50 private Wi-Fi chip companies, whereas now there are often only a couple of start-ups addressing a particular market. In addition, from an M&A perspective there are fewer options for public acquirers, which can increase the value of a private company (e.g., if there is only one private company addressing a particular segment and there are several public companies interested in the technology, then the start-up is well positioned for an attractive M&A exit).



## Select VC Firms That Still Make Chip Investments

The top-tier U.S. venture capital firms historically made many investments in semiconductor start-ups, but now do so only very rarely. There are some exceptions. For example, Sequoia recently invested in Graphcore (AI processor), Barefoot Networks (networking ICs), ThinkForce (AI processor), and Quantum Circuits (quantum processing); and Kleiner Perkins has invested in Ambiq (microcontrollers), mCube (MEMS), Crossbar (resistive memory), and Transphorm (GaN). However, semiconductors account for a very small portion of most top-tier U.S. venture capital firms' portfolios.

There are, however, still a handful of VCs that are fairly active in making semiconductor investments. Many of these are based outside the U.S. Some examples are listed below. In addition, most of the private company profiles throughout chapter 2 include the firms that invested in the companies, which provides some additional information regarding who invests in the semiconductor industry.

- **Atlantic Bridge** – Atlantic Bridge has made a variety of recent investments in semiconductor companies including AnDapt (configurable power management), HiLight Semi (high speed TIAs for optical communications), Lion Semi (power management chips), and Navitas (GaN power chips). It also previously invested in several chip companies that were acquired including BridgeCo, Movidius, Ozmo, and Silicon Blue.
- **BDC** – BDC primarily invests in Canadian companies, but has made a number of semiconductor-related investments, such as Diablo (memory interface), D-Wave (quantum processor/computer), GaN Systems (GaN power devices), and LeddarTech (optical detector chips and software for automotive).
- **Capital-E** – Capital-E has made investments in eoSemi (temperature compensated crystal oscillators), Merus Audio (efficient audio amplifiers), Silicon Line (VCSEL drivers, TIAs, and SerDes for optical communications), Silicon Mobility (automotive EV and hybrid motor control chips), and SureCore (low power SRAM IP).



- **CEA Tech** – CEA is based in France and has invested in a number of chip companies including Apix Analytics (lab-on-a-chip), Asygn (analog design software), Crocus (magneto-resistant chips), EnerBee (energy harvesting), ExaGaN (GaN power chips), Kalray (multicore processors), Presto (semiconductor testing), and Ulis (IR detection). It was also an investor in Tronics (MEMS foundry, acquired by TDK) and Soitec (silicon on insulator technology, public).
- **Gemini Israeli Ventures** – Gemini is based in Israel and has made investments in Autotalks (automotive communications chips), Colorchip (integrated optical chips), and Sckipio (G.fast DSL chips).
- **IPV Capital** – IPV is based in China with an office in Silicon Valley and has made a variety of investments in semiconductor companies in both China and the U.S. Some investments include Brite Semiconductor (ASICs), Giantec (EEPROM), GigaDevice (non-volatile memory), Maxscend (RF components), and Senodia (MEMS).
- **Lux Ventures** – Lux has invested in Everspin (MRAM, public), FlexLogic (embedded FPGAs), Luxtera (optical transceivers and modules), Nervana (AI processor, acquired by Intel), Transphorm (GaN power devices), and other chip companies.
- **NEA** – NEA historically made many semiconductor investments. This has sharply slowed down in recent years, but some examples of recent investments include Aquantia (high speed PHYs, recent IPO), Luxtera (optical transceivers and modules), PsiKick (low power wireless and energy harvesting), and Sentons (touch interface).
- **New Science Ventures** – New Science is based in New York and has made several semiconductor investments. Some examples include: Achronix (FPGAs), Ferric (power management chips), RF Arrays (RF front-end solutions), and Vorago (high-reliability ICs).
- **Pitango** – Pitango is based in Israel, but it has made investments in chip companies around the world. Some examples include: Celeno (Wi-Fi chips), Graphcore (AI processors), Kilopass (memory IP), and Sckipio (G.fast DSL chips).

- **Silicon Catalyst** – Silicon Catalyst is a Silicon Valley based incubator focused specifically on early-stage semiconductor start-ups. Through partnerships, it provides portfolio companies with a variety of in-kind goods and services (EDA tools, foundry wafers, test equipment, etc.), as well as mentoring from a variety of experienced semiconductor executives. It has more than a dozen semiconductor investments including: Ayer Labs (miniature fiber optic transceivers), Aeponyx (micro optical switch semiconductor chips), REX Computing (highly efficient processor), and On Silicon Chip Photonics (MEMS sensors).
- **Walden** – Walden has made numerous semiconductor investments, including chip companies in both China and the U.S. Some chip investments include: 3PeakIC (analog chips), Aquantia (high speed PHYs, recent IPO), Credo (high speed SerDes), GalaxyCore (CMOS image sensors), GigaDevice (memory chips), Innophase (RF chips), Lion Semiconductor (power management chips), MEMSDrive (MEMS actuators), Silicon Mitus (power management), and VeriSilicon (ASICs and IP), amongst many others.
- **WestSummit Capital** – WestSummit has invested in AltoBeam (chips for TVs in China), GigaDevice (non-volatile memory), Movidius (vision processors, acquired by Intel), Tiler (multi-core processors, acquired by EZchip/Mellanox), and VeriSilicon (ASICs and IP).

## More Strategic Investments

While the number of VC firms investing in semiconductors has sharply declined in recent years, some of the slack has been taken up by corporate investors. A variety of public semiconductor companies and OEMs have made multiple recent strategic investments in semiconductor-related companies. Some examples include:

- **Bosch** – Bosch's venture arm has made several investments in semiconductor companies. Some examples include EpiGaN (GaN epitaxial wafers for power devices), Pyreos (MEMS IR sensors), and XMOS (microcontrollers for audio applications).

- Cisco** – Select Cisco investments include Alchip (ASIC supplier), AnDapt (configurable power management), Aquantia (10GBase-T PHYs and 100G chips), Celeno (Wi-Fi chips), CnexLabs (solid state storage controller ASIC), eSilicon (chip services), Netronome (server-based networking NICs and chips), Wilocity (60GHz chips, later acquired by Qualcomm), and Tileria (multi-core CPUs, which was acquired by EZchip, which was acquired by Mellanox).
- Dell Ventures** – Dell has invested in several chip companies including Barefoot Networks (software defined networking ICs), CnexLabs (solid state storage controller ASIC), Diablo Technologies (memory technology), Graphcore (AI processors), and Nantero (carbon nanotube based memory).
- Intel Capital** – Intel Capital has been an extremely active investor, although despite being a semiconductor company only a small portion of its investments have been in semiconductor start-ups. Nevertheless, it is one of the most active semiconductor investors. Some examples include: AnDapt (configurable power management), Aquantia (10GBase-T PHYs and 100G chips, recent IPO), Chronocam (differentiated CMOS image sensors), Grand Chip Microelectronics (Wi-Fi front-ends), Keyssa (high speed, short distance wireless), Netronome (server-based networking NICs and chips), NextInput (MEMS-based touch force sensing), and Trigenice (audio chips).
- Qualcomm Ventures** – Some investments include Cavendish Kinetics (MEMS RF switches), SJ Semi (foundry in China), and Wilocity (60GHz chips) which it eventually acquired.
- Samsung** – Samsung has multiple investment vehicles (Catalyst fund, Samsung Ventures, etc.). Some of its chip investments have included: ALT (driver chips for LEDs), Autotalks (V2X communications), CnexLabs (solid state storage controller ASIC), Graphcore (AI processors), Keyssa (high speed, short distance wireless), Leman Micro Devices (blood pressure and vital sign sensors), and Valens (high-speed communications chips for consumer and automotive).

- **UMC Capital** – UMC is a major Taiwanese foundry and has made a variety of semiconductor investments. A few examples include: Floadia (memory IP), Ineda (wearable processing units), NextInput (MEMS force sensing), and SiFotonics (silicon photonics).

## Some OEMs Making Chip Acquisitions

---

### OEMs Focused on Chips

A few of the major technology OEMs have expanded their semiconductor teams, and occasionally acquire chip companies. While these acquisitions are rare, we are often asked about them. Some examples include:

- **Amazon** – Amazon acquired Israel-based Annapurna Labs in early 2015 for reportedly about \$360 million. Annapurna developed a highly efficient ARM-based processor for servers and networking, although reports indicate it was still essentially pre-revenue.
- **Apple** – Although Apple generally doesn't sell semiconductors externally, many believe that it has one of the largest chip teams in Silicon Valley. Apple has also made a handful of chip-related acquisitions including Passif Semiconductor (Bluetooth Low Energy chips), PA Semiconductor (high performance ARM-based embedded processors), and a Maxim analog semiconductor fab.
- **Cisco** – Cisco has made several chip acquisitions. Most recently, it acquired Leaba Semiconductor (networking chips) for \$320 million. Another example was CoreOptics (DSP-based optical communications chips) which it acquired in 2010.
- **Google** – Google acquired Lumedyne (MEMS accelerometer sensors for indoor navigation) and Agnilux (low power embedded processors).
- **Juniper** – In 2016, Juniper acquired Aurrion (fabless silicon photonics start-up) for approximately \$165 million.
- **Western Digital** – Western Digital acquired SanDisk, a leading supplier of Flash memory controller chips and solutions.

## Semiconductor M&A

---

### Semiconductor Deals in 2016 and 2017

The semiconductor industry remains highly active with respect to mergers and acquisitions. Chapter 3 includes a list of semiconductor M&A transactions (including acquisitions of both public and private companies) for 2016 and 2017. For each transaction, a brief overview of the target's products is provided. When available, valuation data is included. The chapter also includes highlights of some of the larger semiconductor M&A deals that were announced in 2014 and 2015. This provides a sense for semiconductor M&A activity during the past several years.

### Semiconductor M&A Premiums

Chapter 3 also includes a list of acquisitions of public semiconductor companies with each deal's premium (generally relative to the target's stock price just prior to the announcement). As indicated, the median premium for the listed deals is 27%, and the mean is 32%.

## Chapter 2: Semiconductor Technology/Segment Overview

*“The semiconductor business, it's very competitive. You're out there every day slugging it out, no different from players slugging it out on the ice.”*

- Henry Samueli, Broadcom Cofounder

### Introduction

---

This chapter begins with a brief overview of semiconductor technology and explains some of the basic concepts for those that are not already familiar with the industry (e.g., what is a semiconductor?, what is a transistor?, what is CMOS?, etc.). The rest of the chapter discusses several of the major segments of the industry including microprocessors, memory, FPGAs, analog, MEMS, wireless chips, wireline communications chips, optoelectronics, and discrete components. For each section, the report provides a short overview of the technology and recent trends, as well as summaries of several relevant private semiconductor companies, which we believe provides further insight into emerging technologies. In addition to semiconductor chip segments, the chapter also includes brief overviews of the semiconductor capital equipment and EDA sectors. In total, more than 100 private semiconductor-related companies are summarized. The appendix includes an alphabetical list of these private companies.

### Semiconductor Technology

---


#### Intro to Semiconductors

Semiconductor materials have properties that fall between those of conductors (which conduct electrical current well) and insulators (which are very poor conductors). That is, semiconductors can conduct, but only under certain conditions, which makes them extremely useful (e.g., they can be turned “on” or “off”).

Although the term “semiconductor” technically refers to the material, electronic components (microprocessors, DRAM, etc.) that are made from semiconductor materials are generally referred to as semiconductors.

The most common semiconductor material is silicon, although there are many others. Often, impurities (dopants) are added to the semiconductor materials to enhance their electrical properties. For example, adding phosphorous to silicon results in a significant increase in negatively charged free electrons, creating an “n-type semiconductor.” Similarly, adding boron results in a sharp increase in positively charged “holes” (atoms missing electrons), creating a “p-type semiconductor.”

To provide an overview of how a semiconductor operates, consider a semiconductor material in which one part is p-type and one part is n-type (this is shown on the top of the diagram to the right). At the junction in the middle, the free electrons in the n-type combine with the holes in the p-type to create neutral atoms which act as a barrier (depletion region). Subsequently, nothing much happens and no current flows (it is “off”).



In the middle diagram, a positive voltage is applied to the n-type region and a negative voltage to the p-type region. The electrons in the n-type region are pulled away from the junction towards the positive voltage and the holes in the p-type region are pulled away from the junction towards the negative voltage. Nothing crosses the junction, so no current flows (light is off).

In the bottom illustration, the voltage is reversed. The electrons in the n-type region are pulled across the junction to the positive battery voltage, and the holes in the p-type region are pulled across the junction to the negative battery voltage. As a result, current flows (and the light turns on). Thus, the device can be turned “on” or “off” based on the applied voltage.

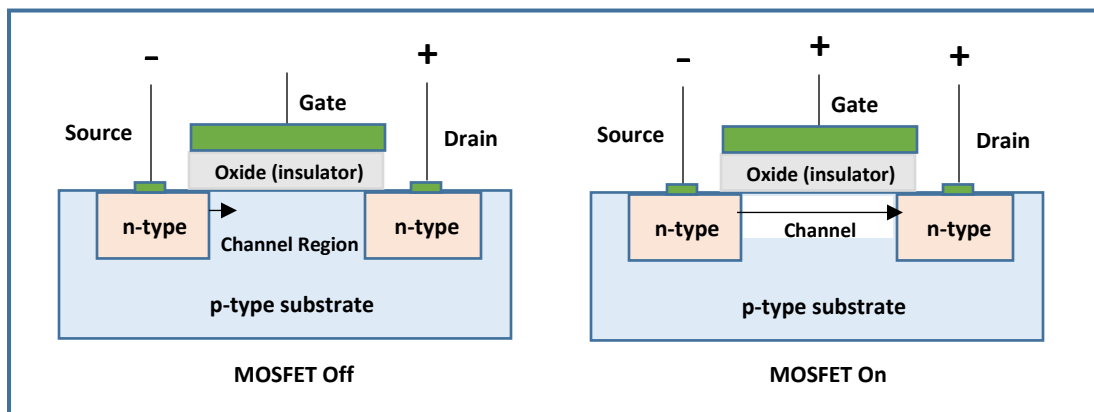


This simple “PN junction” device is also called a diode. It is used in many electronic circuits. One of its uses is it enables current to flow in one direction and not the other. Diodes made from certain materials emit light when turned on and these are referred to as light emitting diodes (LEDs).

## Transistors (MOSFETs)

There are a variety of different types of transistors, but the most popular is a MOSFET (metal oxide semiconductor field effect transistor). A MOSFET has three alternating doped regions (either npn or pnp) instead of two. For example, the diagram below shows an npn MOSFET.


In the diagram on the left, the free negatively charged electrons in the n-type source would like to travel across the transistor to the positive voltage applied at the drain (the plus sign above the drain). However, the free electrons from the source are unable to make it across the transistor, as they combine with the positive holes in the light blue p-type region and are essentially blocked. As a result, no current flows and the transistor is off.



Source: Menalto Advisors

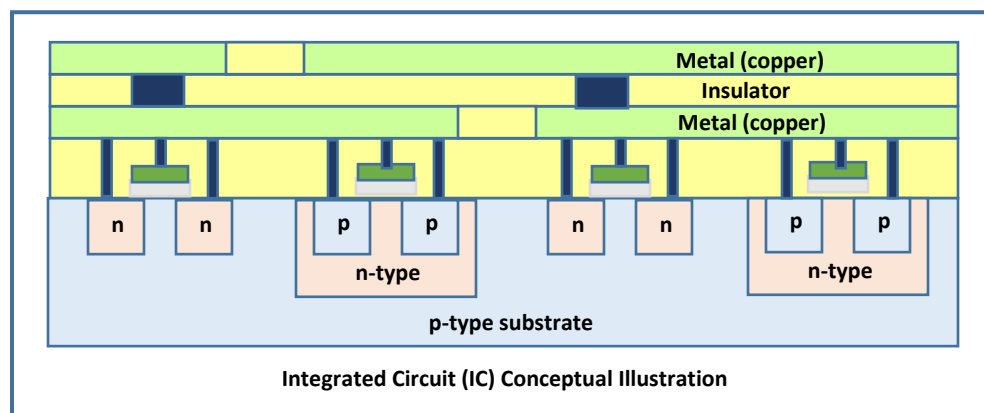
However, if an appropriate voltage is applied to the gate (the plus sign above the gate as shown in the diagram on the right), an electrical field is created that repels the positive holes in the p-type substrate. This creates an open channel between the source and drain (the white region in the diagram) in which free electrons can travel. As a result, there is a flow of electrons (current) from the source to the drain, and the transistor is on. Thus, by changing the voltage on the gate, the transistor can be turned on and off. Some types of transistors (e.g., FinFETs) have two gates to help prevent current from flowing when the transistor is supposed to be off.

The “M” in MOSFET is for metal and the gate of a MOSFET (green area shown in the previous diagram) was historically metal. However, most MOSFETs now use polysilicon rather than metal. The insulator region under this is typically silicon dioxide ( $\text{SiO}_2$ ). The photo to the right is an example of a real transistor. The actual semiconductor material is inside the black “package” and isn’t visible, but you can see the three “legs” that connect to the source, gate, and drain.



## Integrated Circuits (ICs)


One of the most important historical breakthroughs in the semiconductor industry was the development of the integrated circuit (IC). The idea behind an integrated circuit is rather than producing individual transistors and connecting them together with wiring on a printed circuit board, many transistors (potentially billions) can be implemented on a single chip (a single piece of silicon). In addition to the transistors, an integrated circuit also incorporates the “metal wiring” to connect the various transistors together. This is accomplished by adding layers of metal (as well as insulation layers between the metal layers, along with vias to connect different metal layers). This creates conductive paths between various transistors.



Source: Menalto Advisors


This is conceptually illustrated in the above diagram. At the bottom there are four transistors (2 npn and 2 pnp). In addition, the transistors are connected to various metal (copper) layers, with the dark blue “plugs” (typically tungsten) attaching metal layers together. In this way, the transistors can be connected together in any desired manner. [This is meant to illustrate the concept of an IC, as these actual connections wouldn’t make sense.]

The previous diagram showed a piece of silicon with four transistors, but by interconnecting millions or billions of transistors together, very complex functions can be created (e.g., microprocessors, microcontrollers, etc.) which would be nearly impossible using individual transistors. Integrated circuits now represent the vast majority of the semiconductor market, although there is still a sizable market for discrete transistors and diodes. The photo to the right shows an example of an integrated circuit (the actual silicon is inside the black package and is connected to the metal legs).



## Digital Logic and CMOS

Transistors have many uses, but an important one is that several transistors can be combined to create a “logic gate.” A logic gate produces an output based on a logical combination of inputs. For example, an “AND” gate produces a “1” output only if both inputs are “1”. Similarly, an “OR” gate produces a “1” output if either or both of its inputs are “1”. A single OR logic gate could, for example, be used to cause an alarm to activate if a door is open OR a window is open.



By combining multiple gates, mathematical operations (e.g., binary arithmetic) can be implemented, and by using thousands or millions of gates, very complex functionality can be created.

There are a variety of different ways that transistors can be combined to create logic gates and other functions. By far, the most popular method is referred to as CMOS (complementary metal oxide semiconductors). CMOS uses pairs of MOSFETs (one npn and one pnp, which is why the term “complementary” is used) to form logic gates. Relative to other types of digital logic families, CMOS provides many advantages including low power consumption (it generally consumes power only when switching), good noise immunity, and it can more easily scale to smaller sizes than most alternatives. As such, it is now pervasive across digital semiconductor chips.

While CMOS technically refers to a specific family of digital logic, in some cases the term “CMOS” is also used to refer to the manufacturing process used to produce these types of chips (even if the chips don’t actually use CMOS digital logic).

## Digital versus Analog

The vast majority of semiconductor chips (microprocessors, memory, etc.) are digital, meaning that they process and store bits of data (“0”s and “1”s). In a digital chip, a “1” could be represented by a transistor that is on whereas an off transistor could represent a “0”. Alternatively, a bit could be represented as a voltage relative to a threshold. For example, a voltage above 2.5 volts might be considered a “1” and a voltage below 2.5 volts considered a “0”. Thus, it doesn’t matter if the voltage is 2.79 volts, 3.58 volts, or 4.23 volts, it is still a “1.” In contrast, analog chips deal with the continuum of values. An analog amplifier might, for example, double its input voltage, so if the input is 1.43 volts, the output is 2.86 volts (the actual voltage matters). Analog chips are described in much more detail in the Analog section. The term “mixed-signal” is used to refer to integrated circuits that incorporate both digital and analog on the same chip.

## Chip Design and EDA Tools

Designing a complex digital chip typically requires large teams of engineers and a variety of software tools. Engineers will typically describe the desired functionality using a hardware description language (such as Verilog or VHDL). EDA (Electronic Design Automation) software can then be used to help convert this high-level description to a detailed logic gate level design. Once the design is validated and optimized, other EDA tools are used to help create a detailed physical design (i.e., a detailed layout of the various transistors and interconnections). Additional software tools are used to help engineers ensure the chip meets power and timing requirements (EDA tools are discussed in more detail in a later section). After many iterations, eventually the design is finalized (“tape out”). The final design output is a detailed physical description of the chip (often in a format called “GDSII”) which can be used for manufacturing.

## Chip Manufacturing Process


The semiconductor manufacturing process is extremely complex with thousands of steps, many at high temperatures and pressures. Most of the processing occurs at a manufacturing fabrication facility referred to as a “fab.” For leading-edge digital manufacturing processes, fabs now typically cost several billion dollars (e.g., Samsung broke ground on a 10nm fab in 2015 that reportedly cost about \$14 billion). The inputs to a fab are blank semiconductor wafers (typically silicon). Utilizing the GDSII files, the fab performs a variety of processes to implement a given design. Some major steps in manufacturing include the following (more details are provided in the capital equipment section later in this chapter):

- **Deposition** – Deposition involves depositing a layer of material across the surface of a wafer. This can be done using chemical reactions (chemical vapor deposition) or physical means (physical vapor deposition). Deposition can be used to deposit metal, polysilicon, and other materials to help create a semiconductor chip.
- **Photolithography** – With photolithography, the wafer is covered with a material that is sensitive to light (photoresist). An intense light is then directed onto the wafer through a mask (reticle), which is transparent but has opaque patterns on it and is designed to create a desired pattern on the surface of the wafer. The light passes through the transparent portions of the mask (but is blocked by the opaque portions). The light that isn’t blocked causes a chemical change in the photoresist it reaches, which enables it to be easily removed. The net result is a pattern of photoresist that is the same as the opaque pattern on the mask. This can be used to create a variety of patterns on the wafer during the manufacturing process.
- **Etch** – Etch tools remove materials. Etching is typically performed after photolithography, as the photolithography step leaves a desired pattern of photoresist and the etching tools can, for example, etch away at areas that aren’t covered with photoresist (e.g., the photoresist material protects areas from being etched away).


- **Ion Implantation** – Ion implantation produces ions and accelerates and directs them into a target material. It is commonly used for semiconductor doping to create p-type and n-type regions on the surface of a wafer (creating patterns of transistors).

By repeating the above process steps many times (as well as a variety of other process steps not discussed here), transistors and layers of interconnections can be produced to implement a given design. As previously shown, the various transistors are first implemented on the surface of the wafer and then alternating layers of metal and insulating materials are deposited to create the connections among the transistors.

The outputs from the fab are finished wafers, each of which may contain hundreds of copies of the same chip. For example, in the photo on the right, each rectangle on the wafer is referred to as a “die” and is a separate chip. If this were a microprocessor wafer, each die would be a separate microprocessor. The production of these semiconductor wafers is often referred to as the “front-end” of the process.




The finished wafers are then typically shipped to separate facilities for assembly and test, which are referred to as the “back-end” of the process. This includes dicing up the wafer into individual “die” or chips. It would be extremely difficult to directly use these small pieces of silicon on a circuit board, so next they are placed into “packages” which protect the semiconductor device. The silicon die is typically wire bonded to the package to enable communications with other electronics through the metal “legs” of the package. This is illustrated in the figure to the right. The actual silicon “chip” is the square in the middle, which is then wired to the metal contacts around it such that input/output signals are transmitted through the legs of the device (normally the top is opaque not transparent).





Lastly, final testing is performed on the packaged chips to ensure they work properly. After passing final test, the semiconductor chips can be shipped to OEMs for incorporation into equipment. Typically, OEMs incorporate the chips onto printed circuit boards (PCBs) along with other chips and components. The photo to the right shows several packaged chips on a circuit board (the silicon chips are inside the black packages).



### Fabless Model

Prior to the mid-1990s, most semiconductor companies had their own fabs. However, during the past couple of decades, 3<sup>rd</sup> party “foundries” that focus solely on semiconductor manufacturing (e.g., TSMC, GlobalFoundries, UMC, SMIC, etc.) have become increasingly commonplace and most chip companies are now “fabless” and outsource manufacturing. The exceptions are some of the larger semiconductor companies (e.g., Intel, Samsung, Micron, etc.) and many of the analog or non-silicon companies, as analog fabs tend to be much less expensive than digital fabs.

### Moore’s Law and Process Nodes

Moore’s law was already discussed in chapter 1, but this sub-section describes it in a little more detail. As previously noted, Moore’s Law was an observation made by Gordon Moore, the co-founder of Intel, during the mid-1960s about improvements in semiconductor manufacturing technologies. Specifically, he noticed that the number of transistors per integrated circuit was approximately doubling every year and projected that this would continue for a number of years into the future. The “law” has been revised a number of times since then, and is more commonly given as “doubling every two years” or “doubling every 18 months.” However, even at 18 months, it results in an increase of more than a million times over a 30 years period. Moore’s law has generally been viewed as a guide for the industry in terms of improving manufacturing, rather than a rule of nature. A large portion of this improvement comes from significant investments across the industry to shrink the sizes of transistors.



Semiconductor manufacturing processes are often characterized by their process node. This can conceptually be thought of as the length of a transistor (historically the process node referred specifically to the gate length of a transistor, but the term has recently become more ambiguous due to foundry marketing). Process nodes were historically typically reduced by about 30% every 18 to 24 months. For example, some common process nodes have included 250nm, 180nm, 130nm, 90nm, 65nm, 45nm, 32nm, 22nm, 14nm, and 10nm. As one example, Intel's Core i8 microprocessors use a 14nm process and it has announced 10nm versions for 2018. There are already plans for rolling out 7nm in 2018, and expectations are that 5nm processes will be deployed over the next few years. Only certain chips use the leading-edge process nodes and many use older processes.

Reducing the length of a transistor by about 30% approximately doubles the number of transistors that can be implemented in a given area. For example, a transistor with a 10nm length might take up an area of  $100\text{nm}^2$  whereas if that same transistor had a 7nm length it might have an area of only  $49\text{nm}^2$  (less than half the area). This can double the number of chips per wafer, which greatly reduces cost per chip, as the cost of processing a wafer doesn't significantly increase (once yields are high). Alternatively, rather than reducing the size of a chip, the chip size can be kept the same with much more functionality. For example, the amount of memory on a chip can be doubled while keeping the chip size and cost about the same.

In addition to enabling greater functionality and reduced costs, moving to smaller process nodes provides many other advantages, such as increased chip speed (since the transistors are closer together, the chips operate faster). Historically, it also helped significantly reduce power consumption, although more recently this has not always been the case which has become an issue across the industry.

One additional way chipmakers have reduced cost over time has been through increasing wafer sizes, which provides more chips per wafer. For example, some wafer size transitions have included shifting from 100mm (about 4") to 150mm (about 6") wafers, to 200mm (about 8") wafers, and then to 300mm (about 12") wafers. Each of these transitions significantly increased the number of chips per wafer.

For example, the area of a 300mm wafer is 2.25 times more than a 200mm wafer, so the number of die per wafer more than doubles (because of the shapes of the rectangular die and circular wafer, the number of die actually increased by about 2.4x). While processing a larger wafer is more expensive, overall chip costs are significantly reduced due to the much larger number of chips per wafer (once yields of the new equipment improve). This is in addition to any improvements from smaller transistor process nodes.

However, increasing wafer size typically requires a completely new set of processing equipment and initially yields are very low until the equipment is optimized, which can take several years. Historically, wafer size increases occurred about every 10 years, but the industry has been stuck at 300mm for much longer than normal (300mm manufacturing began in the early 2000s). An issue for semiconductor equipment companies is that development of new equipment is very costly, and increasing wafer size actually reduces their market (chipmakers need less equipment). Another issue is that chipmakers are reluctant to be the first to use the equipment since yields are initially poor and any “lessons learned” are incorporated into 2<sup>nd</sup> generation tools that benefit other chipmakers as well, so there is no incentive to be first. Overall, it does not appear that the transition to 450mm wafers is imminent.

As previously noted, there are concerns that Moore’s law may be coming to an end soon as the size of transistors are quickly approaching the size of atoms. While there appears to be a path for at least a few more process node reductions (7nm, 5nm, 3.5nm, etc.), the industry may soon reach a point at which it is no longer possible to further shrink transistors, or these shrinks may take many years rather than just a couple of years. This could have significant negative implications across the technology industry.

## SoCs

One of the most significant trends in the semiconductor industry during the past several decades has been integration. As transistor sizes shrink, more functionality can be squeezed into a single chip. What previously required many chips can subsequently often be accomplished by one. Originally this typically involved integration of digital logic functions, but over the past couple of decades designers have become more adept at integrating many different functions (processors, memory, analog, RF, etc.) into one chip.

In some cases, a large portion of the functionality in a system can be performed by one chip and this is often referred to as a system-on-a-chip (SoC) solution. These SoCs typically incorporate one or more processors as well as a variety of analog, memory, and digital functions.

In some cases, it is not feasible or cost effective to incorporate all major functions on one chip, but to reduce size and cost multiple bare die are incorporated into the same package (from the outside it looks like one chip and it eliminates the need for separate packages for each chip). For example, a processor and memory might both be included in the same package. This is often referred to as “system-in-a-package (SIP).”

## **Many Semiconductor Segments**

---

### **Broad Range of Different Chip Segments**

The semiconductor industry consists of many different segments, each of which has its own set of challenges and competitors. The rest of this chapter discusses several different segments of the industry. We note that a number of the highlighted segments overlap. For example, there are sections on analog, compound semiconductors, and wireless. However, many compound semiconductor chips are analog and a significant percentage are used in wireless applications. We also note that not every segment of the industry is covered. The main focus isn't to segment and cover every sector of the market, but to hit on some of the key major markets and emerging technologies within the industry.

## Processors (and GPUs, DSPs, Microcontrollers, etc.)

---

### Intro to Microprocessors

Microprocessors are the main processing chips in many types of electronics. At a basic level, a microprocessor simply reads software instructions that are stored in memory, executes them, and then moves on to the next instruction.

A simple microprocessor will include a program counter (which keeps track of the memory location of the next instruction to be executed), an instruction register (where the software instruction to be executed is stored), several general registers (which are used to temporarily store small amounts of data), and an arithmetic logic unit or ALU (which can perform simple logical or mathematical operations on data).

Based on the value in the program counter, a microprocessor will obtain a software instruction from the appropriate location in memory. This instruction is loaded into the microprocessor's instruction register. Using logic functions within the microprocessor, the software instruction is read and executed. These instructions are typically very basic. For example, an instruction might be to add the number stored in one register to another register (in which case the ALU is used to add the numbers) or to copy the number stored in a register into memory. After an instruction is executed, the program counter is incremented and the next software instruction is executed. Although each microprocessor instruction is relatively simple, by executing billions of them per second, very complex tasks can be performed. The number of instructions a single processor can execute per second is related to its frequency (a 4GHz processor might execute about 4 billion instructions per second), although with multi-core processors it can be a multiple of this.

The set of instructions that a microprocessor can execute is referred to as its architecture. Software written for one processor architecture won't run on a processor with a different architecture (unless there is some type of intermediate software layer to "translate," which typically slows processing down). However, the underlying processor design and implementation can significantly vary while still supporting the same architecture. For example, Intel's PC processors have used the "x86 architecture" for decades, but Intel has dramatically enhanced its PC microprocessors over time.

## Microprocessor Features

The description above provides a simplified explanation of how a microprocessor works. However, modern processors have a broad range of advanced features to enhance processor performance. For example:

- **Pipelining** – Instead of waiting until one instruction is loaded and fully executed before loading the next instruction, pipelining uses an “assembly line” approach. As a simple example, while one instruction is being executed, the next instruction can be fetched from memory in parallel. This is a simple “two-stage” pipeline example, but real processors can have many pipeline stages.
- **Superscaler** – Superscaler technology involves a single processor (or a single processor core for multi-core processors) executing more than one instruction at a time, which can increase the number of executed instructions per second and enhance performance. However, it also requires having multiple redundant functions on the processor (e.g., to perform computations in parallel) and it adds complexity since some instructions can’t be performed in parallel (e.g., if one instruction requires the result of a previous instruction, they cannot be performed at the same time), so dependencies have to be checked first.
- **Multi-threading** – Multi-threading enables a processor to execute software instructions from several different threads of software. Each thread could be from a different application or a different part of the same application. This is complex, as the processor must store and keep track of all of the data associated with each software thread as it switches between threads (e.g., the data stored in the registers will be different for each thread). As an analogy, if you switch back and forth between reading two books, each time you switch you must remember not only the last page and word you read, but all of the information you learned up until that point in the book.

With coarse grained multi-threading, a processor will process one software thread until it runs into some type of delay, and will then switch to another thread. With fine grained multi-threading, the processor can switch threads each clock cycle (every new instruction).

“Simultaneous multi-threading” combines with superscalar technology to enable multiple threads to be executed simultaneously (in parallel). This essentially allows one processor to “appear” as if it is multiple processors to the outside world. Multi-threading can significantly enhance the efficiency of a microprocessor.

- **Cache Memory** – To enhance performance, microprocessors typically incorporate memory (usually SRAM) directly onto the processor itself. Registers can store only very small amounts of data (a few bytes each), so having memory in the processor for storing and retrieving data and instructions without having to access main memory (which takes much longer) can speed up processing. This type of intermediate fast memory is often referred to as “cache” memory. Modern processors typically include several levels of cache memory (e.g., L1, L2, L3). For example, a processor might have 64KB of L1 cache that it can access in 1ns, 256KB of L2 cache that it can access in 2ns, and 8MB of L3 cache that it can access in 4ns. The processor attempts to keep the data and code it needs most often in the lowest level of cache memory.
- **Multi-Core** – One of the biggest changes in microprocessor design over the past 10 to 15 years has been the shift to multi-core designs in which multiple processors are implemented on a single chip. For many years, the primary focus for processor suppliers was increasing processor frequency (how many instructions per second a single processor could execute), but in many cases performance can be significantly enhanced by using multiple processor cores rather than a single blazingly fast processor. Fully taking advantage of multi-core processors, however, requires dividing up the instructions in such a way that they can be performed in parallel, which adds complexity.

Whether it is multiple processor cores within a single chip or multiple processors within a system, there are different methods for implementing multiple processors. With symmetric multiprocessing (SMP), for example, the processors or processor cores share a single main memory. An issue with this is that as the number of processors increase, memory access can become a major bottleneck.



An alternative is massively parallel processing (MPP) in which each processor (or processor core) has its own memory. The challenge with MPP is data used by one processor might be needed by another, and so communications among the various processors is often a bottleneck. Some solutions use a hybrid of both approaches.

- **Integration** – Many microprocessor solutions now integrate a variety of other functions besides the processor itself. For example, some of Intel's PC processors incorporate a graphics processor. Most microprocessors also integrate a variety of interfaces to communicate with other chips.
- **Power Control** – With the proliferation of mobile devices, there is great interest in reducing power consumption of processors. As such, many processors now incorporate a variety of sleep and power reduction modes to turn off portions of the processor when not in use.
- **CISC vs. RISC** – Over time, the instructions that a processor could handle became more complex, enabling a processor to accomplish more with each instruction. This was expected to improve efficiency and performance, and is sometimes referred to as CISC (complex instruction set code). However, decoding and executing complex instructions takes more time, and the overhead associated with this can slow a processor down even when executing simple instructions (and often the more complex instructions are not used very often). As such, an alternative processor design approach was developed called RISC (reduced instruction set code) in which simple instructions are used and executed very quickly. Another characteristic of "RISC design" involves incorporating a large number of registers. There isn't a strict definition of whether a processor is RISC or CISC but the terms are often used to describe certain processors (e.g., ARM's original name was Advanced RISC Machines).



## Microprocessor Example (Intel Core)

To provide a concrete example of a real microprocessor and its features, in October 2017 Intel formally introduced its 8<sup>th</sup> generation Core desktop processor line (previously nicknamed “Coffee Lake”). These processors are produced on a 14nm process and include three families: the Core i3 (which incorporates four processor cores that each operate at up to 4GHz, and include up to 8MB of L3 cache), the Core i5 (which has four processor cores that operate at up to 4.3GHz with 9MB of L3 cache), and the Core i7 (which has six processor cores that operate up to 4.7GHz with 12MB of L3 cache). The Core i7 also includes Intel’s “Hyper-Threading” technology, which is their version of simultaneous multi-threading (as previously discussed) and allows each processor to essentially act as two separate processors, enabling up to 12 software threads to run simultaneously (6 cores x 2 threads each).

The i5 and i7 processors also include a technology called “Turbo Boost” that enables the processor to run at a much faster than normal frequency if conditions allow (based on power consumption, temperature, and workload). It also integrates an Intel graphics processor (UHD), various memory interfaces (DDR4 DRAM), and PCI Express 3.0 interfaces.

Intel also provides x86 processors for notebooks and servers. The notebook processors tend to run at lower frequencies and consume less power. For example, the i5 and i7 notebook processors are quad-core processors that consume only 15 watts (versus up to 95 watts for some of the desktop processors) and include a variety of power reduction features and sleep modes. Intel’s Xeon processors for servers are larger, more expensive chips. For example, the Xeon E7-4830v4 processor incorporates 14 cores (versus 4 or 6 for desktops) and includes 35MB of L3 cache.

## Microprocessor Background

The microprocessor was invented by Intel in the early 1970s. Nippon Calculating Machine (NCM) Corporation engaged Intel to design a 12 chip solution for a printing calculator. However, Intel decided to integrate most of the non-memory functions into a single-chip central processing unit (CPU). Intel later acquired the rights for the processor from NCM, and launched the 4004 microprocessor in 1971.

Microprocessors soon became pervasive and a variety of other companies introduced processors. In fact, by the late 1970s microprocessors from companies such as IBM (for use in its mainframe computers) and Motorola were generally viewed to be superior to Intel's. However, IBM decided to use an Intel processor (the 8088) in its initial IBM PC computer, which was introduced in 1981. The 8088 processor was a derivative of Intel's previous 8086 processor, and so the architecture (instruction set) of these processors is often called the "x86" architecture. PC software written for the x86 architecture won't run on other processors and so Intel's microprocessors became the standard across the PC industry. Although Intel developed many subsequent PC processors (80286, 80386, Pentium, Core, etc.) the newer processors continued to use the x86 architecture. That is, although many new additional instructions were added over time, the processors generally continued to support the historical x86 instructions to ensure backwards compatibility with older software programs.

### AMD x86 Processors

The only significant competitor Intel has had in the PC processor market has generally been AMD. AMD initially licensed processor designs from Intel and acted as a second source supplier. In 1985, however, Intel stopped licensing its designs to AMD, and AMD had to develop its own processors that could handle the Intel x86 instruction set such that any software that could run on Intel's processors would run on AMD's processors. There were a couple of brief periods (late 1990s, mid-2000s) when AMD gained substantial market share from Intel, but in general AMD tends to have a much smaller portion of the market.

AMD's current major processor lines for desktops and notebooks are Ryzen (including Ryzen Pro and Threadripper) and the lower-end A-Series processors. Its server processor line is Opteron and it has a product focused on high-performance computing (EPYC 7000). As one example, the AMD Ryzen 7 1800X processor includes eight processor cores, each of which can handle two threads (16 threads in total). The clock speed is 3.6GHz, but it includes a Turbo mode that can boost that to 4GHz if conditions allow. The Threadripper is a higher-end processor for gaming that includes up to 16 cores and more cache memory.

## ARM-Based Embedded Processors

The previous processor subsections focused on stand-alone processor chips used in products such as PC and notebook computers. However, one major trend over the past 20 years has been the emergence of ARM-based processing solutions in a broad range of products. ARM doesn't actually sell microprocessor chips, but designs microprocessors and then licenses those designs to other semiconductor companies in exchange for upfront fees and royalties. In most cases, the ARM processors are designed into larger chips that have a variety of other functions. ARM (which was acquired by Softbank) specifically designs its processors for low power consumption and for easy integration into larger chips. The ARM processors have become pervasive in many markets. For example, most of the major smartphone processing chips incorporate one or more ARM-based processors.

ARM's solutions include several families of processors including the Cortex-A (targeting complex computing), Cortex-R (targeting real-time processing), and Cortex-M (targeting low power mobile), and each family includes several different processors. ARM's DynamicIQ technology enables clustering of multiple different ARM processors in the same chip. There has been growing interest in ARM-based microcontrollers and several companies have developed ARM-based server solutions. In late 2017, Microsoft introduced Windows 10 for ARM, which includes an x86 emulation layer (can translate x86 instructions to ARM instructions) and Qualcomm has indicated its ARM-based Snapdragon 835 chip can be used for mobile PCs. As such, there may be growing support for ARM-based PCs.

While ARM has been the most successful provider of processor IP, there are others that license processor designs including Synopsys (ARC), Cadence (Tensilica), and MIPS (which Imagination divested to Tallwood).

## RISC-V Processors

With ARM processors so pervasive, there has been growing interest in alternative processing cores. One of these developments is RISC-V, which is an open instruction set that can be freely used for any purpose. It was originally developed at the University of California, Berkeley but is now promoted by the RISC-V Foundation, whose members include Google,

Qualcomm, Samsung, Microsemi, AMD, Western Digital, Huawei, IBM, Nvidia, and many others. Supporters have indicated that RISC-V can provide higher speed and lower power consumption relative to alternatives. The architecture excludes a variety of instructions that the designers believe often slow processors down (e.g., condition code register). Western Digital recently announced that it plans to transition the company to using primarily RISC-V embedded processors over the next few years. Several examples of RISC-V start-ups are discussed at the end of this section.

## Graphics Processing Units

Graphics Processing Units (GPUs) are processor chips that are specifically optimized for the types of computations required for processing and rendering of graphical images. While a conventional microprocessor can theoretically perform these same computations, general processors are not optimized for graphics and cannot perform the calculations as quickly. In addition, graphics computations can significantly bog down a conventional microprocessor, preventing it from adequately performing other tasks. As such, GPUs were developed to offload much of the graphics processing

One of the major differences between GPUs and conventional processors is that GPUs are highly parallelized, since graphics processing requires performing a large number of complex computations that can often be performed simultaneously. GPUs also typically have very high-speed memory interfaces to rapidly move data into and out of memory. To provide a specific example, Nvidia's GeForce GTX 1080 has a GPU engine that incorporates 2,560 CUDA cores (processing units), which provides a very high level of parallel processing. The GPU is based on Nvidia's "Pascal" architecture and includes a 10Gbps, 256-bit wide memory interface and other functions specifically optimized for graphics.

While GPUs were initially used for solely for graphics, there has been growing interest in utilizing these chips for other applications such as artificial intelligence (AI) and autonomous vehicles. Many of the algorithms typically used for machine learning (neural networks) lend themselves to significant parallel processing and are, therefore, well-suited for GPUs. Nvidia has recently introduced solutions (e.g., Volta, Xavier SoC) that are specifically optimized for AI applications, as opposed to only graphics.

When the graphics accelerator market first emerged in the 1990s, there were a variety of graphics chip suppliers and the leadership crown would change hands every couple of years. However, by the early 2000s, the stand-alone GPU market had largely consolidated to Nvidia and ATI (which was acquired by AMD in 2006), and that hasn't changed much since then. Intel provides its own graphical processors that are integrated into the processor die or into the package (referred to as "processor graphics"), or integrated into the chipset ("integrated graphics"). For many users, the Intel graphics are sufficient, but users that want more advanced graphics capabilities can buy PCs or notebooks with an Nvidia GeForce or ATI Radeon graphics solution.

For smartphones and some consumer devices that don't require the bleeding-edge graphics that are needed in gaming PCs, there is often a GPU core integrated into a larger SoC chip. For example, many of Qualcomm's Snapdragon smartphone processors incorporate its Adreno GPU. Companies such as ARM (Mali GPU) and Imagination (PowerVR) license GPU cores for integration into a variety of different types of chips.

## Digital Signal Processors

Digital signal processors (DSPs) are similar to microprocessors, but whereas microprocessors are "general purpose" and designed to execute a broad range of instructions, DSPs are specifically optimized for computationally intensive applications. As one example, many algorithms used for video compression and communications require repeatedly multiplying two numbers together and then adding the result to a third number. This can take many steps with a traditional processor, but DSPs are designed to perform these multiply and accumulate (MAC) operations very rapidly.

DSP functionality is incorporated into many chips and applications. For example, the core processing in the main communications chip (baseband processor/modem) in a mobile phone is typically performed by some type of integrated DSP functionality. Many video processing and wireline communications chips also integrate some type of DSP function.

However, there is also a large market for stand-alone general-purpose DSP chips. Some applications include base stations, multimedia infrastructure, industrial equipment, medical imaging, and communications equipment.

DSPs are often categorized by bit size (e.g., 16 bits versus 32 bits) and whether it processes fixed-point (integers) or floating-point numbers. As with processors, DSPs can also have multiple cores per chip.

Texas Instruments invented the DSP and remains the market leader in stand-alone DSP chips. It offers several families of DSPs including the C66x (high performance) and C55x (low power), as well as solutions that integrate DSPs with ARM processors (OMAP, 66AK2x). As one example, the TI C6678 DSP incorporates 8 DSP cores, each operating at up to 1.25GHz. It can perform 320 GMACs (billion MACs per second) and includes 4MB of L2 cache, a network co-processor, and a variety of interfaces. Analog Devices is another well-known supplier of stand-alone DSP chips. DSP Group licenses DSP cores for integration into SoCs and ASICs.

## Microcontrollers

Microcontrollers are typically lower performance processors used for control applications, and are utilized in a very broad range of products. In addition to the processor itself, microcontrollers often integrate memory and some analog and interface functionality to provide a complete control solution. Microcontrollers are pervasive and can be found in home appliances, automobiles, industrial equipment, and many other types of electronics.

Microcontrollers are often categorized by the number of bits it processes (4-bit, 8-bit, 16-bit, 32-bit). In general, more complex tasks use higher bit processors, and 32-bit microcontrollers have been the strongest growing segment of the market. Some microcontrollers utilize standard processor architectures (e.g., the 8051 was a processor originally developed by Intel but is now commonly used for microcontrollers), while others have proprietary designs (Microchip's PIC architecture). There has been growing interest in ARM-based microcontrollers, as ARM processors are commonly used in many applications. There are also many differences in the type and amount of embedded memory, interfaces, and analog functions integrated.

As one example, NXP's Kinetis K Series line of microcontrollers incorporates an ARM Cortex-M4 processor, between 32KB and 2MB of Flash memory, between 64KB and 2MB of SRAM memory, and a variety of different optional communications interfaces (Ethernet, USB, CAN, etc.).



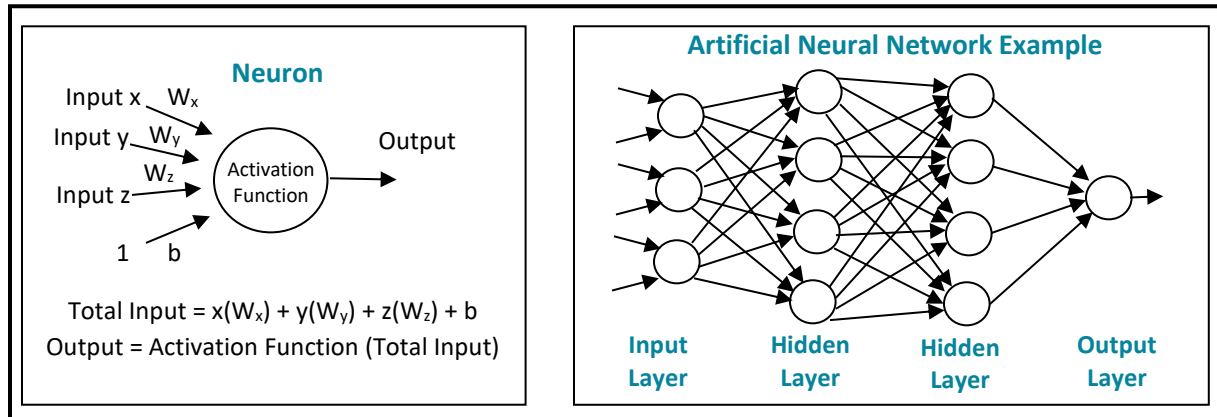
The microcontroller market is approximately \$15 billion. Some of the top microcontroller suppliers include NXP (which acquired Freescale), Renesas, Microchip, Samsung, ST, Infineon, Texas Instruments, and Cypress (through its Spansion acquisition).

## Vision Processors

There is growing interest in analyzing images in video in real-time. For example, automakers are increasingly incorporating ADAS (advanced driver-assistance systems) features in cars with a view towards fully autonomous vehicles. However, analyzing video in real-time to understand what the images in the video are (car, pedestrian, traffic sign, stop light, etc.) is very complex and challenging for conventional processors. As such, several companies have developed processors that are specifically optimized for the types of computations required for vision processing applications. Examples include Mobileye and Movidius (both acquired by Intel) and CogniVue (acquired by Freescale, which was acquired by NXP). As one example, the Movidius Myriad 2 processor integrates twelve 128-bit vector processors (each of which can perform operations on an array of data at one time rather than a single piece of data), a variety of vision processing hardware accelerators, and two RISC processors.

## Artificial Intelligence (AI) Processors

There is great interest in machine learning and neural network technologies (see our previous report on Artificial Intelligence and Machine Learning for a detailed description of machine learning, neural networks, support vector machines, and other related technologies). Technologies such as neural networks typically require enormous processing power and are extremely well-suited for parallel processing. Each “neuron” in a neural network, for example, has to multiply numerous inputs by a variety of weights, add up the results, and then perform some type of mathematical “activation” function on the sum to produce an output. However, there may be a very large number of neurons working in parallel and the output of one layer of neurons becomes input for the next layer, resulting in a huge number of computations, which must be repeated for every new input. This is conceptually illustrated in the following diagram, although a real neural network would have a large number of neurons, many more inputs/weights, and could have many layers.



A Neuron Model and a Neural Network Model. Source: Menalto Advisors

Many AI implementations use high-end server processors (e.g., Intel Xeon), but due to the inherent parallelism of graphics processors, GPUs are now commonly used for AI-type applications and this has been a recent major focus for Nvidia. However, while GPUs generally have advantages for AI applications, traditional GPUs are still not specifically optimized for AI. For example, graphics computations often require much higher levels of precision (e.g., more decimal places) than neural network computations, and graphics calculations differ from the types of calculations needed for many neural networks. As a result, Nvidia has developed GPU-based chips that are specifically optimized for AI, and a number of companies have designed processors from scratch specifically for AI/machine learning applications.

As a few AI chip examples, in October 2017 Intel announced the Intel Nervana Neural Network Processor family (Intel previously acquired Nervana). Intel indicates the Nervana processor has many innovations including massive parallelism and very high bandwidth memory input/output. In addition, it includes software controlled memory management (rather than traditional caching hierarchies), and even uses a new method for representing numbers (Flexpoint) which has advantages for neural networks.




Nvidia introduced its Tesla V100 chip (Volta) in May 2017 which integrates 5,120 CUDA processing cores and was especially optimized for AI-type applications. Its Xavier SoC for autonomous vehicles integrates Volta with vision and deep learning accelerators, video processing, and a custom CPU.

In May 2017, Google introduced its second-generation Tensor Processor Unit (TPU). The Google TPU integrates a matrix multiplier unit that can perform 64 thousand multiplications at one time and a hardwired activation unit to perform neural network activation functions. Thus, it is optimized for the operations used in neural networks (large number of multiplications followed by an activation function). The architecture includes numerous instructions that are specific for neural networks. Some examples of private companies with AI processors are discussed later.

## Quantum Computing Processors

An emerging area of technology is quantum computing. Unlike traditional digital systems in which data is stored and processed as bits (either a “0” or a “1”), quantum computing systems use “qubits” (quantum bits) which are quantum superpositions of “0” and “1” states without necessarily being in either state (unless superposition is removed, in which case it collapses into a traditional digital bit). More specifically, a qubit can be viewed as a linear combination of “0” and “1” states based on the probabilities of each. This can be shown as:  $\alpha |1\rangle + \beta |0\rangle$  where  $\alpha$  and  $\beta$  are probability amplitudes, and  $|\alpha|^2 + |\beta|^2 = 1$  (however,  $\alpha$  and  $\beta$  are complex numbers so the mathematics of even a single qubit, much less multiple qubits, is very complicated). Using superposition and quantum entanglement, it is possible to rapidly compute likely answers to some types of problems (quantum algorithms often don’t return a definite answer but can provide several high probability solutions, which can then be further analyzed by a traditional computer to find the best solution).



There are certain types of problems in which quantum computing can have enormous advantages over traditional computers. One important example is factoring very large integers into two prime numbers, which is the key computation for many public key encryption algorithms (quantum computers can potentially easily crack many commonly used encryption methods, although there are other encryption algorithms that cannot be easily solved by quantum computing).

There are actually many different approaches to quantum computing, so we won't provide details on each here (perhaps a future report). However, the one common theme is that controlling atoms or photons on a quantum level is extremely challenging. To give a sense for what is required, one of the private companies focused on quantum computing is D-Wave Systems. D-Wave's quantum processor is supercooled and operates at just 0.015° above absolute zero (less than -459° F) in an extreme vacuum (pressure is 10 billion times lower than atmospheric pressure) while magnetically shielded (to 50,000 times less than earth's magnetic field). This requires a substantial amount of large precision equipment to support the processor.

Many of the major technology companies (e.g., Google, Microsoft, IBM, Intel, etc.) have quantum computing initiatives. However, each has a significantly different approach. Some initiatives have yet to produce working solutions.

In May 2017, IBM announced a 16-qubit processor (a big leap from its previous 5-qubit processor) and then in November 2017 it announced a 50-qubit processor, and that it is making a 20-bit qubit system available on its cloud. As IBM has indicated it is able to preserve the quantum states for only 90 microseconds, the system is still in development. In October 2017, Intel announced that it delivered a 17-qubit test chip for quantum computing. D-Wave has announced a 2,000-qubit solution, although some have questioned whether its solution truly offers quantum computing.

The achievement of 50 qubits by IBM was considered a milestone, as the general view has been that at least 50 qubits are needed for a quantum computer to surpass a conventional computer for any type of meaningful computation ("quantum supremacy"). While quantum computing is generally expected to surpass conventional computers for solving certain types of complex problems, it remains unclear whether it can be used for more general types of computing applications. Given the level of interest in quantum computing, we anticipate that there will be a variety of interesting quantum processing solution announcements over the next few years.

## Select Processor-Related Start-ups

The following are some examples of processor-related private companies:

- **Ambiq Micro** – Ambiq Micro provides ARM-based microcontrollers. Its Apollo2 family of ultra-low power controllers integrate an ARM Cortex M4 32-bit core, Flash and ROM, power management, serial communications, timing circuits, and a 14-bit data converter for sensors. Its microcontrollers have very low power consumption in both active and sleep modes. Its Subthreshold Power Optimized Technology (SPOT) biases transistors at very low voltages to further improve energy efficiency. It has design wins in Huawei wearables and Fossil smartwatches, and recently announced a low power “always-on” voice control solution. In 2014, it announced a \$15 million C round that included Kleiner Perkins, Austin Ventures, Mercury Fund, and ARM. It is based in Austin, Texas.
- **Cerebras** – Cerebras Systems is a stealth-mode start-up company. Its website does not yet provide any information about its upcoming products, but reports indicate that it is working on an AI processor/software solution, and a significant number of its employees are well-known chip designers. According to an article in Forbes, in three rounds of funding, Cerebras raised \$112 million (with a reported valuation of \$860 million). It is based in Silicon Valley.
- **Codasip** – Codasip develops and licenses RISC-V processor IP. In addition, the company has developed software tools that enable customers to modify and optimize the RISC-V cores (e.g., custom hardware extensions) for particular applications and requirements, which can significantly improve performance. It offers three basic versions of its Berkelium RISC-V processor (1, 3 or 5 stage pipeline) with a variety of other options. In late 2017, it announced a new 64-bit RISC-V processor (Bk5-64). In addition to RISC-V, it offers its own low power Helium family and high-performance Titanium family of processor IP. In November 2017, it announced that Rambus is using Codasip Studio for RISC-V security products. Investors include Credo Ventures and Benson Oak Capital. It is based in the Czech Republic.

- **D-Wave Systems** – D-Wave Systems is focused on developing quantum computing processors and computer systems. Its 2000Q computer includes its quantum processing unit (QPU), which it indicates includes 2,000 superconducting qubits. Its computer uses quantum annealing in which the qubits are initially in superposition and problems are structured such that solutions generate the lowest energy state. When superposition is removed, each qubit collapses to either a “0” or “1” and the resulting state is a likely solution. The QPU is built from a lattice of tiny loops of supercooled metal niobium. It has announced customers such as Google, Volkswagen, Lockheed Martin, Toyota, and NASA. D-Wave has raised over \$200 million. Investors include PSP Investments, Goldman Sachs, Bezos Expeditions, DFJ, In-Q-Tel, BDC Capital, Growthworks, and 180 Degree Capital. It is based in Vancouver, Canada with an office in Silicon Valley.
- **Esperanto** – Esperanto develops energy-efficient processors based on the RISC-V standard. It is developing a multi-core 64-bit RISC-V based solution that is optimized for AI applications, and will reportedly integrate 4,096 small processors (Minions) and 16 larger processors (Maxion). The chip will be produced on a TSMC 7nm process. The company has not yet disclosed specific details about the expected introduction. Reports indicate that it may both sell chips and license its RISC-V processor cores. CEO Dave Ditzel was previously the CEO of Transmeta. Esperanto is based in Silicon Valley.
- **Graphcore** – Graphcore has developed an AI processor. Its IPU (Intelligent Processor Unit) is highly parallel and specifically optimized for machine learning. The IPU processor is graph-focused rather than scalar or vector focused. Graphcore indicates that its solution has over 100X the memory bandwidth of alternatives and can handle more than 14,000 processor threads. In July 2017, it announced a \$30 million Series B round, with investors that include Atomico, Bosch, Dell, Samsung, Amadeus, C4 Ventures, Draper Esprit, Foundation Capital, and Pitango. It subsequently announced a \$50 million capital raise from Sequoia. It is based in Bristol, UK, with an office in Silicon Valley.



- **Greenwaves** – Greenwaves has developed solutions for the IoT processor market. Its GAP8 chip is an ultra-low power IoT application processor that utilizes the RISC-V architecture. It is optimized to enable devices to act on image, motion, and sound data at very low energy levels (which allow years of operation on batteries). The GAP8 incorporates eight cores along with a specialized convolutional neural network accelerator, and uses a shared memory design. It also incorporates many energy-saving features and has only 3uA of standby current. The processor can be used for many applications such as industrial sensor analysis, vision processing, robotics, and a variety of consumer applications. It is based in Grenoble, France.
- **Mythic** – Mythic has developed a unique architecture for AI/neural network processing, which it indicates can provide the power of a desktop GPU in a button size chip. Getting data in and out of memory is a major bottleneck in AI applications. Mythic's approach performs hybrid digital/analog calculations inside flash arrays where processing weights are stored long term. Mythic indicates this provides huge advantages in performance and power consumption. Applications include mobile devices that can benefit from integrated AI processing (e.g., smartphones, smart speakers, drones, and monitoring devices). The company has indicated that it expects to begin shipping chips in 2018. Mythic raised \$15 million in two rounds over a seven-month period. Investors include DFJ, Lux Capital, Data Collective, and AME Cloud Ventures. It is based in Austin, Texas and Silicon Valley.
- **Numascale** – Numascale has developed chip solutions that enable efficient scale-up of servers. There is great interest in scaling up servers to incorporate many processors to handle the growing need for analytics, real-time streaming, and other compute-intensive functions. However, historically it has been difficult to scale up efficiently beyond a few processors as there are typically significant bottlenecks in terms of memory and I/O access. Numascale has developed an innovative node controller (NumaChip) that helps resolve these issues. Its solution integrates a cache coherent memory controller and a high-

performance interconnect fabric. This allows linking of commodity servers together to form a single unified system in which all processors can coherently access and share all memory and input/output. Its technology is available as a standard chip or as an FPGA. It indicates that its solution can reduce total cost of ownership in a data center by 33%. Numascale's investors include Statoil, Investinor, and ProVenture Management. It is based in Oslo, Norway.

- **Quantum Circuits** – Quantum Circuits is focused on developing practical quantum computers based on superconducting devices. It was founded by three experts in quantum devices from Yale University. The team has developed a variety of quantum scientific breakthroughs include a “quantum bus” for entangling qubits with wires, and quantum algorithms and error-correction with a solid-state device. It raised \$18 million in November 2017 from investors that included Canaan, Sequoia, Tribeca Venture Partners, Osage University Partners, and Fitz Gate Ventures. It is based in New Haven, Connecticut.
- **REX Computing** – REX Computing is developing a new, hyper-efficient processor architecture. The company indicates that computing technologies/needs have significantly changed (e.g., energy consumption is more critical and it takes much more energy to move data into or out of memory than to perform extremely complex computations), yet processor design hasn't adjusted to this. As such, it developed a new approach to processors that it indicates results in a 10x to 25x increase in energy efficiency for the same performance level as existing processors/GPUs. Its NEO chip includes 256 cores and a new type of high bandwidth chip-to-chip interconnect. Founders Fund made a \$1.25 million seed investment. It is based in California.
- **Rigetti** – Rigetti is focused on quantum computing processors and solutions. The company announced it developed a 19-qubit quantum processor (19Q) and also has a suite of software to support it. Its Forest solution is a full-stack programming and execution environment for quantum/classical computing and it provides an API for quantum

computing in the cloud. Rigetti has developed solutions using its quantum technology for applications such as machine learning (clustering). It raised \$64 million in Series A and B rounds. Investors include Vy Capital, Andreessen Horowitz, and Y Combinator. It was founded in 2013 and is based in Berkeley, California and in the Silicon Valley area.

- **SiFive** – SiFive is focused on embedded processor IP for the RISC-V architecture. The founding technical team includes several of the key engineers that initially created the RISC-V architecture at Berkeley. SiFive indicates that its E31 Coreplex is the world's most deployed RISC-V core and that it has many advantages in terms of cost, performance, and pricing (no royalties) over the ARM Cortex processors. The company also offers the E51, a higher-end 64-bit RISC-V processor core, and several SoC platforms (Freedom) built around its processor cores. In May 2017, it announced an \$8.5 million funding. Investors include Sutter Hill Ventures, Spark Capital and Osage University. It is based in San Francisco.
- **Silicon Mobility** – Silicon Mobility develops chips and software that control electric motors, batteries, and energy management systems of electric vehicles (EVs) and hybrids. Its OLEA field programmable control units (FPCUs) are optimized for EV control applications. OLEA incorporates its AMEC flexible interface for deterministic control of actuators and sensors and its SiLant functional safety solution. It utilizes multiple ARM processors along with a proprietary Flexible Logic Unit. It also provides a variety of libraries (LIB) and tools (Composer). By optimizing inverter, AC/DC, and DC/DC control, Silicon Mobility indicates that it can dramatically improve the efficiency and range of EVs, while reducing size, weight, and cost. It is working with several major automotive OEMs and Tier-1 suppliers, and is part of GlobalFoundries' new automotive program. Investors include Cipio and Capital-E. It is based in France.

- ThinCi** – ThinCi is developing an advanced deep learning processor that is specifically optimized for vision processing applications. The company has not yet released architecture details but indicates that its highly parallel solution can process an entire image at once (similar to how a brain works), rather than analyzing image portions sequentially. It notes that this results in much higher performance, less memory access, and lower power consumption. The company anticipates applications will include self-driving cars, drones, smart homes, smart cities, social media analytics, and augmented reality devices. ThinCi's investors include automotive supplier Denso, as well as several technology/VC executives. The company is based outside of Sacramento, California with a major office in India.
- ThinkForce** – ThinkForce Electronic Technology is developing an AI acceleration engine, and is focused on combining silicon processing chips that are optimized for AI computations with advanced AI algorithms. In December 2017, reports indicated that the company raised 450 million Chinese yuan (about \$68 million) in a Series A funding round. Investors include Sequoia, YITU Tech, Yunfeng Fund, and Hillhouse Capital. It is based in Shanghai, China.
- ThinkSilicon** – ThinkSilicon provides GPU IP cores for embedded applications. There is growing interest in GPU processing (both for graphics and for AI/vision processing applications). ThinkSilicon has developed several "Nema" GPU cores including the "p" which is only 0.07mm<sup>2</sup> with 0.06mW of power (optimized for SoCs), the "t" (which is optimized for IoT and wearable applications), and the "s" (multi-core solution for high performance needs, including neural network applications). It also offers a variety of advanced display controller solutions and software APIs. The company has announced wins with Lattice (machine learning and sensor security) and Sequans (LTE for IoT SoC), and recently announced a partnership with Synopsys (ultra-low power wearable platform). It is based in Greece.

- **UltraSoC** – As processors and processor-based SoCs become more complex, it is more challenging to determine how well a chip is operating and what is really happening inside. UltraSoC develops IP modules that can be integrated into processors and SoCs to provide flexible on-chip monitoring and analytics. The IP is non-intrusive and works at wire speed. UltraSoC provides over 30 different modules (bus performance monitor, status monitor, debug DMA, processor modules, etc.). This can be used both to debug chips and to improve future designs. It has been used across many processors (ARM, MIPS, etc.) and UltraSoC has had a recent focus on RISC-V. Microsemi recently announced it was using UltraSoC for its RISC-V developments. Investors include Octopus Ventures, Oxford Capital, Enso Ventures, and Atlante Tech. It is based in Cambridge, UK.
- **Videantis** – Videantis licenses vision processing and video-related processor IP cores that it indicates can provide 1000x lower power consumption and 100x better performance per dollar relative to alternative processors/GPUs. The recently announced v-MP6000UDX processor adds instructions optimized for running convolutional neural networks (CNNs), increases the MAC throughput per core eightfold, and extends the number of cores from 8 to up to 256. The company also introduced a variety of software tools to support deep learning and CNNs (e.g., v-CNNDesigner). It is a follow-on to the company's successful v-MP4000HDX processor. Its processors and tools support applications such as computer vision, video compression/decompression, and imaging algorithms. Videantis supports many video standards (H.265, H.264, MPEG-4, etc.) and a variety of image and vision processing functions (object detection, feature detection, image resizing, etc.). The automotive market has been a recent major focus area. Bosch is one of its announced customers. In September 2017, it announced growth financing from eCapital. The company is based in Hanover, Germany.

- **Vorago** – Vorago develops high reliability semiconductor devices (microcontrollers and other types of devices) that can operate under a variety of extreme conditions such as high temperature (200°C+) and high radiation (>300K Rad). Unlike many high-reliability solutions that rely on exotic materials or specialized packaging, Vorago’s “Hardsil” technology can harden any CMOS design, requires only modest modifications to commercial CMOS process flows, and can be implemented quickly with minimal incremental mask and implant steps. Leveraging this technology, Vorago has developed its own hardened chips including ARM-based microcontrollers and SRAMs. It also designs custom chips, and licenses its IP to other chip companies. Vorago initially focused on space/aerospace applications, but has recently seen demand from many other sectors (automotive, data centers, etc.) as there is growing interest in chips that can operate well at elevated temperatures. Investors include New Science Ventures. It is based in Austin, Texas.



## Memory

---

### Intro to Semiconductor Memories

Semiconductor memory chips store digital data. There are other methods of storing digital information (e.g., hard disk drives in which data is stored on rotating metal disks coated with magnetic material), but semiconductor memory is substantially faster, smaller, and has a variety of other advantages. Chip memories are often divided into two categories.

- **Volatile** – With volatile memory, the chip loses all of its stored information when power is removed. The two most common types of volatile memory are DRAM (Dynamic Random Access Memory) and SRAM (Static Random Access Memory).
- **Non-Volatile** – Non-volatile memory retains data even after power is removed and the system is shut down. When the system is powered up, the stored data is readily available to access. The most popular type of non-volatile semiconductor memory is Flash.

Semiconductor memory can be viewed as a matrix of rows and columns with a storage element (which can store a “0” or a “1”) at the intersection of each row and column. Each type of memory has a different kind of storage element. For example, DRAM uses a transistor and capacitor, SRAM typically uses six transistors, and Flash uses a “floating gate” transistor that stores charge. As manufacturing technologies improve and process nodes shrink, chipmakers are able to squeeze more memory into a given size chip. This allows the same amount of memory in a much smaller chip or, more commonly, more memory while keeping chip size and cost about the same.

### DRAM

The storage element in DRAM is a transistor and a capacitor. When the transistor turns on, the capacitor charges up and a charged capacitor can represent a “1” while an uncharged capacitor could represent a “0”; enabling a bit to be stored. However, capacitors store energy for only a short period of time and must be refreshed often. During these “refresh” cycles, data cannot be stored or read from the DRAM, which limits DRAM speed. When DRAM loses power, the refreshing stops and all data is lost.

Major markets for DRAM include PCs, servers, and smartphones, although most complex electronics incorporate DRAM. Traditionally, PCs and servers had two major types of storage – disk drives and DRAM. Disk drives are non-volatile and store data indefinitely as magnetic signals on metallic platters. PC software and data (operating systems, applications, files, etc.) are stored on the disk drive. However, disk drives are much too slow to keep up with a microprocessor (disk drives provide data in milliseconds versus nanoseconds for DRAM and processors). When a traditional PC boots up, key parts of the operating system are copied into the PC's much faster DRAM. DRAM can then provide the software instructions or data to the processor in a timely manner, and will periodically move data to the disk drive so it isn't lost when the system is powered down. (As noted later in the chapter, disk drives are increasingly being displaced by Flash storage.)


Until the mid-1990s, most DRAM was asynchronous, meaning that it was not directly tied to the system clock. A processor would request data from DRAM, and at some point DRAM would place the data on the memory bus for the processor to take, but there wasn't tight coordination between the processor, DRAM, and the rest of the system.

To improve performance, the industry shifted to synchronous DRAM (SDRAM) in which DRAM is synchronized to the system clock (the clock repeatedly emits pulses at precise intervals and the SDRAM can provide data each time it receives a clock pulse). In this way, the microprocessor (which also is tied to the system clock) knows exactly when the data will be available and can obtain and store data faster and more efficiently.

Next, the industry shifted to double data rate (DDR) DRAM. Instead of providing data only when it receives a clock pulse, DDR DRAM can provide data both at the beginning of each clock pulse (rising edge) and at the end of each clock pulse (trailing edge), effectively doubling the speed that data can be stored or provided. Since then the industry has shifted to more advanced versions including DDR2, DDR3, and most recently DDR4.

Each generation further improves performance, but is not backwards compatible. DDR4 can operate at a clock frequency of between 800 and 2133MHz, which is twice that of DDR3. Note that DDR that operates at 800MHz, for example, is typically called DDR-1600.

Currently, 8Gb DDR4 is fairly commonplace and 16Gb DDR4 DRAM is shipping in many systems, with sampling of 32Gb DDR4 DRAM. DRAM chips are often combined on DIMM modules (small printed circuit boards with several DRAM chips, as shown on the right). A typical notebook PC currently might have 8 gigabytes of DRAM memory, which is 64 gigabits (1 byte = 8 bits). That might include eight 8Gb DRAM chips on a DIMM module. The iPhone 8 has 2GB of DRAM while the iPhone 8 Plus has 3GB.



The DRAM market has significantly consolidated during the past couple of decades with three suppliers (Samsung, Micron, Hynix) now accounting for over 90% of the market. It is also one of the most volatile chip segments. For example, DRAM sales declined in both 2015 and 2016 and then increased by more than 70% in 2017 (according to some estimates, as 2017 numbers haven't been finalized). It is also the largest memory segment (likely over \$70 billion in 2017).

## SRAM

SRAM is another type of volatile semiconductor memory. In contrast to DRAM in which the memory element consists of one transistor and a capacitor, most SRAM designs use six transistors (although there have been four transistor designs). The advantage is that by eliminating the capacitor, SRAM doesn't require refresh cycles and can usually operate much faster than DRAM. The disadvantage is that six transistors take up much more area than one transistor and a capacitor. Since the cost of a semiconductor chip is usually directly related to chip size, SRAM is much more expensive per bit of memory than DRAM and is generally used only in applications in which the enhanced speed is needed.

The peak for the SRAM market occurred during the 1995 to 2000 timeframe (about \$6.5 billion in 2000), but the market has sharply declined since then to less than \$500 million. A major reason for this is that for a period of time stand-alone SRAM had been used in PCs. That is, computers typically had DRAM but also included faster SRAM memory. PCs were designed to try to keep the data it used most often in the SRAM to improve overall speed.

However, as previously discussed, processors now often integrate several levels of cache (SRAM) memory on-chip. This, combined with improved DRAM access speeds through DDR technologies, eliminated the need for separate SRAM chips in PCs, which resulted in a sharp decline in the market. SRAM is often integrated into large SoCs, reducing the need for stand-alone SRAM chips. The largest market for stand-alone SRAM chips tends to be high-speed networking/telecom equipment in which packets have to be processed very rapidly. As with DRAM, there are double data rate versions of SRAM, as well as “quad” rate SRAMs that further enhance operation. For example, Cypress’ QDR-IV SRAM includes two separate ports (each double data rate), each of which can read or write data independently.

With the decline of the stand-alone SRAM market over the past decade, many of the traditional suppliers (e.g., Samsung, Sony, Micron, etc.) exited or deemphasized the market. The remaining suppliers include Cypress, Renesas, ISSI, GSI Technologies (which is primarily focused on SRAM for networking), and IDT (which has deemphasized SRAM).

## Basic Non-Volatile Memories

By far, the largest segment of the non-volatile memory market is Flash memory, which is discussed in more detail below. However, there are a variety of older read only memory (ROM) technologies. These non-volatile ROM memories are generally viewed as commodity-type chips and are typically used if only a small amount of memory is required or if the memory contents are unlikely to change over time.


These basic non-volatile memories include: ROM (bits are determined during manufacturing and can’t be modified), programmable ROM or PROM (memory can be programmed in after manufacturing using a special tool, but can’t be changed after that), Erasable PROMs or EPROMs (memory can be programmed in, erased, and reprogrammed using a special tool, but it is a complex process), and Electrically Erasable PROMs or EEPROMs (enables memory contents to be stored, erased, and reprogrammed electrically while the chip is installed in a system, rather than requiring a special tool).

EEPROMs became very popular for a period of time. However, the core architecture is difficult to scale to smaller process nodes and it can be relatively slow (only small sections can be erased at one time). As such, it has generally been displaced by Flash (discussed next), which is substantially less expensive per bit, except for applications when only a small amount of memory is needed. In total, these older ROM memory technologies generate only a few hundred million dollars of sales per year.

## Flash Memory

Flash has become the dominant type of non-volatile semiconductor memory. It is the primary non-volatile memory used in smartphones and tablets, and many notebooks PCs now use Flash memory in lieu of traditional hard disk drives. Flash was originally developed by Toshiba in the early 1980s, although it took a number of years to optimize it for widespread commercial deployments. Technically, Flash could be considered a type of EEPROM in that it is electrically erasable and it is programmable while in-circuit. However, it uses a different storage architecture than historical EEPROMs and is virtually always considered a separate category of memory.

The basic memory cell in Flash is similar to a regular MOSFET transistor, except there is an additional “floating gate” below the regular gate. The floating gate has an insulating oxide layer around it such that electrons that enter the floating gate are generally trapped unless some type of external field is applied. By temporarily applying appropriate voltages/electric fields, it is possible to cause a sharp increase or decrease in the number of electrons trapped in the floating gate, but otherwise the number of trapped electrons remains stable. This can be used for storage (e.g., a low level of stored charge can represent a “1” while a high level can represent a “0”). There are different methods for injecting or removing electrons from the floating gate (e.g., one common approach is a quantum process called Fowler–Nordheim tunneling).



Source: Menalto Advisors

As previously discussed in the transistor overview section, when an appropriate gate voltage is applied, MOSFETs generate an electric field which enables the transistor to turn on and current to flow. However, electrons trapped in the floating gate diminish this electric field. As such, as the number of trapped electrons increases, turning the transistor on becomes more difficult and the threshold voltage required to turn the transistor on increases. By checking whether the transistor turns on or not at mid-range voltages, it is possible to determine whether a significant amount of charge is stored or not. This, in turn, can determine whether a “0” or “1” is stored.

Flash is often divided into two types: NAND and NOR. Both utilize the floating gate memory technology previously described, but the connections among the various memory cells differ and there are some other technical differences. NAND memory is optimized for “disk drive replacement” applications and is



significantly denser and less expensive per bit than NOR. It also has faster write times. NAND is used in Flash memory sticks and in most products that use Flash in lieu of disk drives. NOR is optimized for “code storage” applications. It generally enables faster read times than NAND (write times are longer, but software code is rarely changed). An electronic product can have both NAND and NOR Flash, although the NAND market is significantly larger. Some Flash chips incorporate memory controllers while other are “raw” and require a separate memory controller chip. There are also different types of outputs used (parallel or serial).

In order to improve the density of NAND Flash and further drive down costs, there have been a variety of additional technical enhancements. One of these is multi-bit per cell technology. As previously noted, Flash determines whether a “1” or “0” is stored based on the amount of charge stored in the floating gate. For traditional “single level cell (SLC)” Flash, only two levels of stored charge are evaluated (high and low) which translates into one bit (“0” or “1”). However, if the level of charge can be precisely controlled, each memory cell can potentially store more than one bit. For example, if four different levels of charge can be stored (e.g., very low, low, medium, high)



then each of these levels could represent two bits (“00”, “01”, “10”, “11”). Several of the major Flash companies offer multi-level cell (MLC) NAND Flash that provide two bits per cell, or even triple level cells (TLC) that provide three bits per cell. There have already been announcements about four-bit cells. This can greatly improve density. Flash OEMs often use error correction code technology to help ensure these memories provide accurate data as there is greater likelihood of an error with multi-level cells.

There has also been strong interest in 3D NAND. With a 3D structure, memory cells can be stacked vertically (rather than only across the surface of the chip in two dimensions), which can greatly increase density and potentially eventually drive down costs. Most of the major Flash memory companies have introduced high density 3D NAND solutions, although the technologies differ from company to company. The yields and costs for 3D Flash have not yet been fully optimized, so it accounts for only a portion of the market, but some expect it will eventually become pervasive.

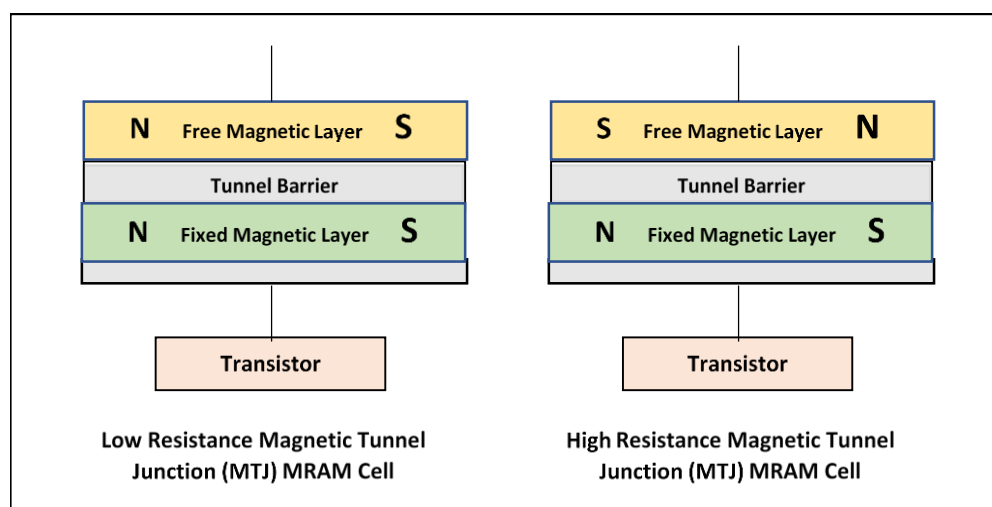
The NAND market had generally been running in the \$30 billion or so range, but increased sharply in 2017. NOR represents a small fraction of that, but is still a nearly \$2 billion market. The Flash memory market is largely dominated by a few companies. The major suppliers of NAND Flash include Samsung, Toshiba (which is divesting its business to a Bain-led consortium), Western Digital (which acquired SanDisk), Micron, SK Hynix, and Intel. The major suppliers in the much smaller NOR Flash market include Cypress (which acquired Spansion), Micron, Macronix, Winbond, and GigaDevice. There are also IP companies that develop and license Flash memory cores for embedding in SoCs and ASICs.

## MRAM

Magnetic RAM (MRAM) is an emerging type of semiconductor memory that uses magnetic storage elements. MRAM is non-volatile like Flash, but is generally much faster than Flash memory (often advertised to have speeds approaching SRAM). It also has other attractive properties such as very high endurance (more reliable than Flash memory after a large number of reads/writes), low power consumption, and the ability to work at high levels of radiation. However, the density of MRAM has been significantly less than Flash and substantially more expensive per bit.

For example, Everspin offers MRAM chips with up to 256Mb of memory, whereas Flash memory chips are available with thousands of times more memory. As a result, MRAM has historically been used only in niche applications, in which relatively small amounts of memory are required.

An MRAM cell typically consists of a regular transistor and a magnetic tunnel junction (MTJ) cell which magnetically stores a bit of data. The MJT has two magnetic layers, separated by a thin insulator layer. One of the magnetic layers has a fixed polarity (one side is always North and one is always South) while the other is a free layer in which the polarity can be electrically flipped. When the two layers are magnetically aligned (both layers have North in the same direction) the resistance is low, otherwise it is high. By changing the polarity of the free layer, a “0” or “1” can be stored, and whether a “0” or “1” is stored can be determined by measuring the resistance. Since the magnetic alignment doesn’t change when power is lost, it is non-volatile.



Source: Menalto Advisors

There have been a variety of different methods for “writing” to MRAM (flipping the free layer). Toggle MRAM uses two write lines around the cell. When current passes through both write lines, an induced magnetic field is created which can be used to “flip” the polarization of the free magnetic layer. However, Toggle MRAM requires a significant amount of current and has not been scalable to high density (generally only 16Mb chips or less).

As such, MRAM developers have focused on spin transfer torque (STT) MRAM. Electric current is generally unpolarized with equal numbers of electrons with “up” spin and “down” spin. However, a magnetic field can alter this to create a “spin-polarized” current. When this type of current passes through a magnetic layer, it can flip the orientation of the layer (e.g., flipping the direction of North and South). In an MRAM, this can be used to flip the free magnetic layer (turning a “0” to a “1” or vice versa). STT-MRAM can scale to higher densities than toggle MRAM, although it is still currently substantially less dense than Flash. While several companies are developing STT-MRAM, implementation details are very different for each.

MRAM was first commercially shipped in 2006 (by Freescale) but is still a niche market. Everspin has generally been viewed as the leader in the market but its revenue during the past year was less than \$40 million, so this is still a small market. However, MRAM density has been increasing and GlobalFoundries recently announced a partnership with Everspin for providing embedded MRAM. Samsung, IBM, Toshiba, and others (including some start-ups noted later) have announced MRAM technologies. It remains unclear if density and yields will ever increase enough to eventually compete with Flash memory in high volume applications.

## XPoint Memory

Micron and Intel have jointly developed a new memory technology called XPoint. Intel markets the memory under the Optane brand while Micron uses the Quantix brand, although both indicate that their marketing strategies for the memory are different. XPoint is a non-volatile memory.

XPoint reportedly has dramatically less latency than Flash memory (nanoseconds versus milliseconds), although it is slower than conventional DRAM. It is also currently less expensive than DRAM (but more expensive than Flash). Intel has indicated that it also has better endurance than Flash and is much denser than DRAM. Unlike NAND Flash (in which pages of data are read or written to at one time), each XPoint cell can be individually addressed, which can provide advantages. Theoretically, XPoint could compete against both DRAM and Flash, although it appears the initial focus is on DRAM replacement for certain applications in the data center.

Intel has indicated that the core XPoint memory cell is “transistor-less” and includes a stacked “memory cell” and “selector.” A grid of perpendicular metal layers, one above and one below each memory cell, is used to select individual memory cells (when the metal layer on both the top and on the bottom of a cell is turned on, that memory cell is selected). These memory “grids” can be vertically stacked to create high density memory. There has been some conflicting information on the core technology that is used in the memory cell, but it is generally believed to use some type of phase change memory (PCM) that provides changes in resistance for differentiating between “0”s and “1”s (phase change memory and resistive RAM memory are discussed in more detail later in this section).

The XPoint memory was initially announced in mid-2015. A few initial XPoint based products were released in 2017 (e.g., Intel announced an Optane solid state drive for the data center), but the general expectations are that there will be higher volume implementations beginning in 2018.

## Emerging Memory Technologies

There are a variety of other non-volatile memory technologies that have been developed with the potential to offer significant advantages over Flash (e.g., lower cost or higher reliability or higher speed, etc.). However, it remains to be seen whether any of these will become commercially successful and challenge Flash. We won’t try to cover all of the emerging memory technologies, but a few are summarized below and several examples of relevant start-ups are highlighted later in the section.

- **Phase Change Memory (PCM)** – The properties of certain materials can considerably vary depending on their state. For example, a material may have a very high resistance in its amorphous state but low resistance when in a crystalline state. By applying heat or laser pulses in a certain manner, it is possible to transform a material from one state to the other and back, but otherwise it will remain in its given state. This can be used to store data (e.g., amorphous/high resistance might represent a “0” while crystalline/low resistance might represent a “1”). Since resistance is easily measured, it is easy to determine if a “0” or “1” is stored. This is essentially how phase change memory

operates, although there are many variations. One of the first materials used for PCM was chalcogenide glass, but since then a number of other materials and techniques for changing states have been developed. One potentially interesting aspect is some materials have numerous discernable states, enabling each state to represent multiple bits rather than just one (e.g., if a material has 8 states, each state can represent three bits – from “000” through “111”).

- **Resistive RAM** – Resistive RAM (ReRAM) is another type of non-volatile memory. The memory element typically consists of two metal electrodes sandwiching a switching layer that contains a dielectric material which is normally not conductive (very high resistance). However, the switching layer material also contains ions that can be manipulated (using heat or an electric field) in such a way that the switching layer becomes much more conductive (low resistance). By applying appropriate voltages, the switching layer can switch from high resistance (which could represent a “1”) to low resistance (“0”), and back. There are many variations using a variety of materials. Several companies have developed ReRAM technologies (e.g., Crossbar, Adesto, Panasonic, HP), although it is still early stage.
- **Nanotube-Based RAM (NRAM)** – Carbon nanotubes are one of the strongest materials and have very high electrical and thermal conductivity. There has been interest in utilizing this technology for electrical applications including non-volatile memory. Two carbon nanotubes conduct well (low resistance) when touching but don’t conduct well (high resistance) when separated. Because nanotubes are very stiff when they are not in contact, they will remain that way unless acted upon by an outside force. Similarly, when they are in contact, Van Der Waals force causes them to remain stuck together. By applying appropriate voltages, the nanotubes can transition between one state or another (touching/low resistance and non-touching/high resistance), but otherwise remain unchanged. This can be used to store bits of data (low resistance = “1”, high resistance = “0”). Nantero (noted below) has been focused on NRAM solutions.

- **3D Technologies** – As previously noted, one of the ways chip companies are attempting to increase density is through stacking memory elements in the vertical direction. This has the potential to greatly increase density and reduce costs, once yields are high. As chips historically implemented transistors only on the wafer surface, 3D requires a number of manufacturing and design changes, but there have been several positive recent developments. There are also technologies for tightly stacking memory chips together in a package to improve density (e.g., Micron’s Hybrid Memory Cube).

### Select Memory Chip Start-ups

The following are some examples of private chip companies focused on memory technologies. Some of these are developing their own chips, while others are IP companies focused on licensing.

- **Avalanche Technology** – Avalanche has developed STT-MRAM technology. Its STT-MRAM memory is based on its proprietary perpendicular magnetic tunnel junction (pMTJ) cells which can be manufactured on standard CMOS 300mm wafer processes. Avalanche’s first STT-MRAM memory products are 32Mb memory chips built on a 55nm-node foundry process. It also offers embedded MRAM. In October 2016, it announced a manufacturing agreement with Sony Semiconductor to begin production of its MRAM in 2017. In February 2016, it announced a \$23 million funding round with investors Thomvest Ventures, Vulcan Capital, Rogers Venture Partners, and VTB Capital. The company is based in the Silicon Valley area.
- **Crossbar** – Crossbar has developed a type of Resistive RAM memory technology. As with most ReRAM, its memory cell consists of two metal electrodes around a switching material. For Crossbar, the switching material is silicon-based, and a metallic nano-filament can be formed when appropriate voltages are applied across the two electrodes. Crossbar indicates that its memory cells can be stacked in 3D arrays, enabling very dense memory arrays, with one transistor driving many memory cells. The technology can also be integrated on



top of CMOS logic to provide high density embedded memory. Crossbar indicates that its memory has numerous advantages over Flash including 100X faster read time and 1000X faster write time, as well as much higher density. Crossbar announced a partnership with SMIC in 2016. In late 2015, it announced a \$35 million Series D (bringing total investment to \$85 million). Investors include Tyche Partners, Artiman Ventures, CBC-Capital, Cheerful Link Ventures, Correlation Ventures, Kleiner Perkins, Korea Investment Partners, Northern Light Venture Capital, Oriza Holdings, SAIF Partners, Tao Invest, and the University of Michigan. It is based in Silicon Valley.

- **eVaderis** – eVaderis provides embedded memory IP for SoCs and other types of chips, with a focus on disruptive new types of non-volatile memory including STT-MRAM and resistive RAM (ReRAM). Its IP includes specialized memory, memory compilers, logic libraries, and optimized subsystem solutions. Target markets include low power applications such as: IoT, mobile, and wearables. It recently demonstrated its MRAM IP in a low power microcontroller produced on a GlobalFoundries' 40nm CMOS process. The company was founded in 2014 and is located in the Grenoble, France area.
- **Nantero** – Nantero has developed carbon nanotube based non-volatile memory technology (NRAM). With NRAM memory, carbon nanotubes remain either conducting or not contacting indefinitely, unless acted upon by an outside voltage. When nanotubes are in contact, resistance is low ("1") and when they aren't resistance is high ("0"). This can be used to store data. By applying appropriate voltages, the nanotubes can be switched. Nantero indicates that its technology is CMOS compatible, can scale below 5nm, is hundreds of times faster than NAND Flash, can retain data for over 1,000 years, and has much lower power consumption than NAND. It announced a \$30 million Series E round in mid-2015 and a \$21 million raise in late 2016. In August 2016, Nantero announced that Fujitsu Semiconductor and Mie Fujitsu (foundry) licensed its NRAM technology and the companies were working towards releasing a product, and that Mie Fujitsu would

offer NRAM to its foundry customers. Investors include Charles River Ventures, DFJ, Globespan, Harris & Harris, Schlumberger, and several undisclosed strategic partners. It is based in Woburn, MA.

- **Spin Transfer Technologies (STT)** – STT has developed Orthogonal Spin Transfer (OST) MRAM technology that it believes has substantial advantages (lower power, faster performance, higher reliability) over MRAM alternatives and will enable MRAM to eventually displace Flash (and DRAM) in both stand-alone and embedded markets. Its OST-MRAM includes a variety of patented technologies including a CMOS circuit engine that provides much higher endurance by optimizing the operating voltage, a spin polarizer that enables higher speeds and stability, and a high-density structure that scales to small sizes. STT announced that it sampled initial chips in January 2017. In late 2017, it announced a partnership with Tokyo Electron and a \$22.8 million bridge facility. It was formed by Allied Minds and New York University. It is based in the Silicon Valley area.
- **SureCore** – SoCs increasingly incorporate large amounts of embedded SRAM, which in many cases can account for the majority of power consumption in a chip. SureCore examined the key parameters that account for SRAM power consumption (bitline voltage swings, number of active gates, etc.) and, taking a “clean sheet” approach, optimized the design of SRAM to minimize power consumption. It provides two families of IP: PowerMiser (low power SRAM IP) and EverOn (single port ultra-low voltage SRAM IP that provides up to 80% savings in dynamic power consumption and an up to 75% reduction in static power). In September 2017, SureCore announced that embedded FPGA supplier Menta is using SureCore’s SRAM IP on TSMC’s 28nm process. More recently, it announced that it was selected for GlobalFoundries’ FDXcelerator program. It is based in Sheffield, United Kingdom.

## FPGAs

---

### Intro to FPGAs

In designing a new product, an OEM will typically try to use off-the-shelf standard semiconductor chips. However, for some portions of the design there may not be any chips on the market that meet its needs. One option is to design a custom chip (an Application Specific IC or ASIC). This has some advantages in that it can provide the exact functionality the OEM is looking for. However, developing and producing a custom chip from scratch is very time consuming (can take many months), is costly (mask costs for leading-edge processes can run into the millions), and has risks (if the chip doesn't work as expected or the requirements change, the chip has to be redesigned resulting in additional costs and delays). An alternative is to use an FPGA (Field Programmable Gate Array). An FPGA enables an engineer to rapidly implement a custom design, without having a new chip produced.

The term Field Programmable Gate Array is based off of a category of legacy chips known as gate arrays. As discussed earlier, the basic elements for most digital circuits are logic gates. These perform simple analysis (e.g., the output of an AND gate is "1" only if both inputs are "1"). By combining thousands or millions of these gates, complex functionality can be implemented. The idea behind gate arrays was to create chips that had a large number of gates, but didn't have the final metal layer which determined how the various gates were connected. When a customer wanted to implement a custom design, the gate array supplier sent the chip through manufacturing again to add the final metal layer (connecting all the gates in the appropriate manner to implement the desired logic). This provide some advantages in terms of cost and time-to-market, but was still not ideal.

FPGAs further improve upon this idea by enabling a custom design to be implemented without any manufacturing. That is, an engineer with a design can buy an off-the-shelf FPGA chip and (using software tools provided by the FPGA company) configure the FPGA to quickly implement the design rather than having to manufacture a custom chip.

An FPGA can be thought of as a gate array that can be programmed, although the architecture of FPGAs typically uses “look up tables” (LUTs) to perform the logic functions rather than connecting a “sea of gates.” With most FPGAs, the desired configuration data is stored in memory (SRAM). When system power is turned on, the configuration data from memory is loaded into the FPGA, which is then configured to implement the desired functionality. Because of this, SRAM-based FPGAs are reprogrammable (a design can be modified simply by changing the data stored in SRAM).

Although FPGAs originally had simple logic gates as basic building blocks, they have become increasingly complex and now often also incorporate a variety of embedded memory, high speed interfaces, and other functions. In some cases, they include embedded processors. As an example, Xilinx’s Virtex UltraScale+ FPGA family provides up to 3.6 million logic cells, but also incorporates up to 500Mb of on-chip memory, up to 8GB of in-package memory, and a variety of interfaces (100G Ethernet, PCI Express, etc.).

FPGAs are a subset of programmable logic devices (PLDs), although it now accounts for the vast majority of the programmable logic market. Another type of programmable logic chip is a complex programmable logic device (CPLD) which has similarities to FPGAs but uses a different type of architecture and isn’t as scalable. CPLDs are still sold, but are generally viewed as a much smaller niche market.

FPGAs were originally often used as “glue logic” in systems and implemented the small amount of custom logic needed to help ensure the other standard chips in a system worked together. FPGAs are also often used for prototyping, in that engineers may initially develop an FPGA version of a chip for testing purposes and then have the design manufactured as a chip once the design has been finalized. However, as FPGAs have become increasingly powerful and time-to-market is increasingly important, they are now often the critical components in many systems. FPGAs are commonly used in telecommunication, networking, and wireless infrastructure equipment. An emerging segment for FPGAs is machine-learning acceleration, since FPGAs are good at performing many computations in parallel (whereas processors tend to execute instructions serially), which is generally important for neural networks and AI algorithms.

There are trade-offs between using an FPGA and having a custom chip produced. FPGAs can provide faster time-to-market and don't require significant upfront costs. There is also less risk since designs can be modified and updated without any additional cost. However, the downside is that FPGAs are not completely optimized for any particular design (typically only a subset of the logic and functionality on a chip is used) and so a custom chip can be much smaller and have better performance. There also might be a variety of functions not available on an FPGA (e.g., analog functions) that could be incorporated into a custom chip. A major factor is often volume. While there is a high upfront cost associated with custom ASICs, for high volumes it is often cost effective. However, for lower volumes, FPGAs often make more economic sense.

## FPGA Suppliers

There are only a handful of FPGA suppliers. Historically, the two leading FPGA companies have been Xilinx and Altera (which was acquired by Intel). Combined, they typically account for more than 85% of the market. The number three player has been Lattice Semiconductor, which historically lagged behind Xilinx and Altera but has had a more recent focus on low-power FPGAs for consumer/mobile applications. Microsemi (through its acquisition of Actel a number of years ago) provides programmable logic chips, with a focus on military and aerospace applications. QuickLogic originally focused on stand-alone one-time programmable anti-fuse FPGAs, but has largely shifted its focus to embedded FPGAs and other products.

There is growing interest in embedded FPGA technology, which enables designers to incorporate small amounts of FPGA functionality into larger chips. For example, a microprocessor might include a small amount of FPGA technology to help implement and accelerate certain functions that can be performed much faster using FPGAs. AI/machine learning, for example, is an area in which highly parallel computations are required and can benefit from embedded FPGA technology. It was generally believed that one of the reasons Intel acquired Altera was to have access to FPGA technology for enhancing server processing. A number of start-ups are focused on this aspect of FPGAs, as noted below.

## Select FPGA Start-Ups

The following are examples of programmable logic private companies:

- **Achronix** – Achronix Semiconductor provides high-performance FPGAs and embedded FPGA (eFPGA) IP solutions. The company initially focused on stand-alone FPGA chips. In 2013, it introduced its Speedster22i FPGA family, which received attention as it was fabricated by Intel on an advanced 22nm process. The Speedster22i FPGAs offer up to a million LUTs, 86Mb of embedded RAM, and include a variety of high speed communications interface IP. In late 2016, the company announced the availability of embedded FPGA technology (Speedcore), which appears to have been a significant catalyst. In June 2017, Achronix announced that its revenue would exceed \$100 million in 2017 (up over 700% Y/Y) and that it would introduce its next-generation FPGA family in 2018. Investors include New Science Ventures, GKFF, and Easton Capital Group. It is based in Silicon Valley, with major offices in India and China.
- **Flex Logic** – Flex Logic is specifically focused on embedded FPGA solutions. The company believes its FPGA technology provides several compelling advantages. Traditional FPGAs can require 10+ metal layers for interconnections. Flex Logic uses a patented architecture which requires about half the area and only 5-6 metal routing layers. It also incorporates reconfigurable I/O, which takes up less area than the hard interfaces used in most FPGAs. Flex Logic has developed embedded FPGA IP for several TSMC process nodes (40nm, 28nm, 16nm). In early 2017, it announced a partnership with DARPA. In May 2017, Flex announced a \$5 million funding with investors that included Lux Capital, Eclipse Ventures, and the Tate Family Trust. It is based in Silicon Valley.



- **Menta** – Menta develops embedded FPGA IP for SoCs and ASICs. Its Oragami software allows customers to customize the embedded FPGA to incorporate various numbers of embedded logic blocks (eLBs), embedded memory blocks (eLMs), and optional embedded customer blocks (eCMs). Its IP is available on several foundry processes (e.g., several TSMC 28nm processes, GlobalFoundries' 14nm FinFET and 32nm SOI processes, etc.) and Menta indicates its IP can be quickly ported to other CMOS processes. It has raised over \$7 million and is backed by FJ Development EN. The company was founded in 2007 and is based in Montpellier, France.
- **NVXL** – NVXL doesn't sell FPGAs, but has developed an FPGA-based architecture and software layer that is optimized for deep neural network and related AI-applications. FPGAs provide a variety of technical advantages in that they can perform many computations in parallel very rapidly (as opposed to sequentially like processors). However, managing, programming, and utilizing large numbers of FPGAs has historically been very difficult, limiting the use of FPGAs in the data center. NVXL has created an architecture and software abstraction layer that enables scaling to thousands of FPGAs to achieve extreme performance while simplifying management. NVXL believes its solution can disrupt the AI and cloud-based computing markets and enable wide-scale use of FPGAs and massive hybrid FPGA/GPU/processor systems. It recently announced a \$20 million round with Alibaba as the lead investor. In conjunction with the investment, the two companies have entered into a collaboration agreement to work together. It is based in Silicon Valley.

## Analog

---

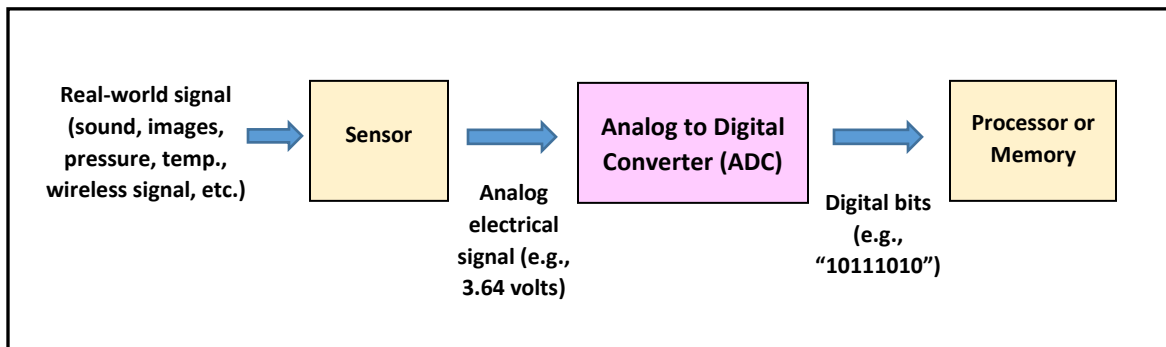
### Intro to Analog

Unlike digital components in which data is stored or processed as digital bits, analog integrated circuits deal with continuously varying voltages and currents. The input or output of an analog chip might, for example, be a voltage that can be any value between zero and 5 volts (e.g., 1.27 volts or 2.52 volts or 3.88 volts) and the exact value of the voltage matters. Some major categories of analog components include data converters, amplifiers, power management, and clocks.

The analog industry has a number of attractive attributes, especially for the high-performance analog component suppliers. There are often slight differences among analog components from different companies so once an analog component is designed into a product, OEMs are often reluctant to change. Because the major analog companies sell hundreds or thousands of different chips, there is generally less dependence on sales of any particular component or customer (although there are some exceptions). Whereas digital chips often become obsolete after a few years, analog chips can often continue to sell for decades. Unlike leading-edge digital chips, many analog chips can be produced on older process nodes and in less expensive fabs. As a result, many of the high-end analog companies have relatively high gross margins and operating margins.

### Data Converters

Real-world signals (images, sound, temperature, pressure, etc.) are inherently analog. As such, whenever there is a need to convert a real-world parameter to digital bits that can be processed by a computing device, or vice versa, a data converter is used. There are two general types of data converters. An analog to digital converter (ADC) converts an analog electrical signal to digital bits. It is commonly used after some type of sensor (e.g., image sensor, microphone, pressure or temperature sensor, etc.). The sensor converts the real-world signal to an analog electrical voltage (e.g., a microphone converts sound waves to an electrical signal), and then an ADC converts the analog electrical signal to digital bits that can then be stored in memory or processed by a microprocessor.



Source: Menalto Advisors

A digital to analog converter (DAC) performs the opposite function in that it converts digital bits to analog electrical signals. For example, to play music stored digitally, the digital representation of the song must be converted from digital to analog. The resulting analog electrical signal is then typically amplified and transmitted to a speaker which converts the analog electrical signals to sound waves.

ADCs and DACs are also used extensively in communications. For example, an incoming wireless signal is inherently analog and must be converted to digital bits for digital processing. However, for many wireless and wireline communications applications, the data converters are integrated into larger chips that contain a variety of other functions as well, rather than having stand-alone data converter chips.

Data converters are often characterized by how many times per second they can perform a conversion. Typically, this is given in terms of samples per second (e.g., a 100Mps ADC can convert analog to digital 100 million times per second). Another major characteristic is the number of bits a data converter uses on the digital side (resolution). An 8-bit ADC, for example, can provide 256 possible outputs ("00000000" through "11111111") whereas a 16-bit ADC could provide over 65 thousand possible outputs which results in greater resolution. For example, with a sensor that measures temperature between 0°C and 125°C, an 8-bit ADC could provide digital data in intervals of about 0.5°C (125 degrees range with just over 250 possible digital values). If you wanted resolution down to 0.1°C intervals, you would need an ADC with more bits. However, for some applications you may not need high resolution and it is helpful only if the sensor itself is highly accurate (if the sensor is accurate to within only 1 degree, for example, then having ADC resolution to 0.1°C isn't helpful).

In general, it is more difficult to develop data converters with higher sample rates and higher resolutions, and they tend to have higher power consumption and are more expensive. As such, engineers try to find the data converters that are best suited for the particular application.

There are a variety of different types of architectures for ADCs and DACs (the underlying technology to implement the data converters). For example, some well-known ADC architectures include SAR (successive approximation register), Pipelined, Flash, and Sigma-Delta. In some cases, data converter suppliers integrate multiple data converters in the same chip (multi-channel) or integrate other functions to provide a complete analog front-end (AFE).

By far, the largest supplier of stand-alone data converters is Analog Devices. Texas Instruments is a distant second, and other suppliers include Maxim and Renesas/Intersil. Although data converters are pervasive, many complex SoCs and communications chips integrate data converters, reducing the market for stand-alone data converter chips.

## Amplifiers

Amplifiers amplify electrical signals, and are used throughout electronic systems for a variety of reasons. For example, amplifiers are often used to amplify audio signals before it reaches a speaker. Amplifiers are also used extensively to amplify signals within electric circuits and between chips. For example, the output of a sensor is often amplified before a data converter converts the signal to digital. When a wireless signal is transmitted, amplifiers are used before it is sent and when it is received.

There are a variety of different types of amplifiers (variable gain, differential, isolation, etc.). One very popular type is an operational amplifier (op-amp) in which the amount of amplification can be modified by changing inexpensive external passive components (resistors, capacitors, etc.) which enables the same amplifier to be used for a variety of different levels of amplification. There are a broad range of different parameters associated with amplifiers such as the amount of amplification (gain), the speed at which the output can change when the input changes (slew rate), the extent to which the signal is distorted by the amplifier, and the range of currents and voltages that it can operate over. Amplifiers that operate at high frequencies (RF amplifiers) are often especially challenging to design.

Some of the major suppliers of amplifier chips include Analog Devices, Texas Instruments, Renesas/Intersil, and Maxim.

## Power Management

Power management is the largest major segment of the general-purpose analog market. In an electrical system, the various semiconductor components require power to operate, and the supplied power must be kept within very tight tolerances. A major role for power management chips is to help provide the appropriate voltages and currents to the rest of the electronics in a system.

There are a variety of different types of power management chips. For example, when an electronic device plugs into a wall outlet, the incoming electrical power is AC (alternating current) and must be converted into DC (direct current) for powering chips. An AC-DC controller chip (along with other components such as a rectifier, transformer, and power transistor) is used to accomplish this conversion from AC to DC.

This resulting DC power, or the DC power supplied from a battery in the case of battery powered devices, is generally not well-suited for sensitive semiconductor chips and must be further “regulated” or modified to make it more stable and to provide the precise voltages and currents needed. Voltage regulators are power management chips that transform an unregulated DC input to a highly stable DC output that is well-suited for powering the chips in the system. In many cases, the output voltage from the regulator will differ from the input voltage. There are two general categories of DC voltage regulators:

- **Linear Regulators** – Linear regulators use relatively simple, non-switching techniques to regulate the voltage output. With linear regulators, the output voltage must be less than the input and the difference between the input and output voltage is referred to as the “dropout.” In general, reducing the dropout makes a regulator more efficient and low dropout regulators (LDOs) are a type of linear regulator in which the chip can regulate even if the desired output is only slightly less than the input voltage.

- **Switching Regulators** – Switching regulators rapidly switch on and off, releasing energy when on. By adjusting the duty cycle (percentage of time it is on), the desired output voltage can be provided. Unlike linear regulators, the output of a switching regulator can be higher (boost) or lower (buck) than the input voltage. Switching regulators tend to be more efficient than linear regulators, and are generally smaller at high power levels. However, they are typically more expensive and can have EMI/noise issues due to the rapid switching.

One recent trend in some types of high volume electronics has been the development of optimized PMICs (power management integrated circuits). A complex product, such as a smartphone, requires dozens of different voltage regulators. OEMs would historically use many “off-the-shelf” chips that were closest to meeting their needs. However, there has been growing interest in customized PMICs that integrate many of the power management functions into a single chip that is specifically optimized for the product.

There are a variety of other emerging power management segments. As one example, there is growing interest in wireless charging in which battery-powered devices (smartphones, tablets, electric vehicles, etc.) can be powered without connecting a cable. There are a variety of different technologies for implementing wireless charging. Two of the major standards are Qi (which uses inductive coupling) and Rezence (which uses magnetic resonance, and also requires the device to have Bluetooth for communicating power needs). Each has advantages and disadvantages (which we won’t describe here) and each has several major technology companies backing it (e.g., Rezence backers include Intel, Qualcomm, IDT, Broadcom and WiTricity). There are also standards being developed for electric vehicle wireless charging (e.g., SAE International J2954).

Most of the major analog companies provide power management chips. Some of the major suppliers of DC regulators include: Texas Instruments, Maxim, Analog Devices (especially since it acquired Linear Technology), Infineon, Renesas (acquired Intersil), Monolithic Power Systems, ON Semiconductor, Microchip (Micrel), and Dialog (focus on PMICs). AC/DC controllers are a much smaller market and some major suppliers include Power Integrations, Infineon, ON Semiconductor, and STMicroelectronics.



## Clocks/Oscillators

Clocks continuously emit a series of pulses at very precise intervals. These pulses are generally used by other chips to “keep track of time.” For example, a microprocessor may be designed to execute an instruction every time it receives a clock pulse. Clock chips can also help synchronize different chips. A key parameter for clocks is jitter, which is a measure of how the actual emitted pulses differ from the ideal. For example, a 1Gbps clock should emit a pulse exactly every one billionth of a second, but some of the pulses may be emitted slightly early or late, which can cause issues if the level of jitter is too large.

The core timing element in most electronics is a crystal oscillator, typically quartz. In a crystal oscillator, a voltage creates an electric field which distorts the crystal, and when the voltage is removed the quartz returns to its previous shape, which generates an electric field and a voltage. A feedback circuit is used in such a way that the oscillator begins producing a periodic output at the resonant frequency of the crystal (or in some cases, a multiple of the resonant frequency).

This output, however, may not be exactly what is needed for an electronic system, so additional types of timing functionality are often added. A PLL (phase-locked loop), for example, can take in a low frequency clock input from a crystal oscillator and create a clock signal at a much higher frequency. Multiplier/divider circuits can create clock signals at multiples or simple fractions of the input frequency. Fanout buffers are used to take a single input clock signal and create multiple identical outputs at the same frequency in order to supply several chips with the same clock signal. Several of these functions can be integrated into the same chip.

As one example of a clock chip, the Silicon Labs Si5332 can use an external crystal (16-50MHz), an embedded crystal, or another type of frequency signal as input. It can then provide up to 12 clock outputs of any frequency up to 250MHz, with very low clock jitter (230 fs rms).

As previously noted, most clock chips utilize crystal oscillators in which a vibrating crystal is used to create an electric signal. However, there have also been MEMS-resonator based solutions that use MEMS devices instead of quartz crystals. For example, MegaChips acquired SiTime, which was focused on MEMS-based timing devices. Many of the analog companies provide clock chips (e.g., Integrated Device Technologies, Texas Instruments, Silicon Labs, Cypress, Analog Devices, etc.).

## Application Specific Analog ICs

The previously noted analog components (data converters, amplifiers, etc.) are classified as general purpose integrated circuits, as they tend to be “off-the-shelf” chips that can be used for a variety of applications. Analog companies generally provide catalogs of these components that engineers can choose from. However, there is also a broad range of analog solutions that are developed or optimized for specific products or applications. In some cases, they may integrate a variety of analog functions into a single chip. These application specific analog chips actually represent a larger market than general purpose analog ICs, but the market is extremely diverse.

## Select Analog Start-Ups

The following are some examples of private analog chip companies:

- **Active-semi** – Active-semi provides power management ICs and intelligent digital motor drive chips. Its products include intelligent motor controllers that integrate analog, power management, and ARM processor cores; PMICs (that integrate multiple DC/DC converters and LDOs for specific applications); and a variety of AC/DC, DC/DC, and battery management chips. In 2017, it indicated that it had shipped more than 1.5 billion power-related components and had over 150 patents. Investors include USVP, LG, Tanaya Capital, Selby Venture Partners, and LDV. It is based in Dallas, Texas.

- **Adsantec** – Adsantec develops high-speed analog and mixed-signal chips. Its products include multiplexers/demultiplexers, amplifiers, drivers, transimpedance amplifiers, high speed logic, data converters, clock generators, and other high performance analog components. It generally focuses on developing high speed solutions that are not readily available from other suppliers. It has over 160 products and is based in Southern California.
- **AnDapt** – AnDapt provides innovative customizable power management chips. There is growing interest in optimizing power management for products. However, developing a custom power management chip from scratch is costly and time consuming. AnDapt had developed power management solutions that can be configured (varying voltage, current, impedance, etc.) much in the same way that FPGAs can be programmed. Its solutions use analog power blocks with scalable integrated MOSFETs and digital interconnect circuitry which can configure the chip to meet the desired specifications. Leveraging this core configurable power management technology, it provides a variety of power management components including regulators and controllers. Investors include Cisco, Intel, Atlantic Bridge, and Vanguard Semiconductor. It is based in Silicon Valley.
- **Aspinity** – Aspinity develops ultra-low power analog signal processing (ASP) chips. Aspinity believes there are a number of applications in which digital processing tasks can be better performed in the analog domain, greatly reducing power consumption and the amount of data that must be processed. Its Reconfigurable Analog/Mixed-Signal Processor (RAMP) enables certain digital signal processing tasks to be replicated in analog. One initial target application is voice/audio processing, as RAMP can significantly reduce power consumption while waiting for “wake words.” Other applications include health monitoring, touch sensing, and various sensor networks. Aspinity was part of the inaugural Amazon Alexa Accelerator, an Amazon program that supports voice-related startups. It is based in West Virginia.

- **Ferric** – Ferric Semiconductor develops DC/DC power chip converters and IP that utilize a unique thin film inductor technology. Ferric has developed a magnetic core inductor process that allows the integration of Ferric inductors with TSMC’s CMOS “back end of line”, enabling very efficient, high density power management solutions. Importantly, its integrated voltage regulator technology enables power management to be integrated into a chip package or even directly into a digital chip (e.g., a microprocessor or SoC), which provides a variety of significant technical, size, and cost advantages. In 2017, it announced a partnership with TSMC. Investors include New Science Ventures. It is based in New York City.
- **IQ-Analog** – IQ-Analog originally focused on licensing data converter IP and it licenses a broad range of ADCs, DACs, and analog front-end cores. However, its more recent focus has been on its Next family of configurable data converters, in which a single data converter chip can be used for a broad range of applications. These configurable ADCs and DACs have resolution ranging from 8 to 16 bits and speed grades ranging from 10Msps to 64Gsps. IQ-Analog also provides integrated solutions such as its NXT-F1000 radar chip (which integrates four 8-bit 64Gsps ADCs, four 8-bit 64Gsps DACs, phased array processors, SerDes and an embedded controller and clock). It has several patented technologies including its “Traveling Pulse Wave Quantizer” for ADCs and its “Current Impulse” DAC technology. It announced a \$4.5 million DARPA contract in 2016 and a strategic investment by Lockheed Martin in early 2017. It is based in San Diego, CA.
- **Kinetic Technology** – Kinetic provides a variety of analog components including power management (DC/DC converters and regulators), LCD and LED drivers, and Power over Ethernet (PoE) controllers for a broad range of markets (consumer, industrial, telecom). In 2017, it introduced the AS1114 (a single chip, highly integrated power solution targeted at 13W PoE). Kinetic is a Cayman Corporation with its major office in Silicon Valley.

- **Linear Dimensions** – Linear provides a variety of custom and semi-custom analog and sensor solutions, and has developed a number of proprietary technologies for applications such as biomedical sensing. Its core PEAL technology uses machine learning engines to separate signals from noise. Applications include Cardiaclock (which can uniquely identify a person through heart monitoring) and various biometric sensing solutions. It also sells a variety of analog components (MEMS drivers, amplifiers, photo detectors, etc.). It is based in Silicon Valley.
- **Lion Semiconductor** – Lion develops power and battery management chips. It indicates its initial two solutions include a DC/DC PMIC, which eliminates the need for many passives, and an innovative battery charger chip. Lion notes that these solutions use a different architecture than alternatives and are ultra-small and fast, with high efficiency. Its technology incorporates robust digital control. Specific product details, however, have not yet been released. Investors include Atlantic Bridge and Walden. It is based in San Francisco.
- **Omni Design** – Omni Design develops ultra-low power analog IP cores, with a focus on high-speed data converters. Its high-performance ADCs utilize its patented SWIFT technology. SWIFT permits the use of simple low power op-amps to be employed in place of highly complex, high power alternatives in high performance circuits, and can be used to increase analog channel speed, reduce power consumption, and dramatically increase circuit performance. Omni's offerings include low power 10/12/14-bit pipeline ADCs, low power SAR ADCs, high linearity DACs, and analog front-ends. It also provides IP for sensor applications and equalization and pre-emphasis IP for wireline communications. It offers its IP at a range of process nodes (e.g., 28nm HPC, FinFET, 40nm LP, etc.). In 2016, Omni announced MegaChips as a partner. Omni indicates that its investors include semiconductor industry leaders. It is based in Silicon Valley.

- **S3 Semiconductor** – S3 Semiconductor designs and sells complex mixed-signal ASICs, primarily to OEM customers. Many OEMs don't have chip expertise, but would like to have single chip SoCs rather than dozens of standard components. S3 Semiconductor designs chips that meet an OEM's requirements (typically integrating a variety of analog, RF, and processors into one chip) and then supplies the chip to the OEM. It indicates that this can reduce an OEM's costs by up to 80%. S3 Semiconductor also has a semiconductor IP business that licenses analog IP (mostly data converters) to large chip companies. S3 Semiconductor is part of S3 Group. Investors include ACT Ventures and management. It is based in Dublin, Ireland.
- **SiBreeze** – SiBreeze provides power management ICs. Inductors are one of major limiting factors in reducing the size and increasing the performance of power chips. SiBreeze has developed advanced technologies (Best Switcher, Slim Switcher) for reducing inductor value by 10X relative to a traditional circuit at the same frequency. This enables much smaller/thinner size as well as faster speed and better performance. It indicates that it can provide better than 4 phase response, efficiency, and noise with a single inductor. Leveraging this core technology, it introduced its Si117, 2.5 volt to 14 volt switching regulator for mobile battery charging. Its roadmap includes an advanced mobile CPU power chip. It is based in Silicon Valley.



## Compound Semiconductors

---

### Gallium Arsenide and Other Materials

The vast majority of semiconductor chips are made from silicon, as the industry has built up considerable expertise in silicon chip manufacturing. However, there are a variety of other materials that can be used aside from silicon. In fact, the first transistor was made from germanium. Silicon is generally the default material for chips unless there are specific reasons why other materials are needed. Most non-silicon chip materials are compounds, and are collectively referred to as compound semiconductors. In most cases, non-silicon components tend to be analog, discrete, or optical/LED components (digital chips are primarily silicon).

Compound semiconductors such as gallium arsenide (GaAs) or indium gallium arsenide (InGaAs) provide some technical advantages over silicon such as higher electron mobility, less noise at high frequencies, and wider bandgap, which provides advantages for certain high frequency applications. Examples of applications for GaAs or InGaAs chips include mobile phone power amplifiers (which must greatly amplify signals so they can reach a base station) and high frequency millimeter wave applications (e.g., satellite communications, radar, etc.).

One segment within semiconductors that extensively uses non-silicon materials is light emitting devices such as LEDs and lasers, since the output wavelength produced is dependent on the type of material used. For example, red LEDs are often made from aluminum gallium arsenide. Similarly, green and blue LEDs are made from various non-silicon materials (this is discussed in more detail in the Optoelectronics section).

### Gallium Nitride (GaN)

One type of material that has received strong interest recently is gallium nitride (GaN). GaN (and related compounds) has historically been used for blue or green LEDs, but there has been growing interest in using the material for power semiconductors and high-frequency RF applications. GaN has several attractive characteristics (very high breakdown voltages, high bandgap) that make it technically well-suited for applications that require high power and/or high frequency.

Historically GaN was very expensive to produce with low yields and a variety of other technical issues. However, with growing interest in advanced power management and high frequency communications, there has been renewed interest in GaN.

While silicon chips are manufactured on silicon wafers and GaAs chips are produced on GaAs wafers, there are a variety of different types of wafers used for GaN chips. “GaN-on-GaN” solutions using GaN wafers are used in some cases, but GaN wafers are very expensive and are generally available only in small sizes. Many LEDs use GaN-on-sapphire, with GaN transistors implemented on sapphire wafers. Silicon carbide (SiC) has attractive thermal properties and GaN-on-SiC is popular for some applications.

There has also been growing interest in GaN-on-silicon technology (GaN transistors on silicon wafers) as silicon wafers are inexpensive and are available in large sizes which could help drive down GaN costs. The challenge is that silicon and GaN have different thermal properties and lattice constants and are not very compatible. At high temperatures during manufacturing, this can cause separation of the materials and poor yields. As such, companies typically must implement some type of innovation (e.g., buffer layers) to help prevent these issues. With TSMC ramping up its GaN-on-silicon foundry services, it is expected to become increasingly popular.

### Select GaAs and GaN Chip Companies

There are a handful of companies that historically focused primarily on compound semiconductor chips, such as gallium arsenide ICs for wireless communications (power amplifiers, low noise amplifiers, etc.). Examples include Qorvo, Skyworks, II-VI, and Avago (which is now part of Broadcom Ltd.). These companies generally have their own compound semiconductor fabs, although more recently there have been foundries, such as WIN Semiconductor, that focus on compound semiconductor manufacturing, and as noted above TSMC now has a focus on GaN manufacturing. Most of these compound semiconductor companies have also developed GaN chips for RF applications, although there are a variety of different approaches. For example, Qorvo has focused on GaN-on-SiC solutions. As one example, in June 2017 it announced a GaN-on-SiC front-end module for 39GHz communications for 5G applications.

For power management applications, GaN often offers higher power efficiency at higher switching frequencies and many of the major analog chip companies have introduced GaN-related power solutions, including Infineon, Texas Instruments, Dialog, ON Semiconductor, and others. Of course, all the LED manufacturers utilize compound semiconductors. With GaN foundry services becoming more widely available, GaN should grow in popularity over the next few years.

### Many Other Materials

While GaN appears to be the non-silicon segment attracting the most interest right now, there are a variety of other materials being used and/or investigated for use in various applications (e.g., silicon carbide, aluminum nitride, graphene, carbon nanotubes, etc.). Due to space constraints we won't describe these in detail, but each has a variety of significant potential technical advantages if costs can be driven down.

### Select Compound Semiconductor Private Chip Companies

Some examples of non-silicon semiconductor companies are noted below:

- **Carbonics** – Carbonics has developed technology for greatly enhancing CMOS by adding small amounts of carbon nanotubes. It indicates that this can dramatically improve performance and linearity, significantly reduce power consumption (by 1000X in some cases), while still being compatible with CMOS manufacturing processes. The linearity of its solutions can be more than 10X that of gallium arsenide. Its Zebra wafers include a 1nm aligned layer of carbon nanotubes on top of the wafers. Its Viper devices and amplifiers include advanced high frequency chips. Stingray is a line of ultra-efficient millimeter wave ICs (using its core carbon nanotube technology) for 5G and other markets that Carbonics expects to introduce in 2018. It is based in Culver City, CA.

- Exagan** – Exagan is developing advanced GaN components. The company has created a GaN buffer technology that enables the growth of high quality GaN material on a silicon substrate, resolving one of the major GaN-on-silicon issues (GaN and silicon have different thermal coefficients so there is often cracking at high temperatures with conventional approaches, resulting in low yields). Its patented “G-Stack” approach (adding insulating and strain management buffer layers between silicon and GaN layers) helps resolve many of the historical GaN-on-silicon issues. Leveraging this technology, Exagan is developing a family of high-power GaN components. Its initial two “G-FET” solutions are GaN-based 650-volt and 1,200-volt fast-switching power devices. In 2017, Exagan announced that, in conjunction with X-Fab, it demonstrated mass-production of high-voltage power devices on 200-mm GaN-on-silicon wafers using standard X-Fab CMOS production. It was founded in 2014 with support from CEA-Leti and Soitec, and is based in Grenoble, France.
- GaN Systems** – GaN Systems manufactures a range of Gallium Nitride high power transistors for a variety of applications. Its solutions use GaN-on-silicon and it has a number of patented technologies including its “Island” topology which improves isolation, reduces inductance, and takes current from the transistors vertically, all of which improve performance. It has also developed advanced near chip-scale embedded packaging (GaNpx). Its initial products include 100 volt and 650 volt families of E-HEMT transistors. In July 2017, GaN Systems announced a strategic investment from BMW i Ventures. Other investors include BDC Capital, Chrysalix Venture Capital, Cycle Capital, RockPort Capital, and Tsing Capital. It was originally based in Ottawa, but moved its headquarters to Silicon Valley.
- Navitas Semiconductor** – Navitas develops GaN power ICs that it indicates enables 100x increase in switching speed combined with a 40% reduction in energy savings. Unlike many GaN providers that develop discrete transistors, Navitas has been able to utilize GaN’s lateral structure and new design approaches to enable the integration

of power FETs, drivers, and logic into GaN integrated circuits, which provides many advantages and resolves several issues (including bottlenecks from silicon drivers that can't keep up with GaN components). In 2017, it announced partnerships with TSMC and Amkor. It recently introduced an ultra-small 65W USB-PD laptop adapter reference design, which uses Navitas' high-efficiency GaN power ICs to deliver 65W in a package that is up to 5 times smaller and lighter than traditional silicon-based designs. It has also announced single and half-bridge GaN Power ICs. Investors include Atlantic Bridge, Capricorn, and MalibuIQ. It is based in El Segundo, CA.

- **Nitride Solutions** – Nitride Solutions has developed solutions for commercial deployments of aluminum nitride technologies. Aluminum nitride (AlN) is ideal for deep UV LEDs (which can be used for sanitizing water and objects) and has many attractive electrical properties (wide bandgap, high thermal conductivity, high voltage resistance), but was historically very difficult to manufacture. Nitride Solutions has developed a broad range of leading-edge aluminum nitride solutions including templates for UV LEDs; advanced coating materials with very high thermal conductivity; and polycrystalline materials for high-temperature piezoelectric devices and sputtering targets. It has a production facility in Kansas, and announced a \$2.75 million round in 2017 with investors that include Nelnet, Leawood Capital Venture Fund, and DK Group. It is based in Wichita, Kansas.
- **Transphorm** – Transphorm develops GaN semiconductor chips, with a focus on high voltage GaN transistors (FETs). Many of its FETs operate at 600 or 650 volts. It indicates that its 3205 650V GaN FET is the first GaN solution to earn AEC-Q101 qualification for automotive applications (it uses GaN-on-silicon manufacturing). Transphorm has over 600 patents. In 2015, it announced a \$70 million capital raise with investors that included KKR, Kleiner Perkins, Foundation, Google Ventures, Soros Quantum Strategic Partners, INCJ, and Fujitsu. In November 2017, it announced a \$15 million investment from Yaskawa Electric, which is also a customer. It is based in Goleta, California.

- **VisiC** – VisiC develops GaN-based high-power transistors and modules. Its initially announced solutions operate at 650 and 1,200 volts, including a 1,200-volt power switch. It indicates that its GaN solutions provide extremely high efficiencies at high voltages. Target markets include motor control, power supplies, solar inverters, and other high-power markets. In early 2017, it announced an \$11 million capital raise. Investors include Genesis Partners and Inversiones Plasma. It recently announced that David French (former CEO of Cirrus Logic) joined its board. It is based in Israel.



## MEMS Components

---

### Intro to MEMS

One type of semiconductor device that has become increasingly popular for sensors and other applications is MEMS (Microelectromechanical systems). With MEMS, semiconductor manufacturing methods are used to create tiny electro-mechanical structures. Traditional semiconductors don't have any moving parts (transistors don't physically move). However, using the same processes that are used for chipmaking (deposition, photolithography, etch, etc.), it is possible to create tiny moveable mechanical structures on semiconductor wafers that can be used to sense or move objects.

MEMS can provide significant advantages in terms of size and reliability over traditional sensors. Since a MEMS sensor is produced using a chipmaking process, it is also possible to integrate other semiconductor functions on the same chip, which can improve performance and reduce size. Additionally, MEMS benefits from semiconductor cost reductions in high volume. Although MEMS manufacturing is available from many foundries, the production processes aren't as standardized as conventional CMOS, and some MEMS devices require specialized steps.

In many cases, MEMS sensors are somewhat similar to conventional sensors but are implemented in a much smaller size. For example, one general method to measure pressure is to use a piezoelectric material. The electrical resistance of this type of material changes when pressure is applied, so by continuously measuring resistance, pressure can be determined. Both conventional pressure sensors and MEMS sensors often use this same approach, but the size of the MEMS sensor can be a tiny fraction of the size of a more conventional sensor, and the MEMS sensor can also directly integrate other functions (e.g., ADCs, amplifiers, etc.).

## MEMS Examples

Some examples of MEMS solutions include:

- **Accelerometers/Gyroscopes** – MEMS sensors are often used for measuring acceleration (accelerometers) and orientation (gyroscopes). MEMS accelerometers became pervasive in cars for air-bag sensors. Subsequently, MEMS accelerometers and gyros have become commonplace in smartphones and tablets (e.g., determining orientation), GPS navigation (“dead reckoning”), game controllers, and many other consumer products.
- **Microphones** – Over the past few years, the industry has shifted to MEMS microphones for most high-volume consumer applications. Smartphones, for example, incorporate several MEMS microphones.
- **Optical Switching** – It is possible to build tiny MEMS mirrors that can be controlled and rapidly electrically tilted in different directions. This can enable switching of light in a variety of directions. One application for this is networking/telecom optical switching, as it can allow switching of fiber optic light signals without having to convert the light to electrical and back to light. MEMS mirror technology has also been used extensively in consumer applications, such as TI’s DLP (digital light projection) technology which uses an array of tiny MEMS mirrors to direct light for projection. There are several MEMS image scanning/projection companies (e.g., MicroVision).
- **RF Switches/Filters** – Another growth segment for MEMS has been RF switches for wireless applications. Cellular handsets typically incorporate a variety of switches to, for example, switch between different frequency bands. MEMS switches can offer many advantages over traditional switching technologies.
- **Pressure/Temperature** – MEMS is increasingly used for traditional sensor applications such as measuring pressures and temperatures.
- **Material Sensing** – An emerging area of MEMS sensors is using MEMS to detect gases, chemicals, or human biometric properties.

## Major MEMS Suppliers

According to Yole (May 2017), the MEMS market is projected to grow from nearly \$15 billion in 2018 to over \$25 billion in 2022. Some of the major suppliers of MEMS components include Bosch (automotive sensors and MEMS mics), Broadcom (Avago RF switches), Texas Instruments (DLP and optical switching), STMicroelectronics (gyros, accelerometers, and other sensors), Qorvo (RF Filters), HP (MEMS print heads), Knowles (MEMS microphones), TDK (InvenSense gyros), Analog Devices (accelerometers and RF switches), and Murata (various sensors).

## Select MEMS Sensor/Actuator Start-Ups

The following are examples of privately-held MEMS companies (we included one or two sensor companies that don't technically use MEMS technology, but this seemed like a good place to include them).

- **Butterfly** – Butterfly developed a highly integrated ultrasound imaging solution that integrates MEMS and deep learning algorithms. The solution uses a single ultra-wideband (UWB) 2D matrix array comprised of thousands of MEMS sensors which are directly overlaid on an ultrasound processing chip. Leveraging this core technology, it introduced its iQ handheld ultrasound scanner in 2017 (works with a smartphone), with plans for shipments to begin in 2018. The price is under \$2,000 and Butterfly indicates that it is cleared by the FDA for 13 applications. One market is poor regions that don't have access to conventional medical imaging hardware. The company has raised more than \$100 million, and has offices in Connecticut and New York.
- **Cavendish Kinetics** – Cavendish develops MEMS-based RF components for wireless applications, including cell phones. Its initial product line of SmarTuners is a family of digitally variable RF MEMS capacitors for tuning smartphone antennas across various bands. It subsequently introduced a variety of additional RF MEMS solutions including transfer switching (switching between an LTE main antenna and a diversity antenna) and other RF switches. In January 2017, it announced that its SmarTuners were being used by 40 different

smartphone models from 10 OEMs including the Samsung Galaxy A8. In 2015, it announced a \$36 million Series F round which included a \$25 million investment from Qorvo/TriQuint. Investors include Qorvo, Qualcomm, Tallwood Venture Capital, and Wellington Partners. It is based in Silicon Valley.

- **EnerBee** – EnerBee has developed a MEMS-based energy harvesting generator, which uses a combination of magnetic and piezoelectric materials. Unlike many energy harvesting solutions that require fast powerful movements, the EnerBee technology works with even low speed and low force motions. The energy can then be used to power sensors or low-power radios (e.g., for sensor networks or other applications) or can be stored for later use using EnerBee’s power management technology. Its technology was developed from a 5-year R&D collaboration that included CEA-LETI and several scientific institutes in Grenoble. In 2015, EnerBee raised €2.5 million from investors that included Robolution Capital and Emertec. EnerBee is based in Grenoble, France.
- **Leman Micro Devices** – Lemman Micro has developed sensor solutions (V-Sensor) that work in conjunction with a smartphone to measure health and biometric parameters including blood pressure, heart rate, temperature, and respiration rate. Blood pressure, in particular, has been challenging to accurately measure without a traditional “cuff” machine, but is critical in detecting many potential medical issues. By enabling consumers to easily measure these parameters anywhere, the company believes many medical issues can be detected early. In addition to the sensor, it has developed an associated smartphone app (Elemdu). For measurement, users press their finger on the V-Sensor (no cuffing) and run through the smartphone app. The sensor integrates MEMS, optical sensing, and other technologies. Lemman expects to obtain medical regulatory approvals in 2018. It indicates that it has two major smartphone companies as strategic investors. It is based in Lausanne, Switzerland.

- **mCube** – mCube develops MEMS accelerometers for consumer and industrial applications (wearables, Internet of “moving things”). Whereas some MEMS solutions use different processes for the MEMS sensor and the companion drive chip, mCube’s MEMS sensors are fabricated on top of the IC electronics in a standard CMOS fabrication facility. It indicates this provides advantages such as reduced cost, smaller size, higher performance, lower power consumption, and the ability to more easily integrate multiple sensors together. The company also introduced a MEMS gyro and eCompass solution. In late 2017, mCube announced it acquired Xsens (sensor fusion and motion-tracking modules) from ON Semiconductor. Investors include Kleiner Perkins, SK Telecom, DAG Ventures, Keystone Ventures, iDVentures America, and MediaTek. It is based in Silicon Valley.
- **MEMSDrive** – MEMSDrive develops MEMS actuator chips, with a focus on image stabilization for cameras. Its OIS actuator provides fast and accurate image stabilization against the motion of a camera. Applications include smartphone cameras, digital cameras, automobiles, drones, and AR/VR devices. In 2016, it announced a partnership with Oppo for smartphone cameras. In late 2016, it announced an \$11 million capital raise led by Walden International. The company is based in Pasadena, California.
- **Micralyne** – Micralyne is a semiconductor foundry that is specifically focused on MEMS manufacturing. The company has a 55,000 ft<sup>2</sup> facility offering wafer fabrication lines, assembly, test, and quality assurance. It fabricates structures on 6” silicon wafers (including standard, single and double SOI), glass and silica, and other materials. Historically, it has had a strong focus on biomedical MEMS (drug delivery), industrial sensors, and telecom (optical MEMS). In 2017, it announced the availability of standard silicon process technologies for gas sensors. Micralyne was acquired by FTC Technologies in 2015. It is based in Edmonton, Canada.

- **NextInput** – NextInput provides MEMS-based force-sensing solutions for touch enabled devices. It addresses a variety of markets (smartphones, consumer, wearables, IoT, automotive, industrial). Its technology enables touch selection without mechanical buttons and it can also provide multi-level touch in which the degree of pressure can be determined, providing an additional user interface dimension. It indicates its touch sensors are more than an order of magnitude more reliable than mechanical buttons and meet automotive standards (AEC-Q100 qualification). In mid-2017, it announced a solution specifically optimized for side buttons on mobile phones. It recently announced a companion analog front end for its sensors. In addition, it also introduced ForceGauge, a new family of high performance strain sensors, further expanding NextInput's market opportunity. It had a Series A/A1 funding of about \$13 million in 2016, with investors that included: Sierra Ventures, Intel Capital, GoerTek Group, Danhua Capital, UMC Capital, and Cota Capital. It is based in Silicon Valley.
- **Peratech** – Peratech has developed force touch sensing solutions which help improve touch controls. Its solution is based on patented Quantum Tunneling Composite technology which transitions from being insulators to providing increasing conductivity as applied force grows. Peratech indicates its solution is multi-touch and highly accurate, with much lower cost than alternatives. It also enables users to make different choices based on the amount of pressure applied, with no moving parts. Its technology can be utilized for many markets, including automotive (its force-based solutions can eliminate unintentional activation and help ensure safety), smart home, consumer, and industrial products. It recently began demonstrating its technology in a smartphone. Peratech announced a partnership with SMK Electronics. In October 2017, it raised \$12.4 million with investors that included Merck Ventures and Arie Capital. It is based in the UK.



- Preciseley** – Preciseley Microtechnology develops advanced MEMS-based optical switching chips. It sells electrostatic-driven MEMS solutions that incorporate tiny micro-mirrors which can be tilted and very precisely controlled for beam steering. Actuation is by ultra-low current analog voltage input and advantages include large mirror size, high switching speed, low power consumption, and stable tilting angle. It initially leveraged this technology for optical communications applications (MEMS mirror arrays for high speed light signal switching) and has numerous major telecom customers with growing interest in data center applications. In addition, it has recently utilized this technology for sensor applications including an advanced MEMS spectrometer solution. The company is based in Edmonton, Canada.
- Pyreos** – Pyreos develops MEMS-based infrared light sensors based on the pyroelectric effect, in which a change in light level produces a change in temperature which, in turn, produces a change in voltage across a crystal. By measuring the voltage, the light level can be determined, and its solutions are especially well-suited for the infrared spectrum. Leveraging this core technology, it has developed a variety of solutions including IR sensors, line array sensors (for applications such as solid-state IR spectrometers and IR imaging devices) and ezPyro digital integrated sensor solutions (for applications such as motion detection and gas sensing). Investors include Bosch, Siemens, Mitsubishi, Braveheart Investment Group, Scottish Investment Bank, and Seraphim. It is based in Edinburgh, Scotland.
- Sheba Microsystems** – Sheba has developed a MEMS-based solution for camera auto-focus and stabilization. Smartphone cameras generally use traditional mechanical devices called voice coil motors (VCMs) to make lens adjustments for auto-focus. VCMs have many limitations in terms of size (hard to shrink) and performance (adjustment speed). For video applications in particular, they significantly limit camera performance. Sheba has developed a unique MEMS-based actuator that can move large masses at high speed, and its first target application is smart camera autofocus and image

stabilization. Unlike alternatives that move the lenses, its solution moves the image sensor which it indicates provides many advantages. It notes that its technology can also be used for 3D touch sensing and other applications. Sheba is based in Toronto, Canada.

- **Tikitin** – Tikitin is focused on utilizing MEMS technology for the timing/clock market. The company indicates that the core technology for its products has been developed through more than a decade of research. However, in 2016 the company determined that the technology was viable for commercial deployments. It believes its MEMS resonator technology has significant advantage over quartz and other MEMS products, and its initial focus is on the 10-50MHz market segment. It expects commercial shipments to commence in 2018. The company is based in Finland.
- **Vesper** – Both conventional microphones and most MEMS mics use capacitive sensing (sound waves move a flexible diaphragm and the capacitance between the moveable diaphragm and a fixed plate is used to capture the audio). In contrast, Vesper uses patented piezo-resistive MEMS technology, in which the sound pressure changes the resistance of a material, which is continually measured and can be used to convert audio to electrical signals. Vesper indicates this approach has many advantages relative to capacitive approaches (e.g., the piezoelectric material is free to move, whereas the air trapped between two capacitive layers prevents free movement which makes it less sensitive). Whereas capacitive mics degrade over time as they are exposed to the environment, the Vesper solution is nearly immune to contaminants, providing better performance over time. Vesper's recently introduced V2000 is specifically optimized for smart speakers/voice interfaces/IoT devices and Vesper indicates it can reduce material costs by 50%. Vesper has a partnership with Synaptics. In 2017, it announced a \$15 million capital raise with investors that include Amazon's Alexa Fund, Accomplice, AAC Technologies, XinGang, and Hyperplane. It is based in Boston.

- **Xsensio** – Xsensio has developed a MEMS-based “Lab-on-a-chip” solution that analyzes tiny amounts of sweat on a person’s skin (including electrolytes, metabolites, small molecules, and proteins) for health/wellness analysis. The sensor is highly miniaturized (10,000 sensor arrays fit on a 0.1mm<sup>2</sup> footprint) with low power consumption (100microWatts). It can be integrated into wearables or patches and can provide real-time feedback about a person’s health and wellness, and can potentially detect certain health issues. In October 2017, it announced that it will lead the XPATCH consortium funded by Eurostars to develop, with four European partners, the next-generation wearable skin patch for ultrasensitive real-time health and wellness monitoring. The company is based in Lausanne, Switzerland.


## Optoelectronics

### Intro to Optoelectronics

Optical components create, detect, or modify light. Some common types of optical components include LEDs, laser diodes, image sensors, and solar cells. This section briefly discusses these types of components.


### LEDs

As previously noted, a diode is a simple semiconductor device (PN junction) which turns on and conducts current when the voltage across it exceeds a certain threshold. A light emitting diode (LED) is similar to a conventional diode but when the diode is on and current is flowing it emits light. Specifically, when the free moving electrons in the LED combine with the free holes (atom with missing electrons), photons are released.



LED suppliers optimize the design of the LED and the packaging around it to maximize the amount of light that is emitted. As shown in the diagram to the right, LED manufacturers will often add a reflecting cavity and lens to enhance and direct the light from the LED chip.

Importantly, the color (wavelength) of the emitted light is dependent on the type of material the LED is made from. For example, red LEDs (around 680nm wavelength) can be made from aluminum gallium arsenide (AlGaAs) or gallium arsenide phosphide (GaAsP); green LEDs (around 540nm wavelength) can be made from aluminum gallium indium phosphide (AlGaInP) or gallium nitride (GaN); blue LEDs (about 475 nm wavelength) might be made from indium gallium nitride (InGaN) or zinc selenide (ZnSe). These are some examples, but a variety of other materials have been used.




High quality white light, which is most common for general illumination, has been challenging to produce. One approach is to combine the output from three different LEDs (red, blue, green) to produce white. However, most “white LEDs” use blue or ultraviolet LEDs with a yellow phosphor coating around it, which produces a white light. The phosphor, however, reduces the output, and over time the phosphor can degrade resulting in changes to the shade of white. A variety of alternative types of white light LEDs have been under development.

## Semiconductor Lasers

The principle behind lasers is an atomic property in which an incoming photon can cause an electron in a high energy state (excited atom) to drop to a lower energy state which, in turn, causes the emission of a second photon that is identical to the first. This is referred to as “stimulated emission.” These photons can then cause the same thing to happen in other excited atoms, producing more identical photons. When this reaches a certain critical level (inversion), “lasing” occurs and a coherent powerful light is emitted.

Semiconductor lasers (laser diodes) are similar to LEDs but produce output that is generally more powerful and more uniform. Whereas an LED’s output is spread over a broad range (e.g., a 540nm LED might produce output ranging from 530nm to 550nm), a laser diode’s output is monochromatic with a tight range around a single wavelength. A



laser’s output is also coherent (all of the light is in phase) and collimated (travels in the same direction). This makes it well-suited for applications such as fiber optic communications or even laser pointers (a laser can produce a clear dot across the room, whereas the light from an LED would spread out and not produce the desired effect). Semiconductor lasers are similar to conventional LEDs except they generally have one or more intermediate layers between the “P” and “N” materials, in which the lasing occurs. During manufacturing, by creating very straight edges or using materials with certain refractive indices, reflective surfaces (mirrors) can be made within the chip, which can confine photons to a small region to produce lasing.

There are a variety of different types of semiconductor lasers such as double heterostructure (one material is sandwiched between two layers of another material which confines the active region to the middle layer), quantum well (a thin middle layer can act as a quantum well which has high efficiency), distributed Bragg reflector (incorporates two mirrors, one is optimized for a specific wavelength, the other is partially reflective to enable emission), and distributed feedback (a diffraction grating optical filter is etched into the chip which filters out non-desired wavelengths, providing an output at a very precise wavelength).

Most semiconductor lasers are “edge-emitting” in that the emitted light is perpendicular to the direction of current flow. In most diagrams, the “P” and “N” layers are on the top and bottom and the light is emitted out of the side (from the “edge”). With a VCSEL (Vertical Cavity Surface Emitting Laser), light is emitted from what is typically shown as the top of the structure, in the same direction as current flow. VCSELs generally produce less output power than edge-emitting lasers, but are typically much less expensive to produce. VCSELs can be tested while in wafer form, whereas edge-emitting lasers cannot, which results in higher yields and lower costs for VCSELs. They also generally use less expensive packaging.

An emerging area for semiconductor lasers is automotive LiDAR, in which the time for a light signal to reach an object and return can be used to determine the distance to the object. The initial LiDAR systems used traditional lasers, but these are very bulky, expensive, and unreliable. As a result, the industry is transitioning to semiconductor-based LiDAR solutions.

## Optical Detectors

The basic principle of optical detection is that when an incoming photon of light is absorbed by an atom, the energy can free the outer electrons of the atom, creating free negatively-charged electrons as well as positively-charged ions. With the appropriate circuitry, these free charges are swept away and create a current, which can be measured. Thus, the current produced is directly related to the amount of light received (incoming photons). In many circuits, this current is converted to a voltage and then to digital bits. The digital data can then be processed and analyzed to determine the amount of light that was received.



There are a variety of different types of optical detector components. The most basic is a photodiode. This is essentially a diode (PN junction) although as with laser diodes, photodiodes often include an intermediate intrinsic region (PIN diode) to improve efficiency. When a photon with sufficient energy is absorbed by the photodiode, an “electron-hole” pair is created, and an electric field causes the electron and hole to move in opposite directions, creating a current.


Photodiodes are commonly used, but have some limitations. For example, typically a relatively small amount of current is produced from incoming photons and, additionally, a small amount of current is produced even when no light is received. As such, it is difficult to discern small changes in light (not highly sensitive) or detect very low levels of light.

One alternative is an avalanche photodiode (APD). Due to an APD’s semiconductor properties, when a photon generates an “electron-hole” pair, it causes an “avalanche” creating dozens of additional “electron-hole” pairs, resulting in a much higher current than for a regular photodiode, increasing the sensitivity. However, some disadvantages are that it requires a high negative operating voltage (typically not desirable), has relatively high-power consumption, and is often more expensive than photodiodes.

## CMOS Image Sensors

Image sensors capture the light reflected from real-world objects and convert the light to analog electrical signals which can then be converted to digital bits that can be stored or processed. CMOS image sensors are the most popular image sensors used in electronics and are commonplace in smartphones, digital still cameras, and other video capture electronics. A CMOS image sensor consists of a matrix (rows and columns) of photo detectors, each of which captures a tiny portion (a pixel) of an image. Each “pixel” in the image sensor includes a micro-lens (which amplifies the incoming light), a color filter (which enables light of only a certain color to pass), a photodiode (which determines how much light is received), and other active pixel transistors (for resetting, eliminating noise, amplifying, and communicating its data). The output from each pixel is amplified and converted to digital bits, which can be processed to create a digital image.

Photodiodes detect only how much light is received and not the color or wavelength of incoming photons. As such it cannot by itself determine color. However, by using a pattern of color filters (such that some photodiodes measure only blue light, some only red light, and some only green light), the color of images can be determined. The diagram to the right illustrates this although a real image sensor has millions of pixels.



Some of the key parameters for CMOS image sensors are the number of pixels (e.g., the iPhone 8 has a 12-megapixel camera), frame rate (e.g., 60 frames per second is common), and pixel size (the larger the pixel size, the more light that can be captured). Fill factor is the ratio of the pixel's light sensitive area to its total area (higher is generally better).

Most CMOS image sensors use what is referred to as a “rolling shutter” in which images are captured one row at a time. This is generally fine, but for images that contain fast moving objects, it can produce image distortion. As such, many suppliers have introduced “global shutter” technology in which entire images are captured at once.


Another trend has been towards backside illuminated (BSI) sensors. As is the case with most ICs, the photodiodes in traditional CMOS image sensors are below the metal layers that provide connectivity throughout the chip. This causes several issues. Having the photodiodes essentially “at the bottom of a well” reduces the amount of light that reaches it, which is especially an issue in low light environments. In addition, the size of each pixel is reduced since a large portion of the surface area is taken up by the metal connections. Backside illuminated sensors resolve this by putting the photodiodes on top of the “metal wiring,” which can greatly increase fill factor and provide better images. However, the manufacturing is much more complex (it typically involves flipping a traditional front-illuminated sensor and then “thinning” the bottom of the wafer such that what had been the bottom of the photodiodes buried in the wafer can receive light). As such, it was initially very expensive, but it has become increasingly cost effective and commonplace in recent years (the iPhone 4 was the first iPhone to use it).

There is also growing interest in 3D stacking. Once you have the photodiodes on top, more electronic circuits can be added underneath without blocking the image. This can be done by stacking and connecting multiple die together (i.e., the image sensor die is on top with other types of die below it). By doing more processing at the pixel level, a variety of advanced features can be implemented. Many other improvements have been implemented in CMOS image sensors to improve sensitivity and increase dynamic range and image quality. There has also been growing interest in near infrared sensors for automotive applications.

Prior to CMOS image sensors becoming pervasive, the most common type of image sensor was charged coupled devices (CCDs). CCDs are similar to CMOS image sensors, but the underlying technology is different. For many years, CCDs had dramatically better image quality than CMOS image sensors, and the major advantage of CMOS image sensors was low cost and higher integration (integrated amplifiers and data converters). Over the past couple of decades, however, CMOS image sensors have dramatically improved and have become pervasive, although CCDs are still commonly used for a variety of applications. According to MarketsandMarkets, the image sensor market (CMOS and CCD) was \$14.2 billion in 2017 and is projected to reach \$24.8 billion in 2023.


## Solar Cells

Solar cells are similar to optical detectors, in that incoming photons cause electron-hole pairs to form (electrons are freed from atoms resulting in free electrons and holes), which creates a current. For optical detectors, this current is measured in order to determine how much light is incident on the detector. For solar cells, however, the current is used for providing power. The power provided by a single solar cell is small, so solar cells are generally combined together into modules, and modules are combined into panels to generate a larger amount



of power. Because the output of a solar cell is DC, and the electrical grid is AC, an inverter is used to make this conversion in solar panels. A variety of other chips are incorporated to optimize power and ensure safe operation.

The primary parameters for solar cells are efficiency (the percentage of sunlight energy that it can convert into electrical power), cost (as it is a highly competitive market), and durability (as installed solar panels are expected to operate for decades). Many different materials have been used during the past few decades to improve efficiency and reduce cost. Some of these include monocrystalline silicon, multicrystalline silicon, cadmium telluride (CdTe), copper indium gallium diselenide (CIGS), and amorphous thin-film silicon. There have also been a variety of different architectures including multi-junction solar cells (multiple PN junctions rather than one to improve efficiency). However, the bulk of the solar cell market is currently silicon-based (multicrystalline or monocrystalline), as a number of the alternative solar cell technology companies went bankrupt (Solyndra, NanoSolar, etc.).



### Optical Component Companies

The overall optoelectronics market is substantial; generally, more than \$30 billion per year. Some major suppliers of LEDs include Nichia, Osram, Samsung, Philips Lumileds, Cree, Seoul Semiconductor, and LG Innotek. Major CMOS image sensor suppliers include Sony, Samsung, OmniVision, ON Semiconductor, and STMicroelectronics. Select photodiode suppliers include Osram, ON Semiconductor, Vishay, Rohm, and Everlight. Some of the major suppliers of solar cells include Hanwha Q-Cells, JA Solar, Trina Solar, JinkoSolar, Motech, Tongwei Solar, Yingli Green, Canadian Solar, and SunPower.

## Select Optoelectronic Component Companies

The following are some examples of private optoelectronic chip companies:

- **ActLight** – ActLight has developed a dynamic photodetector (DPD), that translates light intensity into a time-delay (rather than a photo-current). This provides a variety of advantages over traditional photodiodes. The company indicates DPDs provide performance comparable to avalanche photodiodes, but are built in CMOS and don't require high voltages (operates at 1V), resulting in low cost and much less power consumption. It also provides a digital output (no ADC required). It announced the use of its DPD in a wearable heart rate monitor (which it indicates has 80% less power consumption than alternatives), and notes that there are many automotive and consumer applications. ActLight announced a Series B capital raise in December 2017 that included investiere.ch, pension fund Nest, and Swisscom Ventures. It is based in Lausanne, Switzerland.
- **Chronocam** – Conventional CMOS image sensors generally capture complete static images at constant intervals (frame rate). This is extremely inefficient in that if an object isn't moving at all, you could capture images of that object at a much slower rate, but if an object is moving quickly you would want an even faster capture rate. Chronocam has developed proprietary CMOS vision/image sensors that are dynamically controlled at the pixel level based on what is happening in the scene, rather than a specific frame rate. It indicates that this results in new sensor architectures that enable high speed (100Kfps equivalent), large dynamic range (>120dB), sensor-level video compression (100x) and power efficiency (<10mW). It also believes that the output from its sensors is better suited for vision processing than traditional image sensor output. While its technology has many applications, ADAS and autonomous vehicle vision systems are a major focus for the company. In October 2016, it completed a \$15 million round that included Intel Capital, Bosch, Renault, iBionext, Robolution Capital, and CEA Investissement. It is based in Paris.

- **Exalos** – Exalos is a leading provider of superluminescent diodes (SLEDs). The company indicates that SLEDs have significant advantages relative to laser diodes and LEDs, in that SLEDs provide laser-like output without the speckle issues of traditional lasers. Historically, Exalos focused on medical and industrial applications, but it recently introduced visible light SLEDs (blue and red, with green expected soon). The company expects its visible light SLEDs will be ideal for pico-projection solutions for AR/VR and many other applications. Exalos also believes its SLEDs have attributes that make it attractive for machine vision applications including automotive sensing. It was founded in 2003 and is based in Switzerland.
- **Fasttree3D** – Fasttree3D offers linear and matrix single photon detectors with integrated distance measurement for time-of-flight 3D cameras, which enable vehicles and machines to recognize and locate fast moving objects in three dimensions in real-time. Its solutions include time digital counters, which can be used for time-of-flight detection in conjunction with VCSEL arrays or LEDs in the near-infrared range, and can operate in darkness or adverse illumination conditions. It indicates that one of its primary focus markets is automotive applications including ADAS and autonomous vehicles. It is based in Switzerland.
- **LeddarTech** – LeddarTech develops LiDAR optical ranging and detection solutions (chips and software), with a major focus on the automotive markets (ADAS and autonomous driving). A solid-state LiDAR solution requires emitting laser pulses, measuring the received signals, and analyzing the signals to extract the key information (distance). LeddarTech has developed advanced algorithms to optimize the sequencing of emitted pulses and the signal processing to extract distance information, and indicates it can enhance sensitivity by 25X compared to alternatives. It also enables the use of much less expensive shorter wavelength lasers for longer distances (e.g., 905nm/940nm versus 1550nm). It offers both chips (LeddarCore) and a LiDAR platform. In late 2017, it announced the LCA2, which it



indicated is the world's first 3D LiDAR IC. The chip generates up to 245,000 digitized waveforms per second and, using proprietary signal processing algorithms, performs over 25 billion operations per second to produce a LiDAR data set for autonomous systems. In 2017, it announced a \$101 million capital raise. Investors include Osram, Delphi, Magneti Marelli, IDT, Fonds de Solidarité FTQ, BDC, and Go Capital. It is based in Quebec City, Canada.

- **Plessey Semiconductor** – Plessey has developed unique GaN-on-silicon technology for high and mid power LED applications. It indicates that its surface emitting GaN technology provides a variety of advantages over traditional GaN-on-sapphire LEDs, including reduced operating temperature, smaller heat sinks, higher drive current, more focused output, better uniformity, and higher lumens/watt. Plessey states that its single-chip high power LEDs outperform traditional four-chip solutions by 30%. Its solution includes a patented buffer layer between the silicon and GaN layer to reduce thermal strain and a mirror layer to reflect light (as silicon tends to absorb light). Plessey offers both LED modules and LED Flexi Strips and provides a family of LED solutions for agriculture (Hyperion), which has less power consumption relative to sodium lights. It also has an ECG sensor unit (EPIC) for medical applications. It is based in Plymouth, England.
- **Pmd Tech** – Pmd's PhotonIC chip is a 3D time of flight CMOS image sensor chip. It provides simultaneous recording of depth maps and 2D grayscale images. It incorporates pmd's Suppression of Background Illumination (SBI) technology and can be used in daylight or night conditions. Leveraging this technology, pmd has also developed middleware software and several complete 3D camera designs (Maxx, Flexx, Monstar). Applications include automotive, robotics, AR/VR, and consumer products. It recently announced wins in the Lighthouse smart home camera and several AR headsets (e.g., Rokid), and that it has a variety of automotive partnerships. It was founded in 2002 and has over 150 patents. Pmd is based in Siegen, Germany.

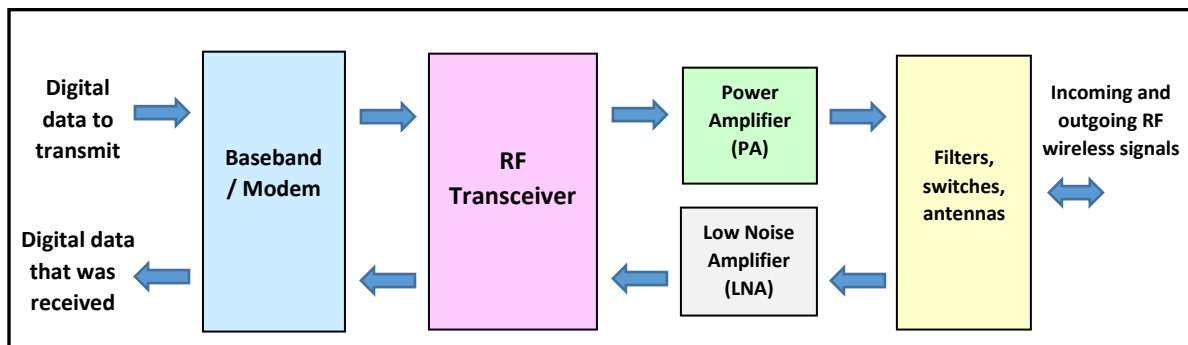
- Sensl** – Sensl develops silicon photomultipliers (SiPMs). Photomultipliers are light detectors that are much more sensitive than traditional photodiodes (e.g., photomultipliers can detect single photons). However, photomultipliers were historically made using legacy vacuum tube technologies. Sensl is the first company to commercialize silicon photomultipliers that can be produced using semiconductor chip manufacturing, which dramatically reduces the cost and size, and greatly improves reliability. Sensl initially targeted medical applications (e.g., PET scanners) and has a number of major medical customers. However, it recently expanded its focus to LiDAR detection and indicates its solutions have substantial advantages relative to traditional LiDAR receivers, including much higher gain and improved sensitivity, enabling better LiDAR performance, shorter laser pulse width, and less power consumption. It is based in Ireland.
- Teramount** – Teramount has developed a unique technology called Photonic-Plug, for connecting optics to silicon. A general issue with optical communications is that aligning a fiber optic cable to an optical chip is difficult, and requires a time consuming active alignment process, which increases costs. Teramount's technology enables rapid passive alignment using standard CMOS assembly lines, increasing throughput and reducing costs, especially for high-speed 100G+ connectivity. It is based in Israel.
- TriLumina** – TriLumina provides advanced VCSEL-based illumination modules for Flash LiDAR applications. TriLumina integrates hundreds of VCSELs into a single chip die, and then uses a patented technology to flip-chip multiple die onto a substrate, with the light emitted out the back. Additionally, TriLumina has technology to directly etch monolithic micro-lenses onto the laser die, which can further enhance the emitted light. This creates a high intensity light output that is very well-suited for LiDAR, while taking advantage of low cost solid-state VCSEL technology and standard industry packaging. TriLumina integrates its VCSEL arrays, along with driver chips and other components, into a complete compact Flash LiDAR illumination

module. It indicates its solution has many advantages including small size, low cost, high resolution, and excellent reliability. TriLumina's solution can also be used for driver monitoring. While TriLumina initially focused on automotive, it believes its VCSEL technology can be used for a broad range of consumer and industrial applications, and plans to address other markets in the future. Strategic investors include Denso and Caterpillar. It announced a \$9 million funding round in 2017, with investors that included Kickstart Seed Fund, Stage 1 Ventures, Cottonwood Technology Fund, Denso, and Sun Mountain Capital. It is based in Albuquerque, New Mexico.

## Wireless Communications Chips

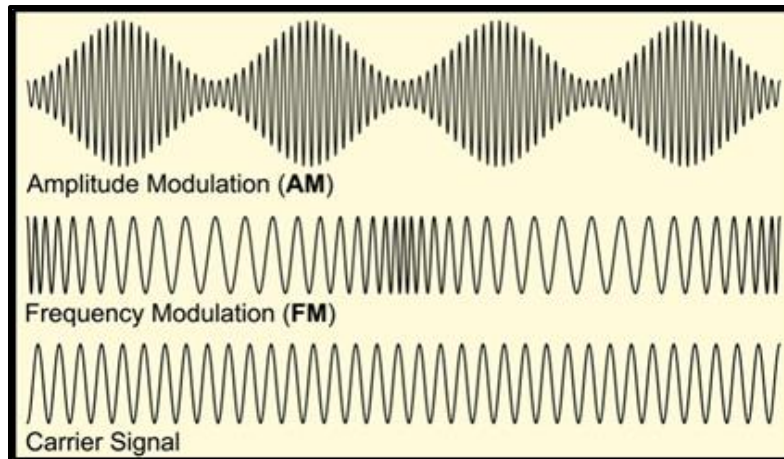
### Intro to Wireless Communications ICs

There are many different types of wireless semiconductor chips. Generically, a wireless solution typically includes a digital baseband/modem chip, an RF transceiver, amplifiers, and miscellaneous filters, switches, and antenna components. On the transmit side, the baseband chip takes the digital data that is to be transmitted and reformats it based on the wireless standard (e.g., LTE, Wi-Fi, Bluetooth, etc.) that is being used. This involves a substantial amount of computations and the baseband chip typically incorporates DSP technology. The resulting digital data must be converted to analog, although this is typically accomplished by data converter technology integrated in the baseband or transceiver chips, rather than using stand-alone data converters.



Source: Menalto Advisors

The resulting analog data signal must then be upconverted to the appropriate transmission frequency (e.g., Wi-Fi typically operates at either 2.4GHz or 5GHz). This is accomplished using some type of modulation scheme which generally involves mixing the analog data signal with a “carrier” signal (pure sine wave) at the desired transmission frequency. One example of a modulation scheme is amplitude modulation (AM) in which the amplitude of the carrier signal is modified based on the data signal (e.g., a high amplitude could represent “1” and a low amplitude “0”). Another example is frequency modulation (FM) in which the frequency of the carrier signal is modified based on the data signal (e.g., a higher frequency could be “1” and lower frequency “0”). These are illustrated in the following diagram.



Most modern wireless systems, however, use more complex modulation schemes. For example, QPSK (Quadrature Phase Shift Keying) is a type of modulation in which only the phase of the signal is used to carry data and there are four possible phases (e.g.,  $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ), each representing two bits of data (“00”, “01”, “10”, “11”). Another modulation scheme is QAM (quadrature amplitude modulation) in which the data signal modifies two carrier signals that are out of phase, and the combination of amplitude and phase changes are used to “encode” the data into the wireless signal. With 16-QAM, for example, there are 16 possible combinations of amplitude and phase and each represents four bits of data (“0000” through “1111”).

Many modern high bandwidth wireless systems (e.g., LTE, Wi-Fi, etc.) also use a technology called OFDM (orthogonal frequency-division multiplexing). With OFDM, the available wireless spectrum is divided up into a number of separate channels, each at slightly different frequencies (e.g., a 200kHz spectrum could be divided up into 20 channels of 10kHz each). Instead of transmitting the data serially (one bit at a time) across the entire spectrum, data is transmitted in parallel across each of the channels. For example, to provide 100Mbps, data could be transmitted at only 10Mbps across each of 10 channels rather than having to transmit 100Mbps across one larger channel. Slowing the transmission rate makes it easier to accurately receive data and helps reduce a variety of technical issues that typically occur. The transmission over each of the individual channels can be done with any one of the previously noted modulation schemes (e.g., QAM, QPSK, etc.).

Before being transmitted, the modulated wireless signal must be amplified so it can reach its destination. The last amplifier is typically referred to as a power amplifier (PA). For applications such as cellular phones, the power amplifier is often made from compound semiconductor materials (e.g., gallium arsenide, indium gallium arsenide, etc.).

On the receive side, the reverse process occurs. The received wireless signal must be amplified. Often the first amplifier is referred to as a low noise amplifier (LNA) since it is optimized for amplifying low power signals while minimizing added noise. The analog data signal and eventually the digital data is extracted from the incoming modulated signal (reversing the modulation process noted above that was used to transmit the signal). The digital bits are then processed by the baseband processor to recover the original data that was transmitted.

In addition to the components noted above (baseband/modem, RF transceiver, amplifiers), there are also a variety of filters (to keep signals only within certain frequency ranges and eliminate the rest), switches (to switch between different frequency bands), antennas, and other RF components.

One interesting type of filter is a surface acoustic wave (SAW) filter in which an electrical signal is converted to an acoustic wave, acoustically filtered based on the dimensions of the cavity it travels across, and then converted back to an electrical signal. A bulk acoustic wave (BAW) filter similarly uses acoustic waves but has a slightly more complex structure and is typically better at higher frequencies. SAW and BAW filters are commonly used in smartphones to help filter the various bands.

The previous few paragraphs were a generic overview of wireless chips, but there are many variations. When new standards are introduced, for example, initial chipsets might include separate RF transmitters and receivers (which eventually are integrated into an RF transceiver). In some cases, the baseband/modem and RF transceiver might be integrated into a single chip (possibly along with the LNA). In smartphones, the baseband modem functionality is often integrated into an SoC with many other functions. There are a variety of other possible combinations.



Semiconductors are used for hundreds of different wireless standards and technologies (Bluetooth, ZigBee, Z-Wave, NFC, etc.). However, just to pick a few, the rest of this section focuses on cellular communications, Wi-Fi, low power wide area networks, and millimeter wave communications.

## Cellular Background

With about 2 billion mobile phones shipping per year, cell phones are an enormous market. The first generation of mobile phones used analog technology (AMPS). The second generation (2G) shifted to digital, with GSM (Global System for Mobile communication). Although GSM was digital, the primary focus was transmitting voice and only very modest amounts of data could be sent (about 9.6kbps). A number of “2.5G” enhancement technologies were later developed to increase data rates including GPRS (General Packet Radio Service) and EDGE (Enhanced Data Rates for GSM). Subsequently, third generation (3G) solutions were rolled out, which were based on a technology called WCDMA (Wideband Code Division Multiple Access). 3G initially provided data rates of up to 384kbps, which was adequate for basic video. Subsequent “3.5G” enhancements such as HSDPA (High Speed Downlink Packet Access) could potentially enable a few megabits of data per second, although actual results varied.

## LTE Overview

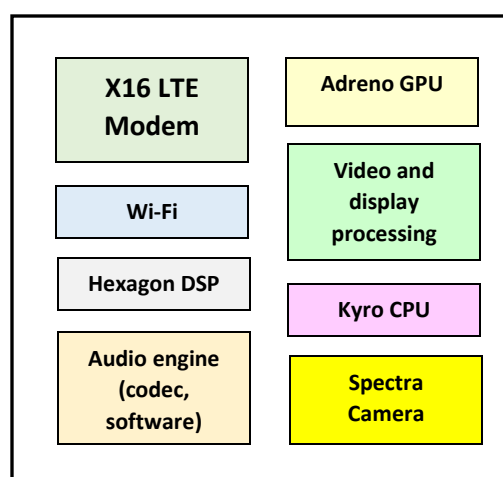
Currently, the industry standard is LTE (Long Term Evolution), which is often referred to as 4G (technically, LTE didn’t really meet the requirements for 4G, but through marketing efforts it was eventually dubbed a “4G” technology). LTE provides much faster rates than 3G technologies. Initially, LTE could provide a theoretical maximum of 300Mbps in downlink and 75Mbps in uplink, but this required 4x4 MIMO (multiple-input and multiple-output). The more typical LTE implementations had significantly lower rates (e.g., one test last year showed typical download rates of about 18 to 36Mbps and upload speeds of about 8 to 16Mbps). There are currently initial deployments of LTE-Advanced, which is often referred to as a 4.5G technology, and it is expected to provide hundreds of Mbps in real deployments. For downlink, LTE uses a version of OFDM (OFDMA) and 64-QAM modulation. For uplink, it uses a different version of OFDM (SC-FDMA) and several possible modulation types (QPSK, 16-QAM, 64-QAM).

LTE provides for two alternative methods for dividing up transmitting and receiving (FDD uses separate frequency bands, TDD uses the same frequency band but transmits and receives at different time intervals). There are actually more than 20 different categories of LTE (including multiple LTE-Advanced categories), each with different uplink rates, downlink rates, and MIMO antenna requirements (up to 8x8).

## LTE Chip Example

As a real-world example, Qualcomm's X16 Modem supports LTE (including both TDD and FDD) as well as legacy standards (WCDMA, GSM/EDGE, etc.). It also supports a broad range of LTE and LTE-Advanced categories including category 16 downlink (about 1Gbps using 256-QAM) and category 13 uplink (about 150Mbps). When introduced, Qualcomm indicated it was the first LTE modem to support gigabit downlink speeds.

Qualcomm subsequently integrated its X16 LTE modem into a Snapdragon 835 SoC chip that incorporates not just the X16 modem, but many other functions including Wi-Fi, a GPU, a camera processor, an audio codec and amplifier, video capture, security, an ARM-based Kryo processor, a "Hexagon" DSP, and many other functions. The Qualcomm WTR5975 is an RF transceiver that is designed to work with the X16 LTE modem. Qualcomm subsequently introduced its X20 LTE modem in early 2017, which supports data rates up to 1.2Gbps downlink (category 18) and 150Mbps uplink (category 13).



Source: Menalto Advisors

## 5G Technology

5G is the next-generation cellular technology beyond LTE and LTE-Advanced. Not surprisingly, there is significant interest in 5G technologies. Many of the major carriers (AT&T, Verizon, etc.) have announced plans to begin rolling out 5G beginning in late 2018/early 2019. However, despite a lot of publicity, and the introduction of a variety of "5G" products, the standards for 5G have still not yet been completely finalized.

Among other things, 5G is expected to provide multi-gigabit per second data rates, and is anticipated to incorporate support for millimeter wave communications to significantly increase bandwidth, but there are still many challenges around this (including that a person's hands can easily block millimeter wave signals). While the 5G details have not yet been completely finalized, that has not stopped companies from announcing a variety of "5G" chips and solutions (some of which have not been 5G, but more marketing spin on LTE enhancements).



As an interim step towards full standalone 5G, the industry has developed an intermediate technology called Non-Standalone (NSA) 5G New Radio (NR). This technology utilizes a conventional LTE network but adds an additional millimeter wave carrier to expand capacity. This is expected to enable commercial deployments by 2019 (possibly late 2018) and provide some technical and commercial feedback before true 5G Standalone New Radio (NR) technologies (which will require substantial new infrastructure) are deployed. However, it will likely be marketed as 5G.

As one example of a 5G chip solution, in October 2017 Qualcomm announced its Snapdragon X50 5G modem chipset which it indicated delivers gigabit speeds and includes a data connection in the 28GHz radio frequency band for Non-Standalone 5G. It also introduced a companion 28GHz transceiver. Qualcomm stated that it is targeting commercial shipments into handsets during the first half of 2019.

### Mobile Phone Wireless Chip Suppliers

The largest supplier of cellular phone communications chips is Qualcomm, which has a substantial share in the baseband/modem and RF transceiver markets and also provides a variety of other smartphone RF components. Other suppliers of these chips include MediaTek, Spreadtrum, and Intel. On the RF side (power amplifiers, LNAs, RF switches, etc.), there are a variety of suppliers including Skyworks, Qorvo, Broadcom/Avago, Murata, and many others.

## Wi-Fi/802.11

The initial wireless LAN products had low data rates and were expensive. However, once the industry was able to develop chipsets that could handle 10Mbps (basic Ethernet speeds) at attractive price points, the technology became pervasive. The IEEE 802.11 standards (typically referred to as “Wi-Fi”) include a broad range of different versions of the technology that have been developed over the years. The first commercialized standard was 802.11b, which could provide up to 11Mbps and operated in the 2.4GHz band. Subsequently 802.11g was introduced which incorporated OFDM technology and improved data rates to over 50Mbps. 802.11n incorporated MIMO technology and provided up to 600Mbps. It also supported dual-bands (at 2.4GHz and 5GHz). The current 802.11ac standard operates at 5GHz and supports up to 3.5Gbps (using 256-QAM and 8x8 MIMO).

The next-generation general Wi-Fi technology is 802.11ax which provides over 10Gbps and operates in both the 2.4GHz and 5GHz bands. Although the 802.11ax standard has not yet been officially approved by the IEEE, companies have already announced 802.11ax solutions (as noted below). The 801.11ax standard incorporates several enhancements including OFDMA technology (which enables a single channel to include data from multiple devices, making it more efficient) and multi-user MIMO.

In addition to these, there has also been developments of 60GHz 802.11 standards. The 802.11ad standard (WiGig) is for very high data rates (about 7Gbps) over short distances (a few meters) at 60GHz. The 802.11ay specification is still in development but is projected to support much higher data rates over longer distances using the 60GHz band.

In the early days of 802.11, solutions often required many chips. However, most of the functionality (baseband, transceiver, etc.) is now usually integrated into a single chip, although many implementations also use separate power amplifiers and LNAs. In many cases, the Wi-Fi functionality is integrated with other functions, such as integrated Wi-Fi/Bluetooth chips.

As an example of a Wi-Fi semiconductor, when Broadcom announced its initial solutions for 802.11ax, it introduced three chips: the BCM43684 (supports up to four streams of 802.11ax for residential Wi-Fi access points), the BCM43694 (similar to the 43684 but for enterprise access points), and

the BCM4375 (supports two streams of 802.11ax Wi-Fi speeds up to 1.429Gbps on mobile devices such as smartphones and tablets, it also integrates Bluetooth 5.0+). Some of the major suppliers of 802.11 chips include Broadcom, Qualcomm, MediaTek, Intel, Marvell, and Realtek.

### Low Power Wide Area Networks (LPWAN)

Cellular and Wi-Fi generally focus on high data rate transmissions. There are a variety of wireless standards for transmitting at lower data rates over shorter distances (e.g., Bluetooth, NFC, ZigBee, etc.). However, with growing interest in IoT, there is now interest in low power wide area network (LPWAN) wireless technologies. LPWAN is focused on transmitting small amounts of data (<5kbps and often much less) over long distances (several miles) with very low power consumption. It targets applications such as meter reading, asset tracking, agricultural monitoring, and certain battery-powered sensor networks. Existing wireless standards generally are not well suited for these applications due to power consumption and/or distance requirements. As such, many new wireless technologies have been developed to address these applications. A couple of examples (LoRaWAN and SigFox) are highlighted below, although there are many others including new Wi-Fi standards (802.11ah/HaLow) and versions of LTE/cellular.

LoRaWAN is supported by the LoRa Alliance (sponsor members include Semtech, IBM, Alibaba, Cisco, Renesas, SK Telecom, STMicroelectronics, and ZTE). LoRaWAN's data rates range from 0.3kbps to 50kbps, and it is expected to have a range of over 15km. It generally operates at about 868MHz in Europe and 915MHz in the U.S. It includes adaptive data rate technology to adjust data rates based on range and battery needs. Semtech originally developed the underlying chip technology, although other chip companies (STMicro, Renesas, etc.) are also expected to provide solutions.

SigFox also uses the 868 and 915MHz bands. It transmits very small amounts of data (12 bytes) over very narrowband channels using BPSK (similar to QPSK but only two phases are used). Because it transmits data very slowly over a very narrow band, the electronics are relatively simple with low power consumption. SigFox has licensed its technology to several chip companies that sell SigFox radio chips. Some of these include Microchip, STMicroelectronics, NXP, ON Semiconductor, TI, and Silicon Labs.



## Millimeter Wave Integrated Circuits (MMICs)

Traditional cellular, Wi-Fi, and Bluetooth solutions operate in the low GHz range (e.g., generally about 2 to 5GHz). As this spectrum has become increasingly crowded, there has been growing interest in using higher frequencies (millimeter wave) with more available bandwidth for wireless communications and other applications. Although there isn't a strict definition of millimeter wave, it typically refers to transmissions in the 30GHz to 300GHz range (from 1mm to 10mm wavelengths). Integrated circuits that operate in this range are often referred to as MMICs (millimeter wave ICs).

MMICs can provide a variety of advantages. Because of the large amount of bandwidth available at higher frequencies, it is possible to transmit at high data rates and high capacity even though communications at those higher frequencies is more challenging (e.g., you can usually get more cars through on a 20-lane highway than on a 2-lane highway, even if the 20-lane highway has many potholes).

In addition, there are a variety of unlicensed millimeter wave frequencies (e.g., 60GHz) and for the frequencies that are licensed, spectrum licensing costs are generally much less. Another advantage is that millimeter wave links can cast narrow beams, which enables greater frequency re-use (the same frequency can be reused in relatively nearby locations without interference, which can increase capacity). Also, at higher frequencies, wavelengths are smaller, and so antenna size can often be reduced.

While there are many potential advantages, millimeter wave communications has major challenges. At these frequencies, the atmosphere (oxygen, water vapor, etc.) can significantly attenuate wireless signals, which makes transmission more difficult. Certain millimeter wave frequencies are much better than others for communications, but there can be a variety of factors. For example, transmissions in the 70-80GHz range are generally not affected by atmospheric oxygen as much as other bands, but tend to have significant losses if there is rain. In addition, humans or other objects can also significantly attenuate or block signals at millimeter wave frequencies (this is a major challenge for millimeter wave in handsets).



In general, millimeter wave communication is much more challenging than communications at lower frequencies, and historically many of the commercial millimeter wave deployments were for fixed “line of sight” communications links with no objects between the two endpoints. During the past few years, however, there has been growing interest in MMICs:

- **5G** – The 5G cellular standard is expected to include both sub-6GHz frequencies as well as millimeter wave frequencies. This is expected to dramatically expand the market for millimeter wave chips.
- **Wireless Backhaul/Small Cell 5G** – As cellular data rates to handsets grow, there is also growing demand for higher capacity wireless backhaul and small cell technologies. Wireless backhaul often historically used “C-Band” technology (about 6GHz uplink and 4GHz downlink), but there is greater interest in backhaul at 60GHz and at even higher E-Band (71-76 and 81-86GHz) frequencies.
- **Automotive** – Radar is critical for many automotive ADAS and autonomous vehicle features, and it uses millimeter wave frequencies (historically around 24GHz or 35GHz, with growing interest in 76-81GHz). Vehicle to vehicle (V2V) communications is expected to operate around 5.9GHz, but there has been interest in using higher frequencies instead for more bandwidth.
- **Consumer** – There is interest in using 60GHz (such as WiGig/802.11ad) for a variety of consumer applications such as transmission of video within a room and for applications such as AR/VR headsets.
- **Data Center** – There is great interest in transmitting large amounts of data rapidly within data centers over short distances (e.g., server to server). One approach has been to use massively parallel millimeter wave connectivity to achieve very high capacity without cabling.
- **Others** – As previously noted, the upcoming 802.11ay specification (successor to WiGig) is expected to provide much higher data rates over longer distances at 60GHz. This could further expand the potential market for MMICs.

There are a broad range of suppliers of millimeter wave chips for various applications. Most of the major RF, compound semiconductor, and analog chip companies sell some types of MMIC solutions. One general trend is that initially high frequency solutions typically use compound semiconductor materials, but over time CMOS technologies can improve to the point where they can displace compound semiconductor chips (lower cost, greater integration). As such, there appears to be growing interest in CMOS millimeter wave components.

### Wireless for Sensing and Location; UWB

While the vast majority of wireless chips are for communications, there are also applications of wireless technologies for sensing. Ultra-wideband (UWB) is a wireless technology in which signals are transmitted over an extremely large frequency range, but in which the energy at any given frequency is very small, so as to not interfere with other signals. At one point, expectations were that it would be used for short distance communications (e.g., the WiMedia standard for PC accessory connectivity), but a number of companies have developed UWB solutions for indoor location tracking and motion detection. In general, there are a growing number of companies using various wireless technologies for sensing applications, including a few that are highlighted below.

### Select Private Wireless Semiconductor Companies

The following are a few examples of private chip companies focused on wireless communications and technologies.

- **Acco** – Acco develops CMOS power amplifiers and RF front-ends for cellular standards. It targets smartphones and IoT devices. In late 2015, it announced production of its 26120 power amplifier, which supports quad-band GSM/EDGE and 12-band 3G/LTE, and uses a low cost CMOS process, rather than gallium arsenide (which is more typical). In early 2017, it announced an integrated RF front-end module for IoT applications. In November 2016, it announced a \$35 million capital raise. Investors include Bpifrance, Foundation Capital, Pond Ventures, Partech Ventures, Omnes Capital, Siparex Group, and A Plus Finance. It has major offices in Paris and Silicon Valley.

- **Alereon** – Alereon develops ultra-wideband (UWB) chips. Its solutions include a digital MAC/Baseband and an RF transceiver. Its chips can transmit over a broad frequency range (3.1-9.5GHz) and support the WiMedia UWB standard. A decade or so ago, WiMedia was expected to become pervasive in consumer devices, but Alereon has more recently focused on military and medical applications. In medical, UWB can provide higher reliability than Wi-Fi or Bluetooth as there is often a lot of interference and wireless traffic at the frequencies those standards operate at (2.4GHz or 5GHz). For military, there are numerous applications including mobile soldiers (eliminating many cables), soldier body-area networks (connectivity among smartphones, head mounted displays, audio headsets, GPS, and other electronics), and in-vehicle video. UWB is difficult to detect as it operates across many frequencies, making it well-suited for soldiers. The company is based in Austin, Texas.
- **Amimon** – Amimon has developed wireless communications chips that are specifically optimized for transmitting video, in contrast to most wireless standards which are designed to transmit any type of data. By focusing on video, Amimon is able to transmit robust high definition video with much higher quality and less latency than alternatives. One element of its solution is that it uses unequal error protection in which the more critical bits have a higher level of protection than other bits, which provides better channel utilization and is more immune to interference. Leveraging this core technology, it has developed its 3<sup>rd</sup> generation chipset as well as a variety of modules and products for applications such as broadcast, drones, and medical, and it recently introduced a solution targeting virtual reality applications. It is based in Silicon Valley with a major office in Israel.
- **Arbe Robotics** – Arbe Robotics has developed advanced imaging radar solutions for autonomous vehicles. It indicates that its solutions provide real-time “4D mapping” and have many compelling advantages over conventional radar including much higher resolution (to better identify objects). Its solutions utilize millimeter wave technology and

offer high resolution at long or short ranges, in both azimuth and elevation, and create a point-cloud format. Its SLAM platform creates full 3D shapes of the objects, tracks them, and classifies them on a map at 25 times per second. Arbe is initially using chips from other suppliers with its advanced software, but plans to introduce its own highly integrated chip solution, which it expects will provide compelling advantages. According to an article in TechCrunch, Arbe is in trials with five major auto-related companies. In November 2017, it raised \$9 million with investors that included O.G. Tech Ventures and OurCrowd. It is based in Israel.

- **Arralis** – Arralis develops high frequency millimeter wave ICs (MMICs), up to 110GHz. Its Leonis Ka band chipset includes the chips required to build a satellite and ground front-end including power amplifiers, low noise amplifiers, clock/VCO, and mixers (generally from 17GHz to 31GHz). Its Tucana family operates around the 94GHz range and includes power amplifiers and other amplifiers, switches, mixers, and RF components. Arralis also provides module level solutions that incorporate its chips, and it has developed proprietary antenna technology for its solutions. Applications include radar, space/aerospace, and communications. It is based in Ireland.
- **Autotalks** – Autotalks is a semiconductor company that develops solutions for V2X (vehicle to everything) communications for cars. Its Craton 2 solution is a dual-ARM core based complete V2X solution that includes a standalone V2X ECU (engine control units) and targets the retrofit and aftermarket. It incorporates an 802.11p modem, and also supports 802.11a/b/g/n/ac for external Wi-Fi connections. Its Sector solution is a secure V2X hardware add-on solution (it integrates with external host CPUs), designed for ECUs that have available CPU processing, and also supports 802.11p and 802.11a/b/g/n/ac. The company has also developed a variety of technologies (e.g., Truly Secure) to help ensure that V2X communications is secure. In 2016, Autotalks announced that Denso was using its chips for several major V2X projects. In March 2017, Autotalks announced a \$30 million

Series D round and in June 2017 it was reported that the round was expanded to \$40 million. Investors include Magma, Gemini, Amity, Mitsui, Liberty Media, Delek Motors, Fraser McCombs, Vintage Investment, Samsung, and Mirai Creation Fund (Toyota and Sumitomo Mitsui). It is based in Israel.

- **Celero** – Celero provides a variety of Wi-Fi chip solutions. Its CL2440 is a 4x4 802.11ac wave 2 single chip solution, with PHYs and MACs that are optimized for dense environments. It includes a built-in “Argus” DSP-based spectrum analyzer and Celero’s OptimizAIR 2.0 suite, which supports beamforming and other advanced features. It also offers 3x3 and 2x2 versions as well as 802.11n solutions. In 2017, it announced that its chips would be used in several Samsung products including a 4K set-top box and a DOCSIS 3.1 gateway. It recently announced its Everest 802.11ax Wi-Fi solution (1024-QAM, up to 4.8Gbps, 14nm FinFET process), and also disclosed a partnership with Realtek for a G.fast/802.11ac gateway. Investors include 83 North, Cisco Systems, Liberty Global, Miven, Our Crowd, Poalim, Pitango, Red Dot, and Vintage Investment Partners. It is based in Israel.
- **Decawave** – Decawave develops chips for indoor location and communication utilizing ultra-wideband (UWB) technology. Leveraging its core impulse-radio UWB (IR-UWB) technology, Decawave developed the DW1000 CMOS chip and associated ScenSor modules, which can be used for very accurate indoor positioning (where GPS generally doesn’t work). The DW1000 operates at data rates of 110kbps, 850kbps, and 6.8Mbps, and can locate tagged objects both indoors and outdoors to within 10cm. Applications range from precise indoor location of people/smartphones for retail, to agriculture, industrial, and automotive. In July 2016, it announced it shipped its one millionth chip and that it had 1,800 customers. In April 2017, STMicroelectronics reportedly invested 2.5 million euros in Decawave. Decawave is based in Ireland.

- **Escape Communications** – Escape provides advanced millimeter wave communication modem solutions for terrestrial (backhaul, point-to-point), satellite, and military communications. Its designs are implemented in FPGAs and it provides complete module solutions. An example is the ESM5008, which is a 2.7Gbps modem for E-band/V-band communications. In 2016, Escape introduced a low-cost 10Gbps E-band (71-76GHz, 81-86GHz) modem for backhaul applications, including 5G backhaul. It also announced partnerships with Analog Devices and Filtronic (including demos of a 10Gbps software-defined E-band backhaul product), and began shipping 10G modem modules in 2017. While its solution is implemented as an FPGA, with higher volumes Escape can transition the design to an ASIC solution. Escape also sells a variety of protected military communications products (modems and signal processing for AEHF ground terminals). It is based in Torrance, California.
- **HMico** – HMico has developed a unique wireless solution specifically optimized for medical sensor applications and patches, which can help eliminate the need for the wires/cables associated with patient monitoring. As conventional wireless standards (Wi-Fi, Bluetooth) are not reliable enough for patient monitoring, HMico developed a solution that integrates three separate wireless technologies (Wi-Fi, a medical-band transceiver, UWB technology) to provide ultra-high reliability. Its “Life Signal Processor (LSP)” integrates three radios, along with power management, sensor interfaces, and an applications processor for a complete medical interface solution. It can intelligently switch between radios for low-power consumption and high link reliability. HMico has subsequently expanded beyond just chips to provide biosensor patch designs, receivers, and software. It announced a chip partnership with STMicroelectronics. In late 2016, it announced a \$10.2 million capital raise from strategic investors, as well as from XSeed Capital, Seraph Group, and Uniquist. It is based in the Silicon Valley area.



- Keyssa** – Keyssa has developed technology that enables very high data rate wireless connections over very short distances, with a focus on replacing connectors in devices. As electronics shrink, traditional connectors are increasingly a bottleneck in terms of size and speed. Keyssa’s “Kiss” technology provides up to 6Gbps at extremely high frequencies. The solution is ultra-compact (the KSS104M is 3 x 3mm), and does not have to be on the edge of the device, freeing up surface space. Kiss uses a combination of semiconductor technology and advanced mechanical design. It can be used for internal communications and/or external connectivity, replacing internal and external cables. It supports many protocols (USB 3.0, SATA, PCIe), eliminating the need for separate drivers. Announced wins include Acer’s Aspire Switch 12 S laptop/tablet. Investors include Alsop Louie, Dolby Family Ventures, Foxconn, SK Hynix, Intel, Nantworks, Neuberger Berman, and Samsung. It is based in Silicon Valley.
- Lime** – Lime Microsystems develops programmable RF transceiver chips. It introduced a multiband LTE transceiver in 2011, and now supplies a variety of programmable transceivers including the 7002M (supports a 100kHz – 3.8GHz frequency range, includes an integrated microcontroller). It has announced the 8001+ (which operates from 100kHz to 12GHz). LimeSDR boards utilize both FPGAs and Lime’s programmable transceivers to create software defined radios. In mid-2017, Lime announced that it will collaborate with Vodafone to develop software-defined cellular radio platforms that support Vodafone’s OPEN RAN vision. Investors include ACT Venture Capital, Intel, Draper Esprit, In-Q-Tel, and Parkwalk. Lime is based in Surrey, UK.
- Maja Systems** – Maja Systems develops highly integrated CMOS millimeter wave digital radio solutions. Maja is able to integrate many analog/RF functions (mixers, A/D converters, D/A converters, LDO, VCXO, digital control, modem, etc.) into a single chip that can be produced using conventional digital CMOS manufacturing. It also provides complete modules that incorporate a CMOS antenna, which it also designed. Maja’s chips are highly configurable and include

patented built in self-test features that enable wafer level testing. Maja believes it can drive lower cost, smaller size, lower latency, less power consumption, and better reliability. Maja's initial solution is a 60GHz CMOS radio, but its roadmap includes MMICs optimized for 5G, IoT, automotive, and other markets. It is based in Silicon Valley.

- **Octasic** – Octasic develops communication chips and systems. Its OCT2224W is a System-on-Chip (SoC) device optimized for GSM, HSPA, and LTE baseband processing. It consumes less than 5 watts while serving 64 users within a range of 20km. Leveraging this core technology, Octasic has developed its OCTBTS family of mini base stations which support a broad range of frequencies and standards for applications such as public safety, military, industrial IoT, and other markets. Separately, Octasic has a family of multi-core DSP-based media gateway chips and software (Vocallo) for media processing (transcoding, conferencing, echo cancellation). Founded in 1998, Octasic is a private company based in Montreal, Canada,
- **Redpine** – Redpine offers a variety of wireless communications chips with a focus on Wi-Fi and microcontrollers that integrate Wi-Fi. This includes SoCs and multi-protocol chips that support Wi-Fi, Bluetooth, and Zigbee. It also has a family of ARM Cortex-M based microcontrollers (WiSeMCU) that integrate Wi-Fi and other wireless technologies. In addition, it offers highly secure FIPS-compliant Wi-Fi solutions; multi-protocol 802.11p chips for the emerging automotive V2X communications market (WAVEcombo); and solutions for wireless asset tracking. It recently launched an ultra-low power wireless microcontroller for IoT devices. It is based in Silicon Valley.
- **Telink** – Telink develops low power radio-frequency and mixed-signal system chips. Its products include low-power 2.4GHz RF SoCs for Bluetooth Smart (Low Energy), Zigbee, and other wireless standards. It also provides touch controllers. Telink targets IoT applications such as smart lighting, home automation, and the smart city. It is based in Shanghai, China.

- **Vayyar** – Vayyar has developed a unique CMOS imaging chip that uses RF over a variety of frequency bands to create 3D images of objects without the use of cameras. The technology can even capture images through walls or other objects. The core technology was originally developed for breast cancer detection, but has been modified to address many markets. Police/fire departments can determine where people are in a house, contractors can identify locations of cables or pipes behind walls, and there are numerous medical and industrial applications. In addition, as the technology can identify if and where people are in a room or car without a camera, it is well-suited for monitoring without privacy concerns (e.g., turning a TV to low resolution if no one is watching, determining how many people are in a car, eldercare in retirement communities, etc.). It can also be used for gesture recognition and for many automotive applications. Vayyar announced a partnership with Softbank for the IoT market in late 2017. In December 2017, Vayyar announced a \$45 million Series C round that included Walden Riverwood, Battery Ventures, Bessemer, ITI, Claltech, Israel Cleantech Ventures, and Amiti. It is based in Israel.
- **Vectrawave** – Vectrawave develops millimeter wave ICs (MMICs) and RF chips and modules for a broad range of applications including wireless, satellite, and optical communications. It provides chips that operate at frequencies up to 110GHz. Its chips include low noise amplifiers, drivers, phase shifters, clocks, and other analog components. It also provides multi-chip module solutions. Investors include Sigma Gestion, Nestadio, and Starquest. It is based in France.

- **XeThru** – XeThru has developed ultra-wideband (UWB) impulse radar technology using extremely low energy gigahertz radio waves. Leveraging this core technology, it has created the X4 SoC that provides sub-millimeter movement sensing accuracy at distances up to 25 meters depending on target size, and can even see through obstacles. It also utilizes this technology for various sensing applications including a presence sensor (which can detect human presence and determine the distance to a person or other object at distances of more than 9 meters) and a respiration sensor (for sleep, health, and baby monitoring). In late 2016, it announced a \$12 million capital raise with lead investor Investinor. It is based in Oslo, Norway.

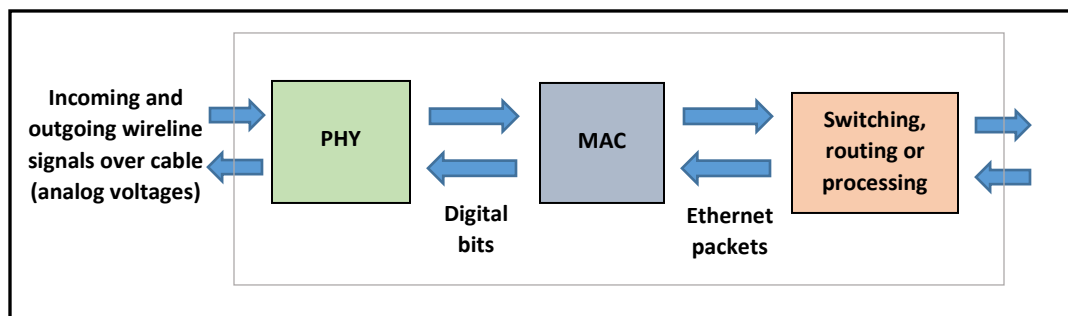
## Wireline Communications Chips

### Intro to Wireline ICs

Wireline communications involves the transmission of data across wires or cables. There are dozens of wireline standards and technologies, although generically the transmission and reception involves digital functionality (that formats the digital data per the particular standard) and analog/mixed-signal functionality (PHY) that is involved in transmitting the electrical data over the specific medium used.

### Ethernet Chip Overview

As an example, Ethernet is the most common communication protocol used within enterprises for networking, and is also increasingly used in telecommunications. When data is to be transmitted using Ethernet, it must be formatted into Ethernet packets which have a specific format (i.e., a preamble of alternating “1”s and “0”s, the destination address of the packet, the source address, the data itself, and error correction data). The resulting digital data must then be converted to appropriate electrical signals that can be transmitted across the networking cable. On the receiving end, the reverse process occurs. From the incoming electrical signals, the digital bits that were transmitted are recovered. These bits are then reassembled into Ethernet packets, which can then be analyzed and processed.



Source: Menalto Advisors

The analog/mixed-signal chip that is involved in transmitting and receiving signals over the physical medium is typically referred to as a PHY. On the receive side, the PHY typically includes a function known as clock and data recovery (CDR) to recover the digital bits. Some chips incorporate multiple PHYs on a chip (e.g., an octal PHY with eight PHYs on a single chip).

Historically the digital Ethernet functionality for controlling transmission and assembling packets was performed by a chip referred to as a MAC (media access control). However, in most cases this digital functionality is now generally integrated into larger chips, performed by some type of processor, or combined with the PHY. For example, many FPGAs incorporate Ethernet MACs, eliminating the need for a stand-alone MAC chip. In other types of communication standards, the digital chip functionality has other names (such as a framer).

Communication equipment (switches, routers) typically have a variety of switching and packet processing chips that can analyze the data packets (e.g., read the destination address in the Ethernet packets) and determine where to forward the packets next to bring it closer to its final destination. As an example, Broadcom's StrataXGS BCM56980 is a highly integrated switching chip that can switch up to 32 ports of 400Gb Ethernet packets and includes 56G SerDes for high-performance networking products.

## Optical Communications

The previous few paragraphs focused on electrical communications, but transmitting at high speed over copper is challenging due to attenuation issues (signals degrade after only short distances), electrical interference, and other technical limitations. Transmitting data via light over fiber optic cable has a variety of advantages and can generally support much higher data rates and significantly longer distance transmissions. However, optical communication is typically much more expensive and requires more components. It is therefore used only when very high data rates are required and/or when data must be transmitted over long distances.

For example, Ethernet with speeds of 1Gbps and below almost exclusively use copper (all electrical). Ethernet at 10Gbps was initially exclusively optical, but a version using copper (10GBASE-T) has become increasingly popular. The vast majority of 40G Ethernet solutions use optical communications although there is a 40GBASE-T standard for 40G over copper for distances under 30 meters. Although 100G is primarily optical, there are 40G and 100G standards for transmission over copper for very short distances (a few meters for transmission over backplanes).




With optical communications, the data that was carried electrically must be encoded in a light signal, which is then transmitted over a fiber optical cable. At the receiving end, the data carried by the light signal must be extracted as an electrical signal which can then be digitally processed.

On the transmit side, a laser is used to generate light at a specific wavelength (lasers were previously discussed in the Optoelectronics section). When an appropriate voltage is applied, semiconductor lasers turn on and emit a beam of light. However, the light must be altered in some way to encode data. This is done using modulation schemes, which are very similar to the previously discussed modulation schemes used for wireless signals in which the data signal modifies a basic carrier wave. With optical, the data signal modulates the light signal such that the amplitude, frequency, phase and/or polarization of the light changes in such a way to “encode” the data in the light signal (which can later be recovered on the receiving end).

There are a variety of ways to modulate the light. As a simple example, by rapidly increasing or decreasing the current to the laser, the intensity of the output can be changed, which can be used for amplitude modulation of the light, which could encode data. When modulation occurs by directly modifying the laser output, it is referred to as directly modulated lasers (DML). At very high speeds, however, direct modulation is very difficult so typically a separate modulator component is used after the laser. That is, the laser emits a plain vanilla light signal at a specific wavelength and then the modulator modifies the light signal (the amplitude, frequency, phase or polarization) based on the data to be transmitted.

There are a variety of different types of modulators. One popular type, for example, is a Mach-Zehnder modulator (typically made from materials such as lithium niobate or indium phosphide). The Mach-Zehnder modulator splits the laser light signal into two paths. By applying voltages to the paths, the phases of the two light signals can be modified. The light signals from the two paths are then recombined (as shown in the next diagram). The difference in phases between the two light signals affects the amplitude (e.g., if the two light signals are 180° out of phase the amplitude is zero, if the two phases are the same the amplitude is maximized) and so the amplitude and phase can be controlled. This enables a variety of modulation schemes.



Source: Menalto Advisors

Another popular type is an electro-absorption modulator (EAM) in which the output intensity can be modified by changing a voltage which generates an electric field that changes the amount of light that is absorbed (the more light that is absorbed, the lower the output intensity).

There are a variety of different modulation schemes used for optical communications. The simplest is on-off key (OOK) in which high intensity light represents a “1” and low intensity is a “0.” A version of this called non-return to zero (NRZ) is commonly used in many optical systems.

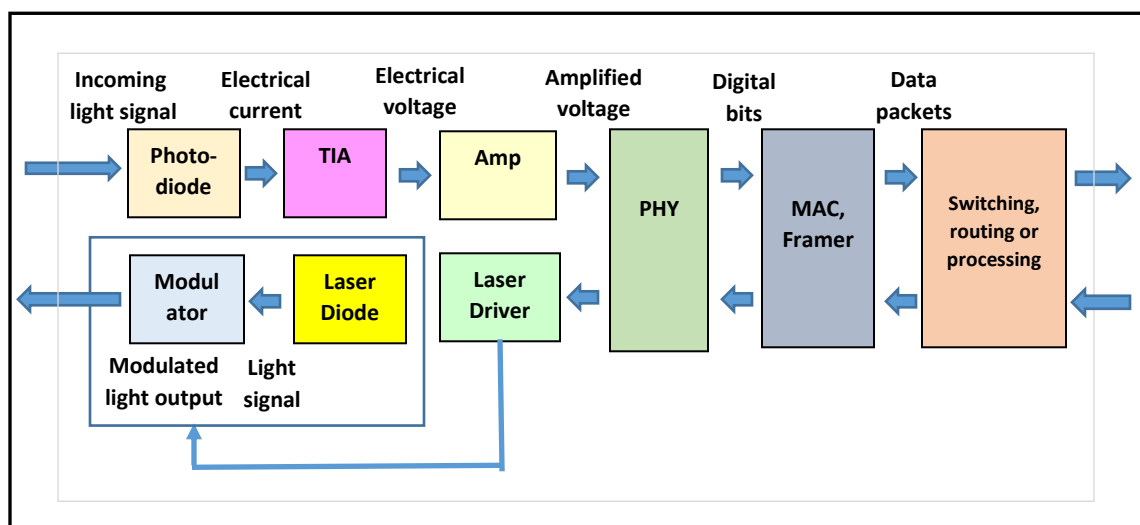
This type of simple amplitude modulation (high or low intensity) was adequate for lower data rates, but at higher rates (generally above 10Gbps) it is problematic due to a number of technical issues. For example, although light travels at a constant speed in free space regardless of its wavelength, this is not the case in a fiber optic cable, as different wavelengths of light travel at slightly different velocities in the fiber. This causes light pulses to spread out over time (chromatic dispersion), which makes it more difficult to recover the data. This, along with other types of dispersion (polarization mode dispersion, etc.), limits the data rate of basic OOK modulation. To increase data rates, more complex types of modulation are often used. Many of these are similar to those already discussed in the Wireless section.

For example, with QPSK (quadrature phase-shift keying) a constant amplitude of light is used, but the phase of light is shifted to one of four phases ( $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ) and each phase represents two bits (00, 01, 10, 11), so light modulated at 25Gbps can provide 50Gbps instead of just 25Gbps. QAM (quadrature amplitude modulation) uses both amplitude and phase, and multiple bits can be provided at a time (e.g., 16-QAM provides

four bits at a time based on 16 different combinations of amplitude and phase of the transmitted light). PAM (Pulse-Amplitude Modulation) uses more than two possible amplitude levels to provide multiple bits (e.g., PAM-4 has four amplitude levels with each level representing two bits). In some schemes the polarization of the light is also used. For example, DP-QPSK (dual polarization QPSK) is similar to QPSK, but also uses the polarization of the light signal in addition to phase, providing more bits at a time.

While these modulation schemes are more efficient and effective at higher data rates, they require complex mathematical operations to analyze and extract the data. This has led to “Coherent DSP” architectures in which immediately after the incoming signal is converted from optical to electrical (analog) to digital bits, a digital signal processor (DSP) is used in the PHY to extract the data and can help compensate for transmission issues.

The following illustrates some key components in a generic optical communications system. Incoming light is converted to an electrical current by a photodiode. This current is transformed to a voltage by a transimpedance amplifier (TIA) and then amplified. The PHY extracts the digital bits from the signal, and then a digital chip creates data packets which can then be analyzed and processed. For transmission, the reverse occurs except a laser driver is used to drive the laser/modulator to modify the outgoing light signal. In some cases, several of these functions may be integrated and there are a variety of variations.



Source: Menalto Advisors

## 100G Ethernet and Beyond

The initial Ethernet standard was for 10Mbps over coaxial cable, but it became pervasive across PCs with 10Mbps over twisted copper cables. Subsequently, the industry shifted to 100Mbps, and then 1Gbps. Networking equipment and servers typically support at least 10Gbps Ethernet and there has been strong demand during the past few years for 40Gbps and increasingly 100Gbps Ethernet.

One of the major changes in the high-speed wireline communications markets is that historically the highest speed communications chips were used almost exclusively for the core of the telecommunications network. However, with the shift to the cloud, and large data centers handling enormous amounts of data that often must be processed in real-time, there has been growing interest in utilizing the highest speed chips in data centers.

The initial 100G solutions were 10x10G (combined 10 parallel fibers each transmitting 10Gbps) but later 4x25G solutions were introduced, and there has more recently been development of 1x100G solutions (100G over a single optical fiber). Some of these solutions use CWDM (Coarse Wave Division Multiplexing) in which several light signals at slightly different wavelengths are transmitted over the same fiber. For example, instead of having four parallel fibers carrying 25G signals, all four signals can be transmitted over the same fiber, which reduces the number of fibers but requires combining and multiplexing the optical signals onto the same fiber.

There are dozens of different 100G variations based on the type of fiber used, the type of connector/interface plug used, and the maximum distance supported (e.g., 100 meters, 1 kilometer, 40 kilometers, etc.). There have also recently been standards developed for very short distance 100G Ethernet over copper for backplane and chip-to-chip communications. In addition, some recent standards incorporate support for forward error correction (FEC) technology to improve performance and reduce errors.

To drive down the cost of 100G, there has been growing interest in shifting from 4x25G (which requires four sets of lasers, modulators, etc.) to 2x50G and ultimately to 1x100G. The 1x100G technology is often referred to as “single lambda” or “true 100G”. It is believed that, in volume, this can help significantly reduce 100G costs.

As one example of a 100G single lambda chip solution, last year Semtech announced (in conjunction with MultiPhy, discussed later) a single lambda FiberEdge 100G platform that uses a 56G laser modulator and PAM-4 modulation (two bits per symbol), which provides over 100Gbps.

Beyond 100G, there are already a variety of 200G and 400G developments. The IEEE just approved specifications for 200G and 400G Ethernet in December 2017. These are similar to 100G standards but use parallel lanes to achieve higher data rates (e.g., 8x56G, 4x112G). It is expected that multilane 100G will be capable of 800G and eventually 1.6T speeds, but that multi-terabit speeds will require additional major technology breakthroughs.

## SerDes

SerDes is short for serializer/deserializer. In some cases, data is transmitted in parallel while in other cases it is transmitted serially. For example, the output of a chip might be eight parallel wires each processing 1.25Gbps (for a total of 10Gbps). However, the communications link on the backplane of the equipment might process 10Gbps as a single serial data link. As such, the data must be converted from one format to the other. A SerDes performs this function. In some cases, a SerDes is a stand-alone chip, while in other cases it is integrated into a PHY or a larger communications IC. As with all communications chips, there is growing interest in higher speed SerDes.

## OTN

Historically, high speed optical communications within the telecommunications network was implemented using a technology known as SONET, with data rates represented by an optical carrier number (e.g., OC-192 provides 10Gbps, OC-768 provides 40Gbps). However, there has also been growing interest in using Ethernet (Carrier Ethernet) and other protocols in telecommunications networks. OTN (Optical Transport Network) is a high speed optical protocol for transmitting any type of data over an optical network, and there has been growing interest in the technology during the past few years.

OTN is often called a “digital wrapper” as it encapsulates any type of data (Ethernet, HDMI, Fibre Channel, etc.) into an optical data unit (ODU) which can be efficiently transmitted across an optical network. At the receiving end, the “wrapper” is removed and the original packet format can be processed. The ODU also integrates advanced forward error correction (FEC) technology to reduce errors and improve performance. While the focus for high speed communications has generally turned to the data center, OTN appears to be the “hottest” segment of the traditional telecommunications chip market with a variety of OTN chip solutions. As one example, last year Microsemi introduced its DIGI-G4 OTN Processor which is a single-chip 4x100G solution for OTN line cards and provides multi-service support for Ethernet, storage, IP, and SONET over OTN and supports 400G switching.

### **Automotive Wireline Communications**

For many years, cars had point-to-point cable connections. However, as the amount of electronics increased, the industry shifted towards networking, most notably the CAN (Controller Area Network) bus. Subsequently additional automotive networking solutions were developed including FlexRay, LIN, MOST, and others (details about these networking standards are discussed in our previous report, Emerging Automotive Technologies). However, with the increase of complex automobile electronics, these legacy networking technologies have become a major bottleneck.

As such, there is growing interest in higher bandwidth car networking. One of these developments is Ethernet for cars. Broadcom developed an automotive Ethernet chip solution (BroadR-Reach) which was the first commercial Ethernet solution to meet the EMI and other requirements for automotive applications. The Open Alliance was formed to support this solution (often called “One-Pair Ether-Net”) and other chip companies began also producing BroadR-Reach chips, which provide 100Mbps over unshielded twisted pair cable. The IEEE has more recently published standards for Gigabit Ethernet over both unshielded twisted pair copper cable and plastic optical fiber. Work has already begun on a multi-gigabit automotive Ethernet standard. Another recent automotive standard is HDBaseT Automotive which provides in-vehicle connectivity for HD video, audio, Ethernet, USB, power, and other signals over a single cable. It can provide up to 6Gbps over unshielded twisted pair cable.



## Many Other Wireline Standards

Most of this section focused on Ethernet, and we briefly touched on OTN and automotive Ethernet, but there are dozens of other wireline communications standards (HDMI, DSL/G.fast, PCI Express, SONET, USB, DisplayPort, Fibre Channel, powerline, GPON, DOCSIS, Thunderbolt, etc.).

## Some Major Wireline Chip Suppliers

There are dozens of chip companies that provide wireline communications components. A few examples of semiconductor companies that have a strong focus on wireline communications include Broadcom, Microsemi (which acquired PMC-Sierra and Vitesse), Inphi, MACOM (which acquired AppliedMicro), Marvell, and Mellanox. There are also many suppliers of optical chips for communications (e.g., Lumentum, Finisar, Neophotonics, Oclaro, Applied Optoelectronics, Acacia, etc.).

## Select Wireline Communications Private Chip Companies

Below are examples of private companies with a focus on wireline chips:

- **Analog Bits** – Analog Bits develops high-performance mixed-signal IP with a focus on high speed SerDes, as well as clocks/PLLs, I/O, sensors, and memory. It provides a broad range of SerDes IP for applications such as PCI Express (Gen4), SAS4, USB, Ethernet, hybrid memory cube, and Fibre Channel, with data rates up to 22.5Gbps and advanced development of 28Gbps and 30Gbps solutions. Its sensor solutions include low power PVT (process, voltage and temperature) sensors. Its clocking IP includes a variety of PLLs and low jitter clocks for a broad range of markets and applications. In 2017, it announced advanced SerDes and PVT sensors for TSMC's 7nm process. It is based in Silicon Valley.
- **Barefoot Networks** – Barefoot Networks has developed a high-speed programmable P4 Ethernet switching chip (Tofino) and a complete software development environment (Capilano) that enables substantial network control, flexibility, and analytics. Tofino uses a protocol independent switching architecture (PISA) and the packet forwarding logic is in software (not silicon), enabling protocol independence and

software defined networking. It is available in versions that support up to 6.5Tbps (e.g., 100G x 65, 50G x 135, etc.). The software enables great flexibility as well as advanced analytics to provide anomaly detection and insights (including its new “Deep Insight” software which provides full visibility into each packet). It has announced partnerships with Cisco and Google, and qualification with Alibaba, Baidu, and Tencent. Investors include Google, Dell, Alibaba, HP Enterprise, Tencent, Sequoia, Andreessen Horowitz, Goldman Sachs, Hermes, and DHVC. It reportedly raised about \$130 million during the 2015/2016 timeframe (Series B and C). It is based in Silicon Valley.

- **cPacket** – cPacket has developed an algorithmic fabric chip (cVu) that enables high speed packet inspection and network performance monitoring at full-line rate. As data rates increase to 100G and above, network monitoring is increasingly important and difficult. cPacket’s cVu-based monitoring nodes can proactively find and help resolve potential network issues before they impact performance. cPacket also provides visualization tools which can provide a detailed view of network operations, and its cStor arrays can store data packet information for forensic analysis, enabling fast troubleshooting. Its cClear platform can analyze thousands of links in real-time. In early 2017, it announced support for 100Gbps networks. It also announced integration with Cisco’s Firepower. It is based in Silicon Valley.
- **Credo** – Credo provides high-speed SerDes and PHYs. The company initially developed and licensed SerDes IP including 28Gbps NRZ and 56Gbps PAM-4 on a TSMC 16nm process. Credo subsequently introduced a variety of its own chip solutions including high speed CDRs (including a 16 x 56G chip) and various connectivity solutions (e.g., conversion between 8x56G and 4x112G). It indicates that it made the first public demonstration of 100G PAM-4. Credo announced a partnership with Foxconn Interconnect Technology for a 100G active copper cable. In September 2017, it announced a single-lane 100G PAM-4 SerDes. In 2015, Credo announced an \$8 million Series A capital raise led by Walden. It is based in Silicon Valley.

- **HiLight Semiconductor** – HiLight Semiconductor provides analog and mixed-signal CMOS components for high speed optical communications, with a focus on transimpedance amplifiers (TIAs) and integrated laser driver combo solutions. Its TIAs support data rates from 1.25Gbps to 12Gbps for applications including Gigabit and 10 Gigabit Ethernet, PON, SONET, and other optical communications markets. In mid-2017, it introduced a “combo” chip that integrates a laser driver, a limiting amplifier, an 8051 microcontroller, and other functionality, and supports 10Gbps communications with very low power consumption. The combo chip can be used in conjunction with its TIAs. It received a grant from the European Commission for 100G component development. In March 2017, it announced the completion of a \$6.1 million capital round. Investors include Atlantic Bridge, Oyster Capital, and Gary Steele. It is based in Southampton, UK.
- **IP Light** – IP Light develops advanced Optical Transport Network (OTN) processors for telecommunications and wireless infrastructure applications. As previously noted, OTN is commonly used in high speed telecommunications (replacing legacy SONET technologies) and is generally expected to be a core technology for 5G front haul networks. Its Apodis II family includes OTN processors that incorporate switching fabric and SERDES. For example, the IPL4002 40G OTM solution supports up to 16 client ports (up to 11G each) and 4 line side ports. It has also developed a family of IP cores (Orion) for CPRI-over-OTN Mobile Fronthaul (MFH). Investors include Cedar Fund, Club 100 Capital, and GPB Capital. It is based in Israel.
- **Kandou Bus** – Kandou Bus has developed a new type of communications signaling technology that enables very high performance, low energy SerDes. Kandou’s core Chord signaling technology transmits correlated signals across multiple wires (similar to differential signals across only two wires) which provides substantial advantages over conventional signaling methods such as ternary or PAM-4. Kandou believes this technology is ideal for many applications

including backplane communications and in-package chip-to-chip and chip-to-memory communications. It has developed IP cores around this technology including a 125Gbps/channel SerDes (Glasswing) which it indicates offers 40% higher throughput over conventional SerDes at the same baud rate over the same number of pins/wires. It also licenses its core technology to SoC chip companies. In early 2016, it announced it licensed Glasswing to Marvell. In July 2016, it announced a \$15 million investment from Bessemer. The CEO had previously invented other well-known communications codes (tornado codes, raptor codes). Kandou is based in Lausanne, Switzerland.

- **MultiPhy** – MultiPhy has developed high-speed DSP-based communications chips with a focus on data center connectivity. In September 2017 it announced the availability of its MPF3101 chip, which it indicates is the world's first 100G single-wavelength PAM-4 DSP-based IC. MultiPhy believes this has significant advantages in terms of cost (fewer optical components) and performance over alternatives, and that it is well positioned for 400G (4x100G). The chip leverages its core FlexPhy technology which utilizes leading-edge digital signal processing and PAM-4 single wavelength technology. It is targeting high speed connections within and between data centers, as well as 5G mobile networks. In January 2016, it announced a \$17 million fund raising. Semtech is a strategic investor in MultiPhy. The company was founded in Israel and has a major office in Silicon Valley.
- **Skipio** – Skipio is focused on advanced G.fast DSL semiconductor chips. DSL (Digital Subscriber Line) is a technology that enables broadband data transmissions over the copper telecom wiring that was historically used solely for voice calls. While DSL typically cannot provide the same data rates as fiber optic cable, installing fiber to the home is extremely expensive so DSL is still commonplace. With growing demand for higher DSL data rates, the G.fast protocol was developed. Skipio's solutions include chipsets for both the CPE (customer premises) and DPU (distribution point unit) equipment. In

September 2017, it announced a family of chipsets which provide up to 2Gbps in both downstream and upstream directions (per Amendment 3 specs). It also announced a 48-port, gigabit-capable G.fast DPU. In October 2017, it demonstrated over 3.1Gbps of download and 900Mbps of upload using G.fast bonding. The company believes its core G.fast technology could also have a variety of applications beyond traditional DSL type applications. Investors include Intel Capital, Amiti, Aviv, Genesis Partners, Gemini Ventures, and Pitango. The company was founded in 2012 and is based in Israel.

- **Semitech** – Semitech Semiconductor provides semiconductors and IP for narrowband power line communications for industrial applications, home automation, and IoT applications. Its power line communications chips support all common OFDM standards and modulation schemes (BPSK, QPSK, 8PSK, coherent 16-QAM) at data rates up to 500kbps. Its chips integrate a PHY, MAC, data converters, and amplifiers. Applications include smart meters, micro-inverters, street light controllers, and other remotely monitored and controlled industrial equipment. Investors include Cleantech Ventures and Get2Volume. It is based in Australia and Southern California.
- **Silicon Line** – Silicon Line provides analog components and modules for high-speed optical communications. Its semiconductor products include transimpedance amplifiers (TIAs), VCSEL laser drivers, and SerDes for general purpose communications, as well as for standards such as HDMI and USB. Many of these chip solutions operate at 10Gbps or more. In 2017, it announced that it would be introducing new products to support up to 90Gbps (18Gbps per lane). The company also provides optical modules for active optical cables (for HDMI and USB). Investors include Capital-E, Munich Venture Partners, WPG Holdings, and LRM.

- **Valens** – Valens is a major supplier of HDBaseT communications semiconductors. HDBaseT is a wireline standard for uncompressed transmission of high-definition video, audio, power, home networking, Ethernet, USB, and control signals over common Cat5 cable. While Valens was initially focused on consumer applications, it introduced HDBaseT Automotive technology in 2016 to address the connected car including bandwidth needed for infotainment and ADAS. In April 2017, Valens announced a \$60 million capital raise that included investors such as Israel Growth Partners (IGP), Delphi, Samsung, Goldman Sachs, and MediaTek. It is based in Israel.



## Video/Audio Chips

---

### Video/Audio Chip Intro

This section briefly discusses semiconductor components for video, audio, and voice-related functions.

### Video Compression Intro

Video is extremely data intensive. Even short high-definition video clips require substantial amounts of data, which makes video difficult to store and transmit. As a result, compression is often used to reduce the size of video files. There are a variety of methods for compressing video. As one example, since not that much typically changes from one frame of video to the next (e.g., in a 60<sup>th</sup> of a second), instead of storing every pixel of every frame, in many cases only the changes from the previous frame are stored, which reduces the amount of data required. By using a variety of different compression techniques, the amount of data required can be greatly reduced, which makes storing and transmitting easier. To play the video, however, the compressed video data must be uncompressed. Some common video compression standards include MPEG-2, MPEG-4, H.264, AVCHD, H.265 (HEVC), and AV1. Compression/decompression is also often used for audio to reduce the amount of data for storing audio files and data.

### Fewer Stand-Alone Digital Consumer Video Chips

During the 1990s and early 2000s, there were a large number of semiconductor chip companies focused primarily on digital video ICs and display chips for end products such as DVD players, digital cameras, DVRs, and LCD displays. However, these markets have generally been commoditized and in many applications the digital video and display functionality is integrated into larger SoCs or performed by a general processor. For example, smartphones can capture and play video, but the compression/decompression is typically one of many functions performed by an applications processor or SoC, rather than by a stand-alone video processing chip. That said, there are still a variety of applications that utilize stand-alone video processing chip solutions.

For example, many security cameras and high definition digital video recorders require advance video processing chips. Many drones can record video, but to effectively transmit high definition video with low power consumption is challenging, and specialized chips may be used. Video processing is also increasingly being integrated into automobiles. Initially, this was with rear-view cameras (showing the driver what is behind the car when backing up), but there is growing interest in including more video features in cars (e.g., digital mirrors). In addition, various types of video technologies are needed for autonomous vehicle applications, as captured video often needs to be enhanced before it can be effectively analyzed.

As one example of a stand-alone video processing chip, the Ambarella S5 is an integrated video encoder SoC chip targeting security cameras. It supports the latest encoding/compression technologies (4Kp60 HEVC/H.265 and 4Kp60 AVC/H.264) and provides functions such as multi-streaming, rate control, transcoding, 360° de-warp, and ultra-low light processing. It integrates a 1.2GHz quad-core ARM Cortex processor. Ambarella also provides similar video SoC solutions optimized for drones, sports cameras, video infrastructure, and automobiles.

## Digital Audio Chips

When digital audio is played, the digital data is converted to analog via a data converter, and then amplified by an amplifier, before a speaker converts the electrical signal to sound waves. Data converters and amplifiers were previously discussed in the Analog section, but there are specialized versions that are optimized for audio. For example, high quality audio data converters typically have high resolution (24 bits) but don't require high data rates. Audio amplifiers are usually designed to be highly efficient (especially in portable devices) since they tend to consume a significant amount of power. Often several ADCs and DACs are integrated into a single chip "audio codec" semiconductor. Many advanced audio functions (e.g., echo cancellation, echo suppression, etc.) are implemented in software. Often, these software functions, as well as audio encoding/decoding, are run on general processors or larger SoCs. However, in some products, a stand-alone audio DSP is used for audio processing.

To provide one example of an audio chip, the Cirrus Logic WM8281 “smart codec” is an integrated audio solution that incorporates six 24-bit ADCs, eight 24-bit DACs, a multi-core audio-optimized DSP/processor, and other related functions. It provides noise reduction, acoustic echo cancellation (AEC), stereo ambient noise cancellation (ANC), and speech enhancement.

### Voice-Assistant Related Components

Voice interfaces have become increasingly pervasive during the past few years. Most smartphones and PCs now incorporate a personal voice assistant, and voice interfaces are now commonplace in cars, appliances, and smart speakers. Many of the major technology companies have developed voice-based personal assistants, and Amazon Alexa and Google Assistant are now incorporated into many devices. While voice assistants are primarily software-based, they have driven demand for related chips to improve the voice interface. Some examples of this are mentioned in the private company subsection at the end of this section.

### Select Video/Audio Chip Companies

Major suppliers of stand-alone video processing chips include Ambarella, Sony, Texas Instruments, Panasonic, Socionext, Realtek, and Sunplus. Suppliers of audio chips include Cirrus Logic (provides the audio codec for the iPhones), AKM Semiconductor, Analog Devices, Synaptics (which acquired Conexant), Knowles (which acquired Audience), Maxim, NXP, Realtek, ST Microelectronics, and Texas Instruments.

### Select Private Audio/Video Chip Companies

The following are some examples of private chip companies focused on audio and/or video chips and technologies.

- **ESS** – ESS Technology develops analog components for audio applications. Its SABRE DACs are high-end solutions for premium home theater equipment and audio products. Its SABRE HiFi DACs are specifically optimized for mobile applications, bringing premium audio quality to mobile products. It also provides headphone audio amplifiers, low noise LDOs, and integrated audio headphone SoCs. In 2017, it announced that its quad-DAC headphone SoC was designed

into LG's V30 and V30+ smartphones. The company was founded in the 1980s and initially focused on speech synthesis. It later became a major supplier of video chips, but subsequently pivoted to focus on premium analog audio components. It is based in Silicon Valley.

- **GEO Semiconductor** – GEO develops chips for image sensor/video processing, with a focus on automotive applications. While cars have a growing number of cameras, the images from image sensors are often distorted and need to be enhanced for ADAS and autonomous system processing. GEO's chips can significantly enhance data from image sensors. As an example, its GW4200 Automotive Camera Processor integrates numerous image and video functions on a single chip (e.g., ultra-wide field of view lens distortion correction, electronic pan/tilt/zoom, digital calibration, a proprietary eWARP Geometric processor, noise reduction, on-screen displays). Applications include backup cameras, mirror replacement cameras, surround-view cameras, and driver monitoring cameras. It recently introduced a more advanced GW5200 automotive processor that integrates a vision processor and a vector graphics engine. It is based in Silicon Valley.
- **Libre Wireless** – Libre develops embedded voice-enabled solutions for mainstream audio/consumer, smart home, and IoT applications. In late 2017, it announced MAVID (Multiprotocol Audio, Voice, IoT Device) – a highly integrated, low power hardware/software system-in-a-package solution for embedded voice/audio applications. It includes Amazon Alexa Voice Service support, providing a complete single component Amazon Alexa solution. It is customizable and integrates wireless technologies (802.11ac dual band Wi-Fi, BT4.1, BT audio streaming, Bluetooth Low Energy) and a variety of voice enhancements (far field voice processing, wake word recognition). MAVID also offers options for adding Zigbee, LTE, and other wireless standards. Earlier in 2017, it announced its LS5BV LibreSync Wi-Fi streaming and voice assistant module featuring optimized far-field smart voice assistant capability as well as extensive music streaming and multiroom audio features. Its website lists many partners (e.g.,

Amazon, Apple, Google, and several major chip and technology companies). It is based in Silicon Valley.

- **Merus Audio** – Audio amplifiers tend to consume significant power, which is especially an issue in mobile and portable devices. Merus Audio has developed a family of amplifiers (eximo) that have a variety of differentiated features including multi-level switching technology (in lieu of conventional single level switching for Class D amplifiers) which enables much higher efficiency, especially at normal operating levels. It can also eliminate the need for large LC filters associated with conventional switching amplifiers, reducing cost and size. A power management algorithm automatically selects the most optimal power mode. It has announced several audio amplifier chips based on its technology and some early customers. Investors include Capital-E and VF Venture. It is based in Denmark.
- **Noveto** – Noveto has developed technology (Sowlo) that can direct audio directly to a person's ears, such that only that person can hear it. It utilizes sensors that constantly locate and track a user's position. Incoming audio signals are processed using Noveto's proprietary DSP engines, and small transducers focus audio signals to a user's ears. Although Sowlo has a variety of applications, automotive has been a target. Its technology can enable, for example, navigation directions, music, or voice messages to be directed only to the driver, without disturbing others (this can also provide privacy for messages and calls). In mid-2016, it announced that it was among the first three selected for Daimler's Startup Autobahn program. It is based in Israel.
- **SoundChip** – SoundChip is a team of acoustic technology experts with considerable experience in optimizing audio solutions and is active noise cancellation (ANC). It develops audio signal processing platforms that enable digitally programmable noise cancelling, audio equalization, and open-ear listening. Its offerings include digitally programmable headphone amplifier chips, intelligent audio control software, and a variety of audio tools. Leveraging its expertise,

SoundChip assists OEMs in optimizing the acoustic performance of products. It announced it licensed IP to STMicroelectronics for innovative noise cancelling audio engines and MEMS microphones that enable headphones and earphones with best-in-class noise cancelling capabilities. In 2015, it announced a strategic partnership with Panasonic Aviation. In 2017, it introduced Aurora, a low-power, hybrid noise cancelling platform for next-generation True Wireless Stereo (TWS) headsets. Aurora delivers leading-edge digitally-programmable noise cancelling, as well as many software-controlled features. The company is based in Switzerland.

- **Tempo** – Tempo Semiconductor sells analog/mixed-signal audio chips, including audio codecs. Many of its chips integrate a variety of audio functions into one chip. For example, its ACS422x00 incorporates a stereo codec, a stereo headphone amplifier, a stereo speaker amplifier, four programmable PLLs, and support for analog or digital microphones. It also provides ultra-low power solutions. Target applications include headsets, headphones, AR/VR devices, docking stations, and mobile phones. The company is based in Austin, Texas.
- **Trigence** – Trigence has developed a unique digital audio technology (Dnote). Unlike conventional digital audio (in which digital audio data is converted to analog and then amplified before a speaker), Trigence’s “all digital” technology provides digital all the way to a multi-coil speaker (no data converter or amplifier). This can provide better audio quality, fewer components, and less power consumption. Trigence provides an audio DSP speaker driver chip and offers complete digital speaker modules. It recently announced a Dnote-based USB sound bar (SHARP Aquos Audio AN-SA1) and an Dnote-based audio headset (from Audio-Technica). Investors include Intel, TDK, INCJ, and Nittoku Engineering. It is based in Tokyo, Japan.



- **USound** – USound had developed advanced MEMS technology solutions for audio speakers, such as earphones, headphones, and device speakers. It uses MEMS and multiple cantilevers to push a piston up and down, replacing the traditional coil (moving in a magnetic field) that is used in conventional speakers. This provides many advantages such as smaller size (as thin as 1mm), better power efficiency (virtually no heat), faster response time, and better audio quality. In addition, other components such as amplifiers can be directly embedded. Reportedly, STMicroelectronics is manufacturing its MEMS solutions. Investors include AWS, eQventure, and Science Park Graz. It is based in Austria.
- **VocalZoom** – While there is strong interest in using voice to control many aspects of cars, one of the issues is that the performance of these systems quickly deteriorates when there is a lot of background noise (e.g., windows open, radio on, people talking, etc.). VocalZoom has developed advanced sensor solutions that can isolate a speaker's voice from the background noise, greatly improving these systems even under adverse conditions. The sensor uses optical technology to detect low level vibrations on the surface of the speaker's face to help isolate the speaker's voice. While VocalZoom's technology has a variety of applications, automotive has been a major focus and it can enhance voice control of infotainment, telematics, or basic car functions. It can also be used to accept commands only from the driver (so words from kids in the back seat, for example, won't affect the car's operation). Two of its announced partners include Honda and Intel. Investors include 3M Ventures, OurCrowd VC, Motorola Solutions, JTI, Fuetrek, SilverFish, and Radiant VC. It is based in Israel.

- **XMOS** – XMOS has developed voice and audio chips and software. An issue with voice recognition is users typically want to provide audio commands while far from an electronic device, and this often causes problems with speech recognition systems. XMOS has developed a far-field voice controller solution (xCore) that integrates a MEMS microphone interface, advanced voice pre-processing (beamforming, noise suppression, echo cancellation, etc.), far-end processing, keyword detection, and a variety of interfaces. Its solution produces high-quality clear speech that is well suited for speech recognition engines, even from 5 meters or more. XMOS also provides a family of microcontrollers that integrate audio processing. Its solutions utilize its xCore multi-core microcontroller technology. In July 2017, it announced the acquisition of Setam, which developed audio signal separation technology. It also introduced its VocalFusion 4-Mic Dev Kit for Amazon Alexa Voice Service (AVS). In 2017, it completed a \$15 million raise with investors that included Infineon, Amadeus, Bosch, Draper Esprit, and Foundation. It is based in Bristol, UK.

## Discretes and Passives

---

### Intro to Discrete Components

While the vast majority of the semiconductor market is integrated circuits, there is still a sizable market for discrete components, which are single semiconductor device chips, such as a transistor or a diode. In general, discretes tend to be relatively low value components (e.g., often a few cents), but there are typically many of them on a printed circuit board, in addition to the more expensive integrated circuits. One example of an application for discretes is to protect the integrated circuits from power surges (which is why they are external). While discretes are generally viewed as commodities, there are several types of higher-value discretes, especially those used for very high power or high frequency applications. Some examples of higher value discrete components include:

- **Power Transistors** – When an electronic circuit has to switch power on or off, or from one line to another, a power transistor is often used for the actual switching. These transistors must be able to handle high levels of current and voltage. They are often classified by the type of transistor (e.g., power MOSFET, bipolar power transistors), the switching speed, and the level of current and voltage they can handle. In some cases, they are made from non-silicon materials (e.g., GaN).
- **IGBTs** – Unlike transistors which have three doped regions, Integrated Gate Bipolar Transistors (IGBTs) have four doped regions. They are typically used for power applications and power switching. Relative to power MOSFETs, they are often preferred when voltages and currents are very high, but switching speeds aren't especially high.
- **RF Transistors** – RF communications systems often include transistors that operate at very high frequencies.



### Discrete Suppliers

Major supplier of discretes include Infineon, ON Semiconductor, Mitsubishi, STMicroelectronics, Vishay, Toshiba, Fuji, Renesas, Rohm, and Diodes.


## Intro to Passives

Passive components are not semiconductor devices, but are used extensively in electronic circuits. Passive technology is also often embedded in integrated circuits (e.g., DRAM cells, for example, integrate capacitors). Passives often sell for a few cents (or even less than a cent), although there are examples of high value passives. Common passives include:

- **Resistors** – Resistors are one of the simplest electrical devices. The voltage across a resistor is equal to its resistance (in ohms) times the current through the resistor. For example, 3 amps of current through a 10-ohm resistor produces a voltage of 30 volts. Resistors have many applications including terminating transmission lines, adjusting signal levels, and dividing voltages. There are many different types of resistors. A potentiometer is a device in which the resistance can be adjusted. Suppliers include Yageo, Walsin Technology, KOA Corporation, Panasonic, Hokeniku Electric, and Rohm Corporation.
- **Capacitors** – A capacitor is a device in which the current is related to the change in voltage across the device. Capacitors can also store energy. A capacitor is typically formed by two conducting layers separated by a dielectric insulating layer. Increasing the surface area of the layers or decreasing the distance between the layers, increases the capacitance. Capacitors are often classified by the materials used for the conducting layers or the dielectric layer. Some common conducting materials include tantalum, aluminum, and niobium oxide, and typical dielectric layers include ceramic, paper, and glass. The picture to the right illustrates typical types of capacitors. Major capacitor suppliers include Murata, Kyocera/AVX, Vishay, and KEMET.



- **Inductors** – Inductors are passive devices in which the voltage across it is related to the change in current passing through it. When current passes through an inductor it induces a magnetic field that creates a voltage. It can be formed by wrapping a copper wire around a plastic or magnetic core. To the right is an example of an inductor. Some major inductor suppliers include TDK, Murata, Vishay, Taiyo Yuden, Sumida, Sunlord, Chilisin Electronics, Bournes, Mitusmi and AVX/Kyocera.



Several passives can be combined to create integrated filter components. Filters are commonly used in a variety of electronic circuits to weed out undesired signals and to let only desired signals pass. In many wireless systems, for example, RF filters are used to allow only the desired frequency signals to pass through and filter out other frequencies. SAW and BAW filters were previously discussed in the wireless section, but most filters consist of combinations of passive components (resistors, capacitors, and inductors). There are also a growing number of “integrated passive” solutions in which a variety of passives are combined into single components for a specific application (to reduce the number of components).

## Semiconductor Capital Equipment

---

### Intro to Semi Cap Equipment

The beginning of this chapter briefly outlined the chip manufacturing process. This section discusses the manufacturing tools in slightly more detail:

- **CVD** – Chemical Vapor Deposition (CVD) equipment deposits thin layers of materials across a wafer using a chemical reaction. Typically, a wafer is placed inside a sealed chamber and the chamber is filled with a gas. Due to a combination of temperature and pressure inside the chamber, the gas reacts with the surface of the wafer causing a desired layer of material (thin film) to form on the surface of the wafer. The remaining gas is removed from the chamber, and the equipment is then ready to process more wafers.

One way to classify CVD tools is by the type of pressure used, such as atmospheric, low pressure, high vacuum, and ultra-high vacuum. Some CVD tools use plasma to enhance the reactions, which can reduce the temperature needed in the chamber (high temperatures can damage the materials). Atomic Layer Deposition (ALD) tools deposit one atomic layer at a time. CVD tools can also be classified by the types of materials they deposit (dielectrics, silicon, tungsten, etc.).

The largest suppliers of CVD tools include Applied Materials, Tokyo Electron, and Lam Research. These three companies account for the vast majority of the market. Other smaller suppliers include Hitachi Kokusai Electric, Aixtron, ASM International, and Veeco Instruments.

- **PVD** – Like CVD, Physical Vapor Deposition (PVD) also deposits a layer of materials across a wafer. The most common type of PVD used for chip manufacturing is called sputtering, in which ions are shot at a target material and cause very tiny pieces of the target material to break off and fall on the wafer, eventually covering the wafer. Sputtering is most commonly used for metal semiconductor layers. Applied Materials is the leading supplier of PVD tools. Others include Ulvac and Canon Anelva.



- **Photolithography** – Photolithography is a critical step in the chipmaking process. A wafer is coated with a material that is sensitive to light (photoresist). An intense light is then directed onto the wafer through a mask (reticle). A mask is partially transparent and partially opaque and is designed to represent a desired pattern on the surface of the wafer. The light is blocked by the opaque portions of the mask, but passes through the transparent portions and reaches the wafer. The light causes a chemical change in the photoresist it reaches, which enables it to be easily removed by subsequently adding a “developer” solution. The net result is a pattern of photoresist that is the same as the opaque pattern on the mask. Many modern photolithography systems use a scanning approach rather than exposing the entire wafer at once. The photoresist is often resistant to etching, which enables removal of material in the areas not covered by photoresist.

A major challenge is that you would ideally want the wavelength of the light to be smaller than the feature size it is creating. However, as transistor sizes have shrunk, this is extremely difficult (at 10nm and below light is an X-ray). For many years, photolithography has been “stuck” at 193nm light, with a variety of very complex “tricks” used to achieve much smaller features sizes, but this has become more difficult. The industry was expected to move to extreme ultra violet (EUV) systems more than 10 years ago, but it ran into many technical hurdles. Last year, however, ASM Lithography announced that it expects volume commercial shipments of EUV systems with 13.5nm EUV light in 2018 for 7nm and 5nm nodes. Major suppliers of photolithography tools include ASM Lithography, Canon, and Nikon.

- **Etch** – Etching removes material and is typically performed after photolithography, as the photolithography step leaves the desired patterns and the etching tools can, for example, etch away at areas that aren’t covered with photo-resist. Etch is often classified by the types of materials it removes. For example, dielectric etch is used to create vias in dielectric layers (these vias are often filled with tungsten and create connections between metal layers). Polysilicon etch is

often used to create gates at the transistor level, and metal/conductive etch is used to remove metal. Plasma etching involves applying energy to a gas which creates positive ions that are pulled towards the wafer and etch away materials. Some major suppliers of etch tools include Lam Research, Applied Materials, and Tokyo Electron.

- **Ion Implantation** – Ion implantation produces ions and accelerates and directs them onto the wafer to create p-type and n-type regions on the surface of a wafer (doping different regions to create patterns of transistors on the wafer). It can also be used to create silicon on insulator substrates. Ion Implantation suppliers include Applied Materials and Axcelis.
- **Inspection and Metrology** – The chip manufacturing process is long with many steps. If problems are found in final testing, it can be difficult to determine exactly where the problems occurred. As such, chipmakers inspect wafers many times during the process to identify if there are problems with manufacturing. This can be done manually, but there are a variety of tools that can automatically inspect and analyze wafers to identify potential issues, which can greatly improve manufacturing yields and reduce costs. Additionally, there are a variety of “metrology” tools that are used to measure various characteristics (layer thickness, uniformity, shape of etched areas, etc.) in order to monitor and improve yields. The leading supplier of wafer inspection tools is KLA-Tencor, which also provides metrology equipment. Other suppliers of inspection and/or metrology tools include Applied Materials, ASML Holding, Hitachi High-Tech, Nanometrics, and Rudolph Technologies.

There are many other types of semiconductor production equipment such as chemical mechanical polishing (to polish each layer to ensure it is level), furnaces (to heat wafers to a desired temperature), and cleaning tools (to ensure any stray particles are removed between steps).

In addition, there are a variety of tools used for the “back end” of the process in which the wafers are diced up into individual die and wire bonded into packages. Some wire bonding tool companies include Kulicke & Soffa, ASM Pacific Technology, and BE Semiconductor. Lastly, there are testing tools, which are often referred to as automatic test equipment (ATE), that perform final testing on chips before they ship to customers. Suppliers of ATE equipment include Teradyne, Advantest, and Xcerra.

### **Semiconductor Equipment Market and Players**

In July 2017, SEMI projected that semiconductor manufacturing equipment sales would increase nearly 20% in 2017 to \$49.4 billion, which would be a record (the previous high was \$47.7 billion back in 2000). That includes \$39.8 billion for wafer processing equipment, \$3.9 billion for test, \$3.4 billion for assembly and packaging tools, and \$2.3 billion for other types of equipment. It projected 7.7% growth in 2018. Final 2017 results are not yet available as of early January 2018.

The semiconductor equipment sector is concentrated with a few companies (Applied Materials, ASML, Lam Research, Tokyo Electron, KLA-Tencor) accounting for a large portion of the market. Lam Research’s acquisition of Novellus in 2012 further consolidated the market. However, two very large deals that were announced but never closed were Applied Materials’ announced acquisition of Tokyo Electron (announced in 2013, but terminated in 2015 due to regulatory issues) and Lam’s announced merger with KLA-Tencor (announced in August 2016 but terminated a few months later due to regulatory issues).

## Select Semi Cap Equipment Start-Ups

There are relatively few semiconductor capital equipment start-ups. Given the complexity of the chip manufacturing process, the large chip manufacturers are generally reluctant to work with start-ups and it is challenging for a private company to displace one of the major semi cap equipment companies. Below is one example.

- **CRC** – CRC has developed a unique technology (“Infinity”) for bonding ceramics with ceramics, metals, and other materials, and its initial target market is semiconductor wafer pedestals. Inside many semiconductor tools, the semiconductor wafers sit on a ceramic object known as a wafer pedestal. These pedestals also often include sensors and heaters to measure temperature and to help heat the wafer to a desired level. Historically, the pedestals were formed as one piece in a furnace as many ceramics are done, and then grinded down to create the pedestal. However, this results in low yields and high costs. In addition, it greatly limits the placement accuracy and number of sensors/heaters that can be integrated. In contrast, the CRC Infinity technology enables bonding of the pedestal pieces, greatly reducing cost and lead times. In addition, it allows many accurately placed sensors and heaters, which can improve yields. It can, for example, enable multi-zone heating in which certain areas of a wafer are heated to higher levels than others (to compensate for the center of a wafer often having different temperatures from the edges). CRC has a pedestal repair business and initially used Infinity to repair pedestals for chipmakers but is planning to begin volume commercial shipments of new Infinity pedestals in 2018. The company is based in Silicon Valley.

## EDA Software

---

### Intro to EDA

Designing a semiconductor chip is extremely complex, and engineers typically use a variety of software tools to help with these designs. The design software tools are often referred to as electric design automation (EDA) software. There are dozens of different types of EDA software. Some of the major segments include:

- **RTL Synthesis** – For digital chips, engineers often initially develop a high-level description of the detailed functionality they want in a chip using a hardware description language (HDL). The two most common HDLs are Verilog and VHDL. RTL (Register Transfer Level) synthesis software converts the software description of the chip into an implementation of the design, typically at the logic gate level.
- **Physical Implementation (Place and Route)** – While RTL synthesis tools create a gate level design, the software doesn't provide any information about how to implement that design in an actual silicon chip. As such, there is a separate category of EDA tools that transform a gate level design into a detailed physical description of an integrated circuit including where the transistors are located and the various layers and interconnections. This is typically referred to as physical implementation or “place and route” software.
- **Verification** – There are a variety of different types of software tools to help verify that a design will work. Functional verification software, for example, checks to make certain that a design operates as expected from a logic perspective (i.e., the chip provides the appropriate output for a given input). There are also tools to help analyze for timing issues, power issues, and physical implementation issues. Emulation tools can emulate how a chip will operate and can be used to test embedded software before a chip is manufactured.

The three categories noted above are some of the major segments of EDA but there are a variety of others and many sub-segments within those categories. Some other examples include design for manufacturing (DFM) software, electronic systems-level (ESL) design (which helps model designs at a high level early in the process), and analog simulation (which helps simulate how analog chips will operate under various conditions).

## EDA Suppliers

The semiconductor EDA software market has generally been an oligopoly with Synopsys, Cadence, and Mentor (which was acquired by Siemens) accounting for the vast majority of the market. Each of these provides a broad range of EDA solutions. For example, Synopsys' products include Design Compiler Graphical (RTL synthesis), IC Compiler II (physical implementation), IC Validator (physical verification), and many other software solutions. Each of the major EDA players encourage customers to use their entire platform of solutions (which can provide certain integration advantages), but many chip companies still take a "best of breed" approach and use combinations of software from different EDA suppliers.

The major EDA companies have all diversified. For example, Synopsys and Cadence have also become major suppliers of semiconductor IP, and Synopsys has made several non-semiconductor software acquisitions. While the EDA market is dominated by the three leaders, there are a handful of other players which typically focus on specific niches. For example, Ansys (Ansoft) and Keysight (EEsof) have solutions focused on RF-related design.

## Select EDA Start-ups

Over the years there have been many EDA start-ups. Typically, a few engineers from one of the three major EDA companies would leave and form their own company which would focus on a specific niche within EDA. The start-up would develop software, secure a few major customers to help validate the solution, and then look to get acquired by one of the major EDA companies (often the same company they previously worked at). More recently, however, there have been far fewer EDA start-ups, although there are still a variety of them. A couple of examples are below.



- **Austemper Design** – Austemper is an EDA company that provides a comprehensive tool-suite to analyze and verify the functional safety features in SoCs and ASICs. As chips become increasingly large and are used in critical applications where faults can have serious consequences (e.g., autonomous driving, certain medical and industrial applications, etc.), ensuring robust design is both difficult and critical. Austemper's solutions include SafetyScope (estimates functional safety metrics), Annealer (adds reliability enhancements to improve fault detection), RadioScope (adds reliability enhancements to the internal logic, to improve fault detection/fault tolerance), and Kaleidoscope (parallel fault simulator with hybrid simulation capability). Austemper indicates Kaleidoscope can achieve a 100x speedup relative to gate-level approaches. The company notes that it is engaged with many automotive OEMs and tier-1 suppliers, as well as many chip companies. Austemper is based in Austin, Texas.
- **NanGate** – NanGate focuses on software for semiconductor libraries. To improve design efficiencies, chip designers want to be able to reuse portions of a design for future designs. As such, there is interest in creating cell libraries which chipmakers can reuse. However, it has historically been challenging to create and validate libraries and migrate and optimize them for a given design (often some tweaks are needed for a specific application). NanGate addresses this with a variety of EDA software solutions including Library Creator (for creation and migration), Analyzer and Design Audit (for analysis and validation of libraries and IP), and Characterizer and NanSpice (for characterization). NanGate also provides some open libraries. The company is based in Silicon Valley with a major office in Denmark.

## Chapter 3: Select Semi M&A Transactions

*“I was running an assembly line designed to build memory chips. I saw the microprocessor as a bloody nuisance.”*

- Andy Grove, former Intel CEO (prior to the late 1970s, Intel was primarily focused on memory not microprocessors)

### Introduction

---

This chapter includes a table of semiconductor M&A transactions since the beginning of 2016. This is not intended to be a comprehensive list (e.g., we excluded a number of small deals), but it provides a good sense for the transactions that occurred during the past couple of years. In some cases, we excluded deals in which the target wasn't a semiconductor company (e.g., Broadcom's acquisition of Brocade or Synopsys' acquisition of Black Duck). We included valuation information when available. In some cases, the valuation data is an estimate from one or more sources (PitchBook, 451Group, news article) and may or may not be completely accurate. When available, the valuation multiple is enterprise value to trailing twelve-month revenue, although in some cases it is deal value to trailing twelve-month revenue. We note that some of these deals have been announced but have not yet closed (and may not close). We excluded a few transactions that were announced but subsequently terminated, and we did not include Broadcom/Qualcomm since a deal has not yet been agreed to (as of January 10, 2018). The chapter also highlights some of the larger M&A transactions from 2014 and 2015. The last section includes a table of public semiconductor company acquisition premiums.

### Recent Semiconductor M&A Transactions

---

#### Semiconductor M&A Deals in 2016 and 2017

The following is a table of semiconductor M&A transactions that were announced in 2016 and 2017 (and through January 10, 2018).

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Jan-18	Synopsys	Kilopass	One-time programmable non-volatile memory IP (part of Kilopass was spun off as TC Lab).			USA
Jan-18	AVX	Ethertronics	Passive and active antenna solutions including active steering solutions for Wi-Fi and cellular.	1.7	\$142	USA
Dec-17	Silicon Laboratories	Sigma Designs	Public company. Z-Wave wireless connectivity and powerline chips. Smart TV and set-top box chips.	1.8	\$282	USA
Dec-17	STMicro-electronics	Atollic	Embedded systems platforms for ARM-based microcontrollers.		\$7	Sweden
Dec-17	AEA Investors	Excelitas Tech [Veritas Capital]	Optoelectronic components including LEDs, detectors, and lighting.			USA
Dec-17	TTM Technologies	Anaren Inc. [Veritas Capital]	Passive RF components, resistive components, and RF solutions.	3.5	\$775	USA
Nov-17	Marvell Technology	Cavium	Public company. Security, server, storage, and video processor solutions and SoCs.	6.8	\$5,974	USA
Nov-17	mCube	On Semi's Xsens unit	Motion tracking technology and IP. Was previously acquired by Fairchild, which ON Semi acquired.			Netherlands
Nov-17	Dialog Semiconductor	ams' LED backlight unit	LED backlight components and IP for TVs and other display applications.			Austria
Nov-17	Siemens	Solido Design Automation	EDA tools for variation-aware design and characterization.			Canada
Nov-17	Cadence	nusemi	High-speed Serializer/Deserializer (SerDes) communications IP.			USA
Oct-17	Microsemi	Knowles' Vectron Timing Unit	Crystal oscillators, frequency translators, clock and data recovery products, SAW filters, BAW filters.		\$130	USA
Oct-17	Project Denver (Carlyle Group)	MACOM's AppliedMicro X-Gene unit	ARM-based server processors (X-Gene) that MACOM obtained in AppliedMicro acquisition.			USA

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Oct-17	Sckorpios	Novati	Heterogeneous integration technology and Austin fab.			USA
Oct-17	Synopsys	Sidense	One time programmable non-volatile memory IP (IP licensing model).			Canada
Oct-17	Dialog Semiconductor	Silego Technology	Configurable mixed-signal ICs (CMICs). OTP memory enables configuration of power chips.	3.4 (3.8)	\$276 (plus \$30.4 earn-out)	USA
Sep-17	Altair Engineering	Runtime Design Automation	EDA tools to design CPUs, GPUs, chipsets, and SoCs for high performance computing.			USA
Sep-17	Bain Capital consortium (SK Hynix, Dell)	Toshiba Memory	Major supplier of memory chips, including NAND Flash memory.		\$17,867	Japan
Sep-17	Canyon Bridge Capital Partners	Imagination Technologies	Public company. Multimedia chip IP including PowerVR GPUs, security, and AI accelerator.	4.0	\$744	UK
Sep-17	Tallwood Venture Capital	Imagination's MIPS Unit	MIPS microprocessor IP. Includes Warrior mobile (M), AI (I) and low power (P) processor IP.		\$65	USA
Sep-17	Synopsys	QuantumWise A/S	Materials modeling for chipmakers (e.g., model how materials impact transistor performance).			Denmark
Sep-17	Keysight Technologies	Scienlab electronic systems GmbH	Customer specific test solutions for electronic components in the auto and industrial markets.		\$62	Germany
Aug-17	Lam Research	Coventor Inc.	SEMulator3D modeling and analysis software for chip manufacturing processes; advanced MEMS design.			USA
Aug-17	MACOM	Luna Innovations' HSOR unit	High speed optical receivers (HSORs) and optical to electrical converters, from Luna.		\$33.5	USA
Aug-17	Littelfuse	IXYS	Public company. Primarily power MOSFETs, IGBTs, and power ICs. Also Zilog microcontrollers.	2.0	\$750	USA
Aug-17	Silvaco	SoC Solutions	Semiconductor IP and services for use in ASICs, SoCs, and FPGAs.			USA
Aug-17	II-VI	Kaia's Newton Aycliffe fab	6" wafer fab in the UK that is capable of producing GaAs, SiC, and InP chips.		\$80	UK

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Jul-17	Littelfuse	US Sensor	High reliability thermistors, probes, and assemblies, and other temperature sensors.			USA
Jul-17	Kulicke and Soffa Industries	Liteq BV	Step and repeat lithography tool for packaging applications such as system in a package.			Netherlands
Jul-17	XMOS	Setam	Signal separation technology that enables separation of speaker's voice from background noise.			USA
Jul-17	Invecas	Lattice's HDMI Design Team	HDMI design team (150 people) and Simplay Labs (standardization subsidiary).			USA
Jul-17	ARM Holdings plc [SoftBank]	Simulity Labs [Foresight]	IoT security software, including software for SIM cards and embedded SIM cards.		£12.0	USA
Jun-17	Semtech	AptoVision	Chips and software for audio/video distribution, including software defined video over Ethernet.		\$28	Canada
Jun-17	Synaptics	Conexant (from Golden Gate Capital)	Low power audio codecs, far field codecs, USB codecs, voice processors, embedded modems.	3.3	\$343	USA
Jun-17	Synaptics	Marvell's Multimedia Solutions Unit	Video and audio processing chips for set-top boxes and streaming media applications.	1.0	\$95	USA
May-17	Zhejiang Wansheng	Jiangxin Zhiben Shanghai (Analogix)	Interface chips including chips for DisplayPort as well as HDMI to DisplayPort and other related chips.		\$545	China
May-17	Sondrel	Imagination's IMG Works unit	SoC design and embedded software integration services.			UK
May-17	Kaiaam	Newton Aycliffe manufacturing	The manufacturing facilities of Compound Photonics (CP) in Newton Aycliffe.			UK
May-17	Altair	MODELiiS	EDA software for circuit modeling, system design and simulation.			France
May-17	Pixelworks	ViXS Systems	Public Canadian company. Video decoding (HD video for H.264/265) and transcoding chips.	1.2	\$20	Canada

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Apr-17	Powertech Technology	Micron's Akita unit	Assembly and test plant in Japan.			Japan
Apr-17	MaxLinear	ViXS Systems' MoCa business unit	ViXS' Xconnex family of MoCA solutions for set-top box applications and patent licenses.		\$5	Canada
Apr-17	Unic Capital Management	Xcerra	Public company. Testing solutions including LTX/ Credence semiconductor test products.	1.4	\$580	USA
Apr-17	Advanced Micro Devices	Nitero (IP and staff assets)	60GHz wireless communications chips for mobile devices including AR/VR applications.			USA
Apr-17	Mercury Systems	Delta Microwave	RF and MMIC sub-assemblies and components for military and space applications.	3.2	\$40.5	USA
Mar-17	Analog Devices	OneTree Microdevices	GaN and GaAs amplifiers for broadband applications including cable TV and fiber to the home.			USA
Mar-17	MaxLinear	Exar Corporation	Public company. Power chips (LDOs, regulators, PMICs), interfaces (UARTs), and compression chips.	3.0	\$700	USA
Mar-17	Micron	Cando assets	Backend equipment company (acquired in bankruptcy).		\$89	Taiwan
Mar-17	TDK	ICsense NV	ASIC and IC design services. Focus on sensor and MEMS interfaces, high voltage, and power.			Belgium
Mar-17	Murata	Arctic Sand Technologies	"TIPS" power solutions with high switching frequency and efficiency for LED backlights, PC regulators.			USA
Mar-17	ams AG	Princeton Optronics	VCSEL lasers including high-power VCSELs and bare VCSEL die for depth cameras, LiDAR, and medical.	5.3 (12.8)	\$53.3 (plus earn-out of up to \$75)	USA
Mar-17	Intel	Mobileye N.V.	Public Company. Major supplier of vision processing chips and software for ADAS and autonomous vehicles.	41.0	\$15,300	Israel
Mar-17	ON Semiconductor	IBM's mmWave unit	IBM's millimeter wave chip unit in Haifa Israel which was focused on radar chips for automotive.			Israel



Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Mar-17	SkyWater Technology Foundry	Cypress Semi's fab in Minnesota	A 200mm wafer fab in Bloomington, Minnesota.		\$30	USA
Feb-17	KEMET Corporation	NEC TOKIN	Supercapacitors, inductors, relays, filters, and other passive components.		\$149.2 (for shares it didn't own)	Japan
Feb-17	ARM Holdings plc [SoftBank]	Mistbase AB and NextG-Com	Both companies have expertise and technology related to NarrowBand-IoT wireless solutions.			Sweden, UK
Feb-17	Amkor	Nanium	Wafer-level fan-out (WLFO) semiconductor packaging solutions.			Portugal
Feb-17	Integrated Device Technology	GigPeak (fka GigOptix)	Public company. High speed drivers, amplifiers, and TIAs for optical comm. Video encoding chips, ASICs.	3.9	\$250	USA
Feb-17	MaxLinear	Marvell's G.hn Unit	Marvell's G.hn home networking unit which can operate over phone lines, coax, or power lines.		\$21	USA
Feb-17	Veeco Instruments	Ultratech	Public company. Lithography equipment for packaging and LEDs, wafer inspection tools.	2.8	\$815	USA
Jan-17	SK Holdings (SK Hynix)	LG's Siltron silicon wafer unit (51% stake)	Silicon wafers used for semiconductor manufacturing.		\$527	South Korea
Jan-17	Silicon Laboratories	Zentri	Modules, software, APIs and tools for rapid development of secure IoT solutions.			USA
Jan-17	Novatek Micro	Faraday's Surveillance business unit	DVR and IP camera SoCs, and other related chip solutions and IP.		\$21	Taiwan
Dec-16	Panasonic	Qualtre	MEMS gyroscopes that utilize a differentiated approach based on BAW technology. 3-axis solutions.			USA
Dec-16	TDK	Invensense	Public company. MEMS gyroscopes for smartphones and other markets. Also MEMS microphones.	3.5	\$1,300	USA
Dec-16	APAT Optoelectronics	NeoPhotonics' low-speed transceiver unit	GPON and GEAPON transceiver products at up to 10G data rates, specialty transceiver products.		\$26.4	USA

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Dec-16	ams	Incus Labs	Intellectual property (IP) for digital active noise cancellation in headphones and earphones.			UK
Dec-16	WiLAN	Globalfoundries select patents	Select patents (originally from IBM) related to semiconductor chip manufacturing.			USA
Dec-16	Teledyne Technologies	e2v technologies	Public company. RF power, high-end CMOS image sensors, CCDs, data converters, other analog ICs.	2.6	\$780	UK
Nov-16	Nanoco Group	Eastman Kodak's Quantum Dot	Patents related to quantum dots for use in electroluminescent displays.			USA
Nov-16	Samsung Electronics	QD Vision	Quantum dot technology, over 250 patents (some reports indicate Samsung acquired only IP).		\$70	USA
Nov-16	MACOM	AppliedMicro	Public company. High speed comm. chips (PHYs, OTN, 100G), X-Gene ARM-based server ICs.	4.2	\$770	USA
Nov-16	Analog Devices	Vescent Photonics' laser beam steering	Non-mechanical laser beam steering technology for LiDAR applications.			USA
Nov-16	Siemens AG	Mentor Graphics Corporation	Public company. One of the three largest suppliers of EDA software. Also software for PCBs.	4.1	\$4,029	USA
Nov-16	WiLAN	Panasonic select MEMS patents	Patents related to MEMS for applications such as fitness trackers and handsets.			Japan
Nov-16	Inphi	ClariPhy	Coherent DSP chips for high-speed (100G+) optical communications.		\$275	USA
Nov-16	Solomon Systech	Microchip's mobile touchscreen ICs	Certain maXTouch chip products (touch controllers for mobile devices/OLED displays).		\$23	USA
Oct-16	Analog Devices	Innovasic	Deterministic Ethernet chip solutions for industrial applications for industrial IoT and other markets.			USA
Oct-16	Qualcomm	NXP Semiconductors	Public company. Major supplier of microcontrollers, auto chips, sensors, security chips, other ICs.	5.5	\$39,187 (ent. value of \$47,000)	Netherlands

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Oct-16	ams AG	Heptagon	Time of flight 3D sensors, VCSEL infrared emitters, Bluetooth transceivers, and other ICs.	6.3 (9.5)	\$570 (plus \$285 earn-out)	Singapore
Oct-16	Murata	IPDiA S.A.	Integrated passive devices including 3D, high reliability, and ultra broadband capacitors.			France
Oct-16	Infineon	Innoluce BV	MEMS-based mirror scanning technology for LiDAR applications.			Netherlands
Oct-16	Mentor Graphics	Galaxy Semiconductor	Test data analysis and defect reduction software for the semiconductor industry.			Ireland
Oct-16	Nokia	Eta Devices	Power amplifier envelope tracking for reducing power consumption in handsets and cellular infrastructure.			USA
Sep-16	Acacia Research	Renesas select semiconductor patents	Patents related to power mgmt., SoCs, and chip manufacturing.			Japan
Sep-16	Beijing Shanhai Capital	Analogix Semiconductor	Interface chips including chips for DisplayPort as well as HDMI to DisplayPort and other related chips.		\$500	USA
Sep-16	Tessera	DTS	Public company. Audio technology and IP.	5.8	\$850	USA
Sep-16	Applied Materials	DFMSim	Simulates critical IC design and processing steps to find issues that could reduce yield or increase costs.			USA
Sep-16	X-Fab	Altis Semi fab assets	Foundry in Paris area with 200mm wafer processing. Was in insolvency proceedings.			France
Sep-16	Renesas Electronics	Intersil	Public company. Major supplier of analog components (power management, amplifiers, etc.).	5.8	\$3,219	USA
Sep-16	Moschip Semiconductor	elitePLUS Semiconductor Technologies	Mixed signal and lower-power semiconductor verification.			India
Sep-16	Moschip Semiconductor	Orange Semiconductors	Chip integration services and embedded software.			India

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Sep-16	Supernova Investment	Actions Semiconductor	Public company. Portable multimedia and mobile internet system-on-a-chip (SoC) solutions.	2.4	\$97	China
Sep-16	Intel	Movidius	Myriad 2 vision processing unit incorporating array processing and hardware accelerators.			USA, Ireland
Aug-16	Littelfuse	ON Semi select auto electronics components	Transient voltage suppression diodes, switching thyristors, and IGBTs for ignition applications.	1.9	\$104	USA
Aug-16	TDK	Tronics	Public company. MEMS foundry with a focus on MEMS sensors.	6.8	\$54	USA
Aug-16	Juniper Networks	Aurion	Silicon photonics including tunable lasers and laser arrays combining InP on silicon wafers.		\$165	USA
Jul-16	STMicro-electronics	Ams' NFC and RFID business assets	NFC and RFID reader solutions, as well as about 50 people.		\$77.8 (plus \$37 earn-out)	Austria
Jul-16	ams AG	MAZeT GmbH	IC and filter design with expertise in optical color and spectral sensors for medical/industrial.			Germany
Jul-16	Analog Devices	Linear Technology	Public company. High-performance analog chips (power management, amplifiers, other).	9.4	\$14,880	USA
Jul-16	Tsinghua Unigroup	Wuhan Xinxin Semiconductor Manufacturing	Semiconductor foundry with an initial focus on NOR Flash and CMOS image sensors, and IoT chips.			China
Jul-16	SoftBank	ARM Holdings	Public company. Leading supplier of embedded microprocessors (several Cortex families) and other IP.	20.9	\$32,434	UK
Jun-16	Rambus	Inphi's memory interconnect unit	Memory interconnect and buffer chip solutions that enable more and faster memory capacity.		\$90	USA
Jun-16	SMIC	LFoundry Europe GmbH (70% stake)	Specialized analog foundry in Italy and Germany. 40K wafer/month capacity with 90nm processes.		\$55	Germany
Jun-16	ams	Cambridge CMOS Sensors	MEMS-based "hot plate" sensors for detecting a variety of gases (carbon monoxide, carbon dioxide, etc.).			UK

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Jun-16	ASML Holding	Hermes Microvision	Public in Taiwan. E-beam Inspection tools for fabs. Technologies for sub-10nm defect inspection.	16.5	\$3,087	Taiwan
Jun-16	Cavium	QLogic	Public company. High-speed storage interfaces including Fiber Channel controllers and modules.	2.2	\$1,360	USA
Jun-16	Jianguang Asset Mgmt. / Wise Road	NXP's Standard Products unit	Discrete components including MOSFETs, bipolar transistors, and diodes.	2.3	\$2,750	Netherlands
Jun-16	Rambus	Semtech's Snowbush IP assets	Semiconductor IP including SerDes, PHYs and other IP for high-speed communications applications.	1.9	\$32.5	USA
Jun-16	Silvaco	Ipextreme	Semiconductor IP including interfaces, processors, and others. Also, biometric analysis software.			USA
Jun-16	Beijing E-Town Chipone Technology	Exar's Integrated Memory Logic	Power management and color calibration chips for the display and LED lighting markets.		\$136	Taiwan
May-16	CML Microsystems	Wuxi Sicom Technologies	Digital radio ICs, vocoders, and digital radio baseband signal processors.		\$11	China
May-16	ARM Holdings	Apical	Imaging and embedded vision IP for smartphones and other vision applications.		\$350	UK
May-16	Advanced Semiconductor Engineering	Siliconware Precision Industry (SPIL)	Public company (Taiwan). Major semiconductor assembly and test provider.			Taiwan
May-16	Synopsys	Gold Standard Simulations	EDA simulation solutions for design technology co-optimization for advanced process nodes.			UK
May-16	GlobalWafers	Topsil's wafer business unit	Silicon wafer manufacturer for high power devices, including wafer factories in Denmark and Poland.	1.2	\$48.5	Denmark
May-16	WiSeKey International	INSIDE Secure's IoT chip business	VaultIC security chips (secure RISC CPU, crypto-accelerators, random number generators, etc.).	0.4	\$13	France
May-16	NavInfo	AutoChips	SoCs and communications chips for automobile connectivity.		\$600	China

Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
May-16	MaxLinear	Broadcom's Wireless Backhaul unit	Wireless backhaul communications chips (most of which came from Broadcom's Provigent acquisition).	2.7	\$80	Israel
Apr-16	Cypress Semiconductor	Broadcom's Wireless IoT Unit	Wireless chips for IoT applications (WiFi, Bluetooth, Zigbee, WICED), not for other markets.	2.9	\$550	USA
Apr-16	MaxLinear	Microsemi's broadband wireless unit	Wideband RF transceivers and synthesizers for cellular base stations.		\$21	USA
Apr-16	POET Technologies	DenseLight Semiconductors	Compound semiconductor optoelectronic devices and photonic integrated circuits.	4.0	\$10.5	Singapore
Apr-16	Qorvo	GreenPeak Technologies	Supplies IEEE 802.15.4 and ZigBee wireless communications chips and software.			Netherlands
Apr-16	FOGALE Nanotech	Altatech SAS	Wafer inspection and material deposition tools for semiconductor manufacturing.			France
Apr-16	Cadence	Rocketick	Multicore parallel technology. Can accelerate RTL and DFT simulations using standard x86 servers.		\$40	Israel
Apr-16	Molex (Koch Industries)	Interconnect Systems	Design and manufacture of high density silicon packaging with interconnect technologies.			USA
Apr-16	Intel Corporation	YOGITECH SpA	Functional safety technology for semiconductors.			Italy
Apr-16	GigOptix	Magnum Semiconductor	Video encoders, decoders, and transcoders for H.264/H.265/MPEG-2 for infrastructure applications.	2.9	\$55	USA
Mar-16	Mercury Systems	Microsemi RF, microwave, security units	Select embedded security, RF, and microwave board-level systems and packaging products.	3.0	\$300	USA
Mar-16	Analog Devices	SNAP Sensor SA	Technology improves optical sensors by ensuring accurate image detection.			Switzerland
Mar-16	Cisco	Leaba Semiconductor	Networking semiconductor chip company. Was still in stealth mode.		\$320	Israel



Announce Date	Acquirer	Target	Target's Products	Val./Rev.	Total Deal Amt. (\$M)	Target Country
Mar-16	Synopsys	WinterLogic	Fault simulation software used in automotive/safety environments for design/verification of SoCs.			USA
Feb-16	FormFactor	Cascade Microtech	Major supplier of probe cards and related technologies for chip testing.		\$352	USA
Feb-16	Sigma Designs	Bretelon	ASICs and software for mobile IoT applications.		\$22	USA
Feb-16	e2v technologies	Signal Processing Devices	Provides high speed analog to digital converters. Also provides ADC IP.		\$14 (plus \$4 earnout)	Sweden
Jan-16	Sony	Altair Semiconductor	LTE baseband modem chips, and other wireless communications chips.		\$212	Israel
Jan-16	Ciena Corporation	TeraXion's high-speed photonics	High-speed InP and silicon photonics technologies and IP.		\$32	Canada
Jan-16	II-VI	ANADIGICS	Public company. GaAs semi chips including power amplifiers, LNAs, and RF amplifiers.	1.0	\$61	USA
Jan-16	II-VI	EpiWorks	Compound semiconductor (GaAs, InP, etc.) epitaxial wafers for optical components and wireless devices.	3.1	\$43	USA
Jan-16	Microchip Technology	Atmel	Public company. Microcontrollers (ARM-based, 8051, AVR), automotive ICs, touch solutions.	3.3	\$3,560	USA
Jan-16	Bourns	Murata select potentiometer business assets	Trimming potentiometers (variable resistors) for circuit adjustments in a variety of electronic equipment.			Japan

Source: 451Group, PitchBook, press releases, public news articles.

## Select Semiconductor M&A Deals from Late 2013 through 2015

The following are brief summaries of some chip transactions from late 2013 through 2015 with a focus on the larger (mostly public) acquisitions:

- **TDK/Micronas** – In December 2015, Japan-based TDK agreed to acquire Micronas for about \$217 million. Micronas provided microcontrollers and sensors for automotive and industrial markets.
- **Micron/Inotera** – In December 2015, Micron agreed to acquire Inotera for about \$4 billion (Micron already owned about 35% of Inotera). Inotera was a Taiwanese supplier of DRAM chips and was originally formed as a joint venture between Nanya and Infineon.
- **Microsemi/PMC-Sierra** – In November 2015, Microsemi agreed to acquire PMC-Sierra for about \$2.5 billion. PMC-Sierra provided a broad range of mostly wireline communications chips.
- **ON Semiconductor/Fairchild Semi** – In November 2015, ON agreed to acquire Fairchild Semiconductor for \$2.4 billion. Fairchild sold a variety of analog and discrete components, with a focus on power management applications.
- **IDT/ZMD** – In October 2015, IDT announced it agreed to acquire ZMD for \$310 million. ZMD provided sensing and digital power semiconductor solutions, with a focus on automotive and industrial applications. It was based in Germany and privately-held.
- **Mellanox/EZchip** – In September 2015, Mellanox agreed to acquire EZchip for about \$811 million. EZchip provided networking processors for communications and was based in Israel.
- **Diodes/Pericom** – In September 2015, Diodes agreed to acquire Pericom Semiconductor for about \$400 million. Pericom provided a variety of mixed-signal chips.
- **Qualcomm/Ikanos** – In August 2015, Qualcomm agreed to acquire Ikanos for about \$47 million. Ikanos sold DSL chip solutions including G.fast chips.

- **Intel/Altera** – In June 2015, Intel agreed to acquire Altera for about \$16.7 billion. Altera was a major supplier of FPGAs and other programmable logic solutions.
- **Parade/Cypress TrueTouch Unit** – In June 2015, Cypress Semiconductor announced that it was divesting its TrueTouch mobile touch controller unit for \$100 million to Parade.
- **Uphill/ISSI** – In June 2015, Uphill Investment announced a deal to acquire ISSI (which sold niche memory chips) for \$765 million. It initially offered \$639.5 million a few months earlier but the price was increased due to subsequent offers from Cypress.
- **Microchip/Micrel** – In May 2015, Microchip agreed to acquire Micrel for about \$839 million. Micrel sold a broad range of analog components, with a focus on power management (LDOs, regulators).
- **Avago/Broadcom** – In May 2015, Avago agreed to acquire Broadcom (the resulting company was renamed Broadcom Ltd.). Broadcom was a leading supplier of integrated communications chips (Ethernet, Wi-Fi, set-top box chips, etc.). The deal size was about \$37 billion.
- **JAC Capital/NXP's RF Power Unit** – In May 2015, NXP agreed to divest its RF Power unit to Beijing-based JAC Capital for \$1.8 billion.
- **Hua Capital/OmniVision** – In April 2015, China-based Hua Capital agreed to acquire OmniVision, a major supplier of CMOS image sensors, for \$1.9 billion.
- **Knowles/Audience** – In April 2015, Knowles agreed to acquire Audience for about \$129 million. Audience provided audio chips and software for mobile devices.
- **Microsemi/Vitesse** – In March 2015, Microsemi agreed to acquire Vitesse for about \$389 million. Vitesse provided high speed wireline communications chips.

- **NXP/Freescale** – In March 2015, NXP agreed to acquire Freescale for about \$11.8 billion. Freescale sold a variety of chips including microcontrollers and embedded processors. It had been a leader in chips for the automotive market.
- **MaxLinear/Entropic** – In February 2015, MaxLinear agreed to acquire Entropic for about \$287 million. Entropic provided integrated solutions for set-top boxes and connectivity.
- **Intel/Lantiq** – In February 2015, Intel agreed to acquire Lantiq from Golden Gate Capital for \$383 million. Lantiq had previously been the wireline communications unit of Infineon.
- **Amazon/Annapurna Labs** – In January 2015, Amazon reportedly acquired Israeli-based Annapurna Labs for about \$360 million. Annapurna was developing an advanced networking chip.
- **Lattice/Silicon Image** – In January 2015, Lattice Semiconductor agreed to acquire Silicon Image for about \$600 million. Silicon Image was a major supplier of HDMI chips and IP, and other interface chips.
- **Cypress/Spansion** – In December 2014, Cypress agreed to acquire Spansion for about \$1.6 billion. Spansion was a provider of microcontrollers and Flash memory (mainly NOR Flash).
- **MACOM/BinOptics** – In November 2014, MACOM announced the acquisition of BinOptics for \$230 million. BinOptics sold indium phosphide based lasers for telecommunications applications.
- **Koch Industries/Oplink** – In November 2014, Koch Industries agreed to acquire Oplink for about \$445 million. Oplink developed optical communications components and modules.
- **MegaChips/SiTime** – In October 2014, Japan-based MegaChips announced the acquisition of private MEMS timing chip company SiTime for \$200 million.

- **Qualcomm/CSR** – In October 2014, Qualcomm agreed to acquire CSR for about \$2.5 billion. CSR was a leader in Bluetooth chips, but previously divested its mobile phone Bluetooth unit to Samsung.
- **Murata/Peregrine** – In August 2014, Japan-based Murata agreed to acquire Peregrine Semiconductor for about \$465 million. Peregrine provided silicon-on-insulator chips, primarily for wireless applications.
- **Intel/LSI's Axxia Unit** – In August 2014, Intel agreed to buy LSI's Axxia networking chip business unit for \$650 million.
- **Infineon/International Rectifier** – In August 2014, Infineon agreed to acquire International Rectifier (IRF) for about \$3 billion. IRF provided discrete and analog power components.
- **Vishay/Capella** – In July 2014, Vishay agreed to acquire Capella for about \$205 million. Capella was based in Taiwan and supplied ambient light sensors, optical encoders, and related components.
- **Inphi/Cortina** – In July 2014, Inphi agreed to acquire Cortina Systems for \$126 million. Cortina sold a variety of communications chips.
- **ON Semi/Aptina** – In June 2014, ON Semiconductor agreed to buy Aptina Imaging from TPG/Riverwood Capital for \$400 million. Aptina was a major supplier of CMOS image sensors.
- **Shanghai Pudong/Montage** – In June 2014, Shanghai Pudong agreed to purchase Montage Technology for \$693 million. Montage developed chips for set-top boxes and other applications.
- **Synaptics/Renesas SP Drivers** – In June 2014, Synaptics agreed to purchase Renesas' SP Driver unit for \$475 million.
- **Analog Devices/Hittite** – In June 2014, Analog Devices agreed to acquire Hittite Microwave for about \$2.45 billion. Hittite provided a variety of high-performance analog and mixed-signal chips.
- **Seagate/Avago's LSI Flash Unit** – In May 2014, Seagate announced that it agreed to acquire the LSI Flash Components and Accelerated Solutions units from Avago for \$450 million.

- **Microchip/ISSC** – In May 2014, Microchip agreed to acquire ISSC for about \$328.5 million. ISSC was publicly traded in Taiwan and provided Bluetooth chips and other related components.
- **Cirrus Logic/Wolfson** – In April 2014, Cirrus Logic agreed to acquire Wolfson for about \$488 million. Wolfson was based in Scotland and provided low power audio chips for mobile devices.
- **Exar/iML** – In April 2014, Exar agreed to acquire Integrated Memory Logic (iML) for about \$223 million. iML was publicly listed in Taiwan and provided power management chips for TFT LCD displays.
- **Cadence/Jasper** – In April 2014, Cadence agreed to acquire Jasper Design Automation, an EDA software company focused on verification, for \$170 million.
- **RF Micro Devices/TriQuint** – In February 2014, RF Micro Devices agreed to acquire TriQuint for about \$1.6 billion. Both companies supplied gallium arsenide components for applications such as cell phone power amplifiers. The combined company was renamed Qorvo.
- **Microchip/Supertex** – In February 2014, Microchip agreed to acquire Supertex for \$394 million. Supertex provided a variety of analog chips including high-voltage ultrasound components and telecom ICs.
- **Avago/LSI** – In December 2013, Avago agreed to acquire LSI for about \$6.6 billion. LSI provided a broad range of chips for storage and communications applications.
- **Tsinghua Holdings/RDA Micro** – In November 2013, Tsinghua agreed to acquire RDA Microelectronics for \$910 million. RDA sells a variety of power amplifiers and transceivers for cellular applications.
- **MACOM/Mindspeed** – In November 2013, MACOM agreed to acquire Mindspeed for \$272 million. Mindspeed provided analog communications, VoIP chips, and wireless infrastructure solutions (MACOM later divested the wireless infrastructure unit to Intel).



## Public Semiconductor M&A Premiums

---

### Semiconductor Public Company Premiums

This section includes a table of acquisitions of public semiconductor companies, including the premium paid by the buyer (typically relative to the target's stock price immediately prior to the announcement).

For some transactions, the premium is not clear. For example, in a number of cases, news of a potential deal was reported well before the actual deal announcement date. In these situations, the target's stock typically had a big run-up prior to the formal announcement, so the premium relative to the day before the announcement isn't relevant. In other cases, there was a bidding war among two or more buyers over a period of time. Depending on the situation, we excluded some of these or used the premium from when the news of a deal first emerged. As a few examples:

- **Microchip/Micrel** – In May 2015, Microchip announced an agreement to acquire Micrel for about \$839 million. This was only a 3% premium to Micrel's stock just prior to the announcement. However, it represented a 30% premium to when Starboard Value disclosed a 12% stake in Micrel and pushed for an exit several months earlier. Based on all the variables, we did not include this transaction in the table.
- **Microsemi/PMC-Sierra** – In November 2015, Microsemi reached an agreement to acquire PMC-Sierra. Skyworks had originally agreed to acquire PMC-Sierra, but was subsequently outbid by Microsemi. Microsemi paid about a 77% premium to where PMC-Sierra had originally traded just before the Skyworks announcement, but some time had passed. As such, we didn't include this deal.
- **ON Semi/Fairchild** – In November 2015, ON Semiconductor announced that it agreed to acquire Fairchild for approximately \$2.4 billion in cash. This was only a 12% premium to Fairchild's stock price, but reports about a potential Fairchild sale were reported a few weeks earlier. The premium to when these reports first emerged was 41%. Given the short period, we included the deal and the 41%.

- **Renesas/Intersil** – In September 2016, Renesas agreed to acquire Intersil for approximately \$3.2 billion. This was only a 14% premium to Intersil’s most recent stock price, but reports about the deal emerged in late August 2016. Relative to Intersil’s stock price just prior to these reports, the premium was about 44%. As such, we included the deal and the 44% premium.

There were a number of other deals that were excluded for related reasons (e.g., Marvell/Cavium, II-VI/Anadigics, ISSI’s acquisition, etc.).

### **Semiconductor Premium Table**

The following table summarizes the premiums for a number of public semiconductor acquisitions for the past few years. The median premium of these deals is 27% and the mean is 32%.

Announce				
Date	Acquirer	Target	Total Deal Amt. (\$M)	Premium
Dec-17	Silicon Laboratories	Sigma Designs	\$282	26%
Sep-17	Canyon Bridge Capital Partners	Imagination Technologies	\$744	42%
Aug-17	Littelfuse	IXYS	\$750	44%
May-17	Pixelworks	ViXS Systems	\$20	48%
Mar-17	MaxLinear	Exar	\$700	22%
Mar-17	Intel	Mobileye N.V.	\$15,300	35%
Feb-17	IDT	GigPeak (fka GigOptix)	\$250	22%
Dec-16	TDK	Invensense	\$1,300	20%
Dec-16	Teledyne Technologies	e2v technologies	\$780	48%
Nov-16	M/A-COM	AppliedMicro (AMCC)	\$770	15%
Nov-16	Siemens AG	Mentor Graphics	\$4,029	21%
Oct-16	Qualcomm	NXP Semiconductors	\$39,187	12%
Sep-16	Renesas Electronics	Intersil	\$3,219	44%
Sep-16	Supernova Investment	Actions Semiconductor	\$97	50%
Aug-16	TDK	Tronics	\$54	78%
Jul-16	Analog Devices	Linear Technology	\$14,880	24%
Jul-16	SoftBank	ARM Holdings	\$32,434	43%

Announce				
Date	Acquirer	Target	Total Deal Amt. (\$M)	Premium
Jun-16	ASML Holding	Hermes Microvision	\$3,087	17%
Jun-16	Cavium	Qlogic	\$1,360	14%
Jan-16	Microchip Technology	Atmel	\$3,560	52%
Dec-15	TDK	Micronas Semiconductor	\$217	63%
Dec-15	Beijing E-Town Capital	Mattson Technology	\$300	23%
Nov-15	ON Semiconductor	Fairchild Semiconductor	\$2,400	41%
Sep-15	Mellanox Technologies	EZchip Semiconductor	\$811	16%
Sep-15	Diodes	Pericom Semiconductor	\$400	40%
Jun-15	Intel	Altera	\$16,700	56%
May-15	Avago Technologies	Broadcom	\$37,000	28%
Apr-15	Hua Capital and Other Investors	OmniVision Technologies	\$1,900	18%
Apr-15	Knowles	Audience	\$129	-9%
Mar-15	Microsemi	Vitesse Semiconductor	\$389	36%
Mar-15	NXP Semiconductors	Freescall Semiconductor	\$11,800	7%
Jan-15	Lattice Semiconductor	Silicon Image	\$600	24%
Median				27%
Average				32%

Source: 451Group, PitchBook, press releases, news articles.

## Appendix: Select Private Semiconductor Companies

### Introduction

---

This Appendix includes an alphabetical list of the private companies highlighted in Chapter 2. This is not meant to be a comprehensive list of private chip companies as there are several hundred start-ups that could have been included. We tried to include a cross-section of different types of companies including some of the higher profile companies that have had large capital raises as well as a number of smaller less well-known start-ups. Similarly, we tried to include companies addressing a variety of different semiconductor segments. We note that we have included only publicly available information about these companies, even if we have met with management and have additional information. As noted in the Introduction chapter, we finalized the report contents in early January 2018, so company news items after that are not included.

### Brief Summaries of Select Private Companies

---

#### Alphabetical List of Private Companies Highlighted in the Report

The following is an alphabetical list of the private chip companies listed in this report. More details about each company are included in Chapter 2.

- **Acco** – Acco develops CMOS power amplifiers and RF front-ends for cellular standards for products such as smartphones and IoT devices. It has major offices in Paris and Silicon Valley.
- **Achronix** – Achronix provides high-performance FPGAs and embedded FPGA (eFPGA) IP, and indicated in mid-2017 that it would exceed revenue of \$100 million in 2017. It is based in Silicon Valley.

- **Active-semi** – Active-semi provides power management ICs and intelligent digital motor drive chips (that integrate analog, power, and ARM processors). It is based in Dallas, TX.
- **ActLight** – ActLight has developed dynamic photodetectors that translate light intensity into a time-delay which provides advantages over traditional photodiodes. It is based in Lausanne, Switzerland.
- **Adsantec** – Adsantec develops high-speed mixed-signal chips, including Mux/Demux ICs, amplifiers, drivers, TIAs, and other high performance analog components. It is based in Southern California.
- **Alereon** – Alereon develops UWB chips. It has recently had a focus on military (soldier body area networks and cable elimination) and high reliability medical applications. It is based in Austin, TX.
- **Ambiq Micro** – Ambiq Micro provides low power ARM-based microcontrollers which include a variety of proprietary technologies to reduce power consumption. It is based in Austin, TX.
- **Amimon** – Amimon has developed wireless communications chips and modules that are optimized for transmitting video, providing high quality and low latency. It is based in Silicon Valley and Israel.
- **Analog Bits** – Analog Bits licenses high-performance mixed-signal IP including SerDes (for PCI Express Gen 4, Ethernet, and more), clocks/PLLs, I/O, sensors, and memory. It is based in Silicon Valley.
- **AnDapt** – AnDapt provides unique power management chips that can be configured (varying voltage, current, impedance, etc.) similar to how FPGAs can be programmed. It is based in Silicon Valley.
- **Arralis** – Arralis develops high frequency millimeter wave ICs (as well as modules and antennas) that operate up to 110GHz for radar, space, and communications applications. It is based in Ireland.



- **Arbe Robotics** – Arbe develops differentiated high resolution “4D” imaging radar solutions for autonomous vehicles. Its solution includes both advanced software and its own unique chips. It is based in Israel.
- **Aspinity** – Aspinity develops ultra-low power analog signal processing (ASP) chips, which have advantages over digital processing in certain applications, such as voice processing. It is based in West Virginia.
- **Austemper Design** – Austemper is an EDA company that provides a tool-suite to analyze and verify the functional safety in SoCs and ASICs. Austemper is based in Austin, TX.
- **Autotalks** – Autotalks develops ARM-based chip solutions for V2X communications (802.11p) with a focus on autonomous driving applications. It is based in Israel.
- **Avalanche Technology** – Avalanche has developed STT-MRAM technology, based on proprietary perpendicular magnetic tunnel junction (pMTJ) cells. It is based in the Silicon Valley area.
- **Barefoot Networks** – Barefoot has developed a high-speed programmable Ethernet switching chip and software platform that enables software defined networking. It is based in Silicon Valley.
- **Butterfly** – Butterfly developed a single-chip ultrasound imaging solution, and announced its iQ handheld ultrasound scanner which is expected to ship in 2018. It has offices in Connecticut and New York.
- **Carbonics** – Carbonics has developed technology for improving performance and reducing power consumption of CMOS by adding small amounts of carbon nanotubes. It is based in Culver City, CA.
- **Cavendish Kinetics** – Cavendish develops MEMS-based RF switches for wireless applications, and its solutions have been incorporated into a variety of smartphones. It is based in Silicon Valley.

- **Celero** – Celero provides a variety of Wi-Fi chip solutions, including several types of 802.11ac and ax semiconductors (which include its DSP-based spectrum analyzer and software). It is based in Israel.
- **Cerebras** – Cerebras Systems is a stealth-mode start-up company. Reports indicate that it is working on an AI processor and related software. It is based in Silicon Valley.
- **Chronocam** – Chronocam has developed CMOS image sensors that are controlled based on what is happening in the scene, enabling ultra-high speeds and dynamic range. The company is based in Paris.
- **Codasip** – Codasip develops and licenses RISC-V processor IP, including a new 64-bit RISC-V processor. Its tools enable customized RISC-V extensions and processors. It is based in the Czech Republic.
- **cPacket** – cPacket has developed an algorithmic fabric chip (cVu) that enables packet inspection and network performance monitoring at full-line rate (including 100G). It is based in Silicon Valley.
- **CRC** – CRC has invented a unique technology for bonding ceramics with ceramics and metals, which enables enormous advantages for its semiconductor wafer pedestals. It is based in Silicon Valley.
- **Credo** – Credo provides high-speed SerDes and PHY chips and IP. It has 28G and 56G IP and has demonstrated 100G solutions. Its chips include a 16x56G CDR. It is based in Silicon Valley.
- **Crossbar** – Crossbar has developed a type of 3D Resistive RAM memory technology that it indicates is significantly faster than Flash with higher density. It is based in Silicon Valley.
- **Decawave** – Decawave develops UWB chips, primarily for indoor location and positioning, as GPS generally does not work indoors. It can locate objects to within 10cm. It is based in Dublin, Ireland.

- **D-Wave Systems** – D-Wave Systems develops quantum computing processors and computer system. Its 2000Q computer includes 2000 superconducting qubits. It is based in Vancouver, Canada.
- **EnerBee** – EnerBee has developed a MEMS-based energy harvesting generator and related power management chips. EnerBee indicates it works even at low speed/force. It is based in Grenoble, France.
- **Escape** – Escape Communications provides millimeter wave RF modems (e.g., a 10Gbps FPGA-based E-band modem) for markets such as backhaul and 5G. It is based in Torrance, CA.
- **Esperanto** – Esperanto develops energy-efficient processors based on the RISC-V standard. Its initial solution reportedly has over 4,000 cores and targets AI applications. It is based in Silicon Valley.
- **ESS** – ESS Technology develops analog components for audio applications. This includes high-end audio DACs, headphone audio amplifiers, and audio SoCs. It is based in Silicon Valley.
- **eVaderis** – eVaderis provides non-volatile memory IP for chip companies. This includes new memories such as embedded MRAM and Resistive RAM. It is based in Grenoble, France.
- **Exagan** – Exagan is developing advanced GaN-on-silicon high power components, including 600V and 1,200 volt fast-switching power devices. It is based in Grenoble, France.
- **Exalos** – Exalos provides superluminescent diodes. It historically focused on medical/industrial markets, but has recently developed visible light (RGB) SLEDs for AR/VR. It is based in Switzerland.
- **Fasttree3D** – Fasttree3D offers linear and matrix single photon detectors with integrated distance measurement for time-of-flight 3D cameras. It is based in Switzerland.

- **Ferric** – Ferric Semiconductor sells DC/DC power chip converters and IP that utilize a unique thin film inductor technology, enabling integrated power management. It is based in New York.
- **Flex Logic** – Flex Logic licenses embedded FPGA solutions that it indicates requires less area and fewer metal layers than regular FPGAs, reducing size and cost. It is based in Silicon Valley.
- **GaN Systems** – GaN Systems manufactures GaN-on-silicon high power transistors. Its patented “Island” topology improves isolation and performance. It is based in Silicon Valley (originally in Ottawa).
- **GEO Semiconductor** – GEO develops chips for image sensor video processing, with a focus on automotive applications. This includes automotive camera processor chips. It is based in Silicon Valley.
- **Graphcore** – Graphcore has developed a graph-focused intelligent processor unit (IPU) that is specifically optimized for AI/machine learning applications. It is based in Bristol, UK.
- **Greenwaves** – Greenwaves’ GAP8 chip is an ultra-low power IoT application processor that utilizes the RISC-V architecture, and integrates neural network acceleration. It is based in Grenoble, France.
- **HiLight Semiconductor** – HiLight provides TIAs and laser driver/combo chips for high-speed optical communications markets (10G Ethernet, PON, etc.). It is based in Southampton, UK.
- **HMicromicro** – HMicromicro has developed a highly reliable/redundant wireless chip solution for medical monitoring applications, as well as complete biosensor patch solutions. It is based in the Silicon Valley area.
- **IP Light** – IP Light develops advanced Optical Transport Network (OTN) processors for telecom and wireless infrastructure, and IP cores for mobile fronthaul applications. It is based in Israel.

- **IQ-Analog** – IQ-Analog has developed a family of configurable high-performance data converter chips (Next). It also offers high-speed data converter IP. It is based in San Diego, CA.
- **Kandou Bus** – Kandou has created a new signaling technology that enables very high performance, low energy SerDes, including 125Gbps/channel SerDes IP. It is based in Lausanne, Switzerland.
- **Keyssa** – Keyssa has developed technology that enables very high data rate wireless connections over short distances, with a focus on replacing connectors and internal cables. It is based in Silicon Valley.
- **Kinetic Technology** – Kinetic provides a variety of analog components including power management and LED/LCD drivers. It is a Cayman corporation with its major office in Silicon Valley.
- **LeddarTech** – LeddarTech develops LiDAR optical ranging and detection solutions (receiver chip and software), with a major focus on the automotive markets. It is based in Quebec City, Canada.
- **Leman Micro Devices** – Lemman has developed a sensor (for smartphones) that can accurately measure biometric parameters, including blood pressure and heart rate. It is based in Switzerland.
- **Libre Wireless** – Libre's MAVID is a highly integrated, low power hardware/software solution for embedded voice/audio applications (integrated Amazon Alexa). It is based in Silicon Valley.
- **Lime** – Lime Microsystems develops programmable transceiver chips (its 8001+ supports from 100kHz to 12GHz). It also has a family of FPGA-based software defined radios. It is based in Surrey, UK.
- **Lion Semiconductor** – Lion develops ultra-small, fast, and efficient power and battery management chips, which incorporate advanced digital control. It is based in San Francisco, CA.

- **Linear Dimensions** – Linear Dimensions provides custom and semi-custom analog and sensor solutions including biometric identification and health sensing solutions. It is based in Silicon Valley.
- **Maja Systems** – Maja has expertise in highly integrated CMOS MMICs. Its initial solution is a 60GHz CMOS radio, but it plans to address a variety of markets/frequencies. It is based in Silicon Valley.
- **mCube** – mCube develops MEMS accelerometers which can be fabricated on top of the drive chip to reduce cost and size. It also announced gyros and eCompass chips. It is based in Silicon Valley.
- **Menta** – Menta licenses embedded FPGA IP for SoCs and ASICs. Its Oragami software allows customers to customize the embedded FPGAs for their needs. It is based in Montpellier, France.
- **MEMSDrive** – MEMSDrive develops MEMS actuator chips, with a focus on image stabilization for cameras (including smartphones, automobile cameras, and AR/VR). It is based in Pasadena, CA.
- **Merus Audio** – Merus has created a family of audio amplifiers (eximo) that enables higher efficiency (less power consumption) and smaller size. It is based in Denmark.
- **Micralyne** – Micralyne is a MEMS semiconductor foundry with a strong focus on biomedical MEMS, industrial sensors, and telecom (optical MEMS). It is based in Edmonton, Canada.
- **MultiPhy** – MultiPhy sells high-speed (including 100G single wavelength) DSP-based communications chips with a focus on data center connectivity. It has offices in Israel and Silicon Valley.
- **Mythic** – Mythic has developed a unique architecture for AI/neural network processing, which can provide the power of a desktop GPU in a button size chip. It is based in Austin, TX and Silicon Valley.



- **NanGate** – NanGate has developed EDA software for IP libraries, including tools for creating, analyzing, migrating, validating and characterizing libraries. It is based in Silicon Valley and Denmark.
- **Nantero** – Nantero has developed carbon nanotube based non-volatile memory (NRAM) which it indicates is CMOS compatible and is faster than Flash with better endurance. It is based in Woburn, MA.
- **Navitas Semiconductor** – Navitas developed GaN power ICs that it indicates enables 100x increase in switching speed combined with a 40% reduction in energy savings. It is based in El Segundo, CA.
- **NextInput** – NextInput provides MEMS-based force-sensing solutions for touch applications including smartphones, wearables, and automotive applications. It is based in Silicon Valley.
- **Nitride Solutions** – Nitride Solutions has developed advanced aluminum nitride technologies for deep UV LEDs, thermal coatings, and sensors. It is based in Kansas.
- **Noveto** – Noveto's Sowlo technology can focus audio directly to a person's ears, such that only that person can hear it, using sensors that track a user's position. It is based in Israel.
- **Numascale** – Numascale has developed advanced node controllers that enable efficient scale-up of servers (large numbers of processors), greatly reducing cost of ownership. It is based in Oslo, Norway.
- **NVXL** – NVXL has created a disruptive architecture and software abstraction layer that enables scaling to thousands of FPGAs for AI and neural network applications. It is based in Silicon Valley.
- **Octasic** – Octasic has developed highly integrated base station SoCs and mini base stations. Separately, it has a family of multi-core DSP-based media gateway chips. It is based in Montreal, Canada,

- **Omni Design** – Utilizing its patented Swift technology, Omni provides ultra-low power analog IP cores, with a focus on advanced data converter IP. It is based in Silicon Valley.
- **Peratech** – Peratech has developed force touch sensing solutions (using Quantum Tunneling Composite technology) for touch controls. Automotive and consumer are target markets. It is based in the UK.
- **Plessey Semiconductor** – Plessey has developed unique GaN-on-silicon technology for high and mid power LED applications. It also has a family of LEDs for agricultural markets. It is based in the UK.
- **Pmd Tech** – Pmd has developed 3D time of flight CMOS image sensor chips, which provides depth and images, as well as complete 3D camera modules. It is based in Siegen, Germany.
- **Preciseley** – Preciseley Micro sells advanced MEMS-based optical switching solutions for telecom/datacom. It has also developed a tiny MEMS spectrometer. It is based in Edmonton, Canada.
- **Pyreos** – Pyreos provides MEMS-based infrared light sensors. It has a variety of products for motion detection, spectrometers, gas detection, and other applications. It is based in Edinburgh, Scotland.
- **Quantum Circuits** – Quantum is developing quantum computers based on superconducting devices. It was formed by quantum experts at Yale University. It is based in New Haven, CT.
- **Redpine** – Redpine offers wireless communications chips including Wi-Fi ICs, microcontrollers with integrated Wi-Fi, and 802.11p chips for automotive V2X communications. It is based in Silicon Valley.
- **S3 Semiconductor** – S3 Semiconductor designs and sells complex mixed-signal ASICs, primarily to OEM customers. It also provides analog IP to chipmakers. It is based in Dublin, Ireland.

- **Sckipio** – Sckipio offers advanced G.fast DSL semiconductor chips, including chipsets for both customer premises and distribution point units. It is based in Israel.
- **Semitech** – Semitech provides chips and IP for narrowband power line communications for industrial applications, home automation, and IoT products. It is based in Australia and Southern California.
- **Sensl** – Sensl develops highly sensitive silicon photomultipliers (SiPMs). Its solutions are used in medical equipment and it has recently expanded into the automotive sector. It is based in Ireland.
- **Sheba Microsystems** – Sheba has developed advanced MEMS-based actuator technology and its initial focus is MEMS for camera auto-focus and stabilization. It is based in Toronto, Canada.
- **SiBreeze** – SiBreeze provides power management chips with significantly reduced inductor values, providing smaller size, better performance, and thermal advantages. It is based in Silicon Valley.
- **SiFive** – SiFive is focused on embedded processor IP for the RISC-V architecture and its founders include several of the original engineers that created RISC-V. It is based in San Francisco, CA.
- **Silicon Line** – Silicon Line provides TIAs, laser drivers, and modules for high-speed optical communications (including HDMI, USB, and various optical standards). It is based in Munich, Germany.
- **Silicon Mobility** – Silicon Mobility's OLEA solutions (chips and software) control and optimize electric motors, batteries, and energy management of electric vehicles and hybrids. It is based in France.
- **SoundChip** – SoundChip develops audio signal processing platforms and chips that help optimize audio systems. It has significant expertise in audio, noise cancelling, and equalization. It is based in Switzerland.

- **Spin Transfer Technologies** – STT has developed Orthogonal Spin Transfer MRAM technology that it believes has lower power and faster performance than alternatives. It is based in the Silicon Valley area.
- **SureCore** – SureCore has developed and licenses ultra-low power SRAM IP for SoCs, ASICs, and FPGAs. It indicates its IP can reduce power consumption by up to 80%. It is based in the UK.
- **Telink** – Telink develops low power RF chips including low-power 2.4GHz RF SoCs for Bluetooth Smart (Low Energy), Zigbee, and other wireless standards. It is based in Shanghai, China.
- **Tempo** – Tempo Semiconductor sells analog/mixed-signal audio chips, including audio codecs and integrated audio SoCs. The company is based in Austin, TX.
- **Teramount** – Teramount's Photonic-Plug technology enables fast optical alignment of optical components to fiber cable, eliminating costly and time consuming active alignment. It is based in Israel.
- **ThinCi** – ThinCi is developing an advanced deep learning processor that is specifically optimized for vision processing applications. It is based near Sacramento, CA.
- **ThinkForce** – ThinkForce is developing an AI acceleration engine, with a focus on combining silicon processing chips that are optimized for AI with advanced AI algorithms. It is based in Shanghai, China.
- **ThinkSilicon** – ThinkSilicon provides GPU IP cores for embedded applications, including GPUs optimized for low power graphics and for neural network applications. It is based in Greece.
- **Tikitin** – Tikitin developed MEMS technology for the timing/clock market and indicates its MEMS resonator has significant advantages over quartz and other MEMS resonators. It is based in Finland.

- **Transphorm** – Transphorm develops GaN semiconductor chips, with a focus on high voltage GaN transistors (FETs) that operate at 600 volts or more. It is based in Goleta, CA.
- **Trigence** – Trigence has developed chips that can directly drive speakers with digital data (rather than converting to analog first), improving audio and reducing power consumption. It is based in Japan.
- **TriLumina** – TriLumina provides advanced VCSEL-based illumination modules for Flash LiDAR applications. Its solution integrates hundreds of VCSELs into a single chip. It is based in New Mexico.
- **UltraSoC** – UltraSoC develops IP that can be integrated into processors and SoCs to provide non-intrusive on-chip monitoring and analytics. It is based in Cambridge, UK.
- **USound** – USound has developed MEMS actuators for audio speakers (earphones, headphones, speakers) which enables a variety of advantages (small size and better audio). It is based in Austria.
- **Valens** – Valens is a major supplier of HDBaseT communications chips (uncompressed video/audio/Ethernet) with a recent focus on HDBaseT Automotive for car networking. It is based in Israel.
- **Vayyar** – Vayyar has developed an RF imaging chip that can create 3D images of objects, even through walls or other objects, without the use of cameras. It is based in Israel.
- **Vectrawave** – Vectrawave sells millimeter wave ICs and modules for wireless, satellite, and optical communications. Chips include LNAs, drivers, phase shifters, and clocks. It is based in France.
- **Vesper** – Vesper has developed piezo-resistive MEMS microphones that it indicates have significant advantages relative to conventional capacitive MEMS microphones. It is based in Boston, MA.

- **Videantis** – Videantis licenses vision/image processing and video-related processor IP and related software tools. Automotive and deep learning are recent focuses. It is based in Hanover, Germany.
- **VisIC** – VisIC develops GaN-based high-power transistors and modules. Its initial solutions operate at 650 and 1,200 volts, including a 1,200 volt power switch. It is based in Israel.
- **VocalZoom** – VocalZoom has developed advanced sensor solutions that can isolate a speaker's voice from the background noise, greatly improving voice and speech recognition systems. It is based in Israel.
- **Vorago** – Vorago has technology that enables high reliability semiconductors (high temperature/radiation operation). Its chips include ARM-based microcontrollers and SRAM. It is based in Texas.
- **XeThru** – XeThru has developed ultra-wideband (UWB) impulse radar technology that can sense sub-millimeter movement and presence. It is based in Oslo, Norway.
- **XMOS** – XMOS provides advanced voice/audio chip and software solutions (including a far-field voice controller) for the consumer electronics market. It is based in Bristol, UK.
- **Xsensio** – Xsensio has created a MEMS-based “Lab-on-a-chip” solution that analyzes tiny amounts of sweat on a person's skin for wearables/health monitoring. It is based in Lausanne, Switzerland.



## Notes



Menalto Advisors, LLC (“Menalto Advisors”) is a preeminent middle market technology M&A advisory firm. Investment banking services are provided by Menalto Advisors, a registered broker-dealer and member of FINRA and SIPC. Menalto Advisors provides financial advisory services related to mergers, acquisitions, private capital financings, and other similar transactions. Menalto Advisors is regulated by FINRA, but does not provide research or analyst reports or solicit or carry accounts for, or offer or sell securities products to, retail customers.

The information and opinions in this document have been prepared by Menalto Advisors. This document is provided for informational purposes only. It is not a research report, as such term is defined by applicable laws and regulations. Nothing contained herein is to be construed as an offer to buy or sell or a recommendation to buy or sell or a solicitation of an offer to buy or sell any type of financial instruments or securities. The information contained in this document should not be construed to represent an analysis of any company or security to form the basis for any type of investment decision. This document should not be construed to provide legal or tax advice, and you are encouraged to seek legal counsel for any legal or tax related questions or issues.

We believe the information contained herein to be reliable, but Menalto Advisors makes no representation as to the accuracy or completeness of the information. Menalto Advisors does not guarantee that it represents a complete or accurate summary of the available data. This document contains information from third party sources that we believe to be reliable, but that information has not been verified by Menalto Advisors to be accurate, and we therefore cannot be certain whether the information is accurate or complete.

Menalto Advisors or its bankers may have provided in the past, may currently be providing, or may in the future provide advisory services to one or more of the companies mentioned in this document.

The contents of this document are under copyright to Menalto Advisors. No part of this document may be copied or duplicated in any form, or redistributed, without Menalto Advisors’ prior written consent.

Copyright© 2018 Menalto Advisors, LLC. All rights reserved. Member FINRA, SIPC.

### **Menalto Advisors, LLC**

---

2550 Hanover St.  
Palo Alto, CA 94303  
+1-650-453-5815  
[info@menaltoadvisors.com](mailto:info@menaltoadvisors.com)