McKinsey on Semiconductors

Number 4,
Autumn 2014

Highlights include:

McKinsey on Semiconductors is written by experts and practitioners in McKinsey & Company's semiconductors practice along with other McKinsey colleagues.

To send comments or request copies e-mail us: McKinsey_on_Semiconductors@McKinsey.com.

Editor: Roberta Fusaro
Art Direction and Design: Leff Communications
Managing Editors: Michael T. Borruso, Venetia Simcock
Editorial Production: Runa Arora, Elizabeth Brown, Heather Byer, Torea Frey, Shahnaz Islam, Katya Petriwsky, John C. Sanchez, Sneha Vats

Cover illustration by Andrew Baker

McKinsey Practice Publications
Editor-in-Chief: Lucia Rahilly
Executive Editors: Allan Gold, Bill Javetski, Mark Staples

McKinsey on Semiconductors

Number 4, Autumn 2014

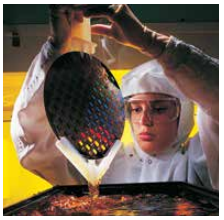

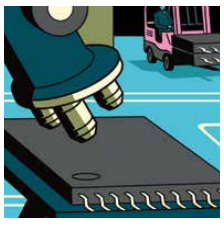

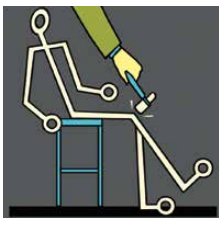

Introduction

Welcome to the fourth edition of *McKinsey on Semiconductors*. This year, we have expanded both the range and number of articles in the issue to provide a broader analysis of an industry that continues to be in transition.

In recent surveys, business leaders have told us that, above all, the Internet of Things has the potential to create new sources of growth in the semiconductor industry. We kick off this issue by reflecting on the progress this trend has made so far and, as the development and adoption of Internet of Things applications gains momentum, its potential implications for semiconductor companies and customers. Another article in the issue explores how trends in big data and connected consumer products could shape the market for microelectromechanical-systems technologies.

From there, we move to China and an update from our colleagues in Asia on developments in the fastest-growing semiconductor market in the world. As policy changes take hold, how will components manufacturers need to respond?

Continuing on a theme we established last year, we have put hardware and software development into focus with several articles—exploring first how semiconductor companies can improve their software-development capabilities and organization and second how they can address the cost and process challenges associated with verification of complex system-on-a-chip devices. We also take a quantitative look at R&D productivity in the industry, and we consider the evolution of advanced-packaging technologies and where companies are placing their next round of bets.

From an operations perspective, we pose several performance-related questions in this issue: What does excellence now look like in wafer production? And how can fab owners continue to reset the bar in performance by challenging established practices and implementing new data-centric techniques? Two articles address these questions—one is a case study of a fab that was able to realize significant process improvements through a lean approach, and the other offers a short perspective on how to rein in the costs of indirect materials.

Rounding out the issue is a consideration of how semiconductor companies can seek opportunities for growth beyond the core, examining their own supply chains for sources of innovation and market access.

This year, we are delighted and honored to also include executive perspectives from four highly respected business leaders—Joep van Beurden of CSR, Steven Mollenkopf and Murthy Renduchintala of Qualcomm, and Vincent Roche of Analog Devices. We thank each of them for sharing their time and insights.

McKinsey on Semiconductors is written, first and foremost, for industry executives who are passionate about their organizations' development and success. We hope that you find these perspectives helpful and a source for discussion and debate about where this industry is headed.

Harald Bauer
Director

Mark Patel
Principal

Nick Santhanam
Director

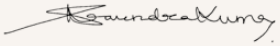

Florian Weig
Director

Bill Wiseman
Director

Andrew Baker

The Internet of Things:
Sizing up the opportunity

This connectivity trend is now recognized as a source of growth for semiconductor players and their customers. Here we consider the opportunities and constraints for components manufacturers.

Harald Bauer, Mark Patel, and Jan Veira

The semiconductor industry has been able to weather the fallout from the global financial crisis and realize several years of healthy growth—in part because of the widespread adoption of smartphones and tablets, which created demand for mobile and wireless applications. The industry's average annual growth rate between 2010 and 2013 was about 5 percent. Could the same sort of growth result from widespread adoption of the Internet of Things? Many semiconductor players have been asking themselves just this question.

The Internet of Things refers to the networking of physical objects through the use of embedded sensors, actuators, and other devices that can collect or transmit information about the objects. The data amassed from these devices can then be analyzed to optimize products, services, and operations. Perhaps one of the earliest and best-known applications of such technology has been in the area of energy optimization: sensors deployed across the electricity grid can help utilities remotely monitor energy usage and adjust generation and distribution flows to account for peak times and downtimes. But applications are also being introduced in a number of other industries. Some insurance companies, for example, now offer plans that require drivers to install a sensor in their cars, allowing insurers to adjust their premiums based on actual driving behaviors rather than projections. And physicians can

use the information collected from wireless sensors in their patients' homes to improve their management of chronic diseases. Through continuous monitoring rather than periodic testing, physicians could reduce their treatment costs by between 10 and 20 percent, according to McKinsey Global Institute research—billions of dollars could be saved in the care of congestive heart failure alone.

In each of these cases, the connected devices that transmit information across the relevant networks rely on innovations from semiconductor players—highly integrated microchip designs, for instance, and very low-power functions in certain applications. The semiconductor companies that can effectively deliver these and other innovations to original-equipment manufacturers, original-device manufacturers, and others that are building Internet of Things products and applications will play an important role in the development of the market. That market, in turn, may represent a significant growth opportunity for semiconductor players.

Indeed, semiconductor executives surveyed in June 2014 as part of our quarterly poll of the components-manufacturing market said the Internet of Things will be the most important source of growth for them over the next several years—more important, for example, than trends in wireless computing or big data. McKinsey Global Institute research supports that belief, estimating that the impact of the Internet of Things on the global economy might be as high as $6.2 trillion by 2025.[1] At the same time, the corporate leaders polled admit they lack a clear perspective on the concrete business opportunities in the Internet of Things given the breadth of applications being developed, the potential markets affected—consumer, healthcare, and industrial segments, among others—and the fact that the trend is still nascent.

In this article, we take the pulse of the market for Internet of Things applications and devices. Where along the development curve are the enabling technologies, and where can semiconductor players insert themselves in the evolving ecosystem? We believe components manufacturers may be able to capture significant value primarily by acting as trusted facilitators—it is their silicon, after all, that can enable not just unprecedented connectivity but also long-term innovation across the Internet of Things.

Sizing the opportunity

Three years ago, industry pundits and analysts predicted that, by 2020, the market for connected devices would be between 50 billion and 100 billion units. Today, the forecast is for a more reasonable but still sizable 20 billion or 30 billion units. This leveling off of expectations is in line with what we have seen in past introductions of new technologies. Throughout the late 1990s and early 2000s, for instance, there was much discussion in the semiconductor industry about the potential benefits and implications of Bluetooth technology, but the inflection point for Bluetooth did not happen until 2003 or 2004, when a large enough number of industry players adopted it as a standard and pushed new Bluetooth-based devices and applications into the market. The market for Internet of Things devices, products, and services appears to be accelerating toward just such an inflection point, based on four critical indicators.

Supplier attention. Internet of Things developer tools and products are now available. Apple, for instance, has released HealthKit and HomeKit developer tools as part of its latest operating-

system upgrade, and Google acquired Nest to catalyze the development of an Internet of Things platform and applications.[2]

Technological advances. Some of the semiconductor components that are central to most Internet of Things applications are showing much more functionality at lower prices. Newer processors, such as the ARM Cortex M, use only about one-tenth of the power that most energy-efficient 16-bit processors used only two years ago. This leap forward in technological capabilities is apparent in the evolving market for smart watches. The first such products released in 2012 boasted 400-megahertz single processors and simple three-axis accelerometers. Now a typical smart watch will include 1-gigahertz dual-core processors and high-end, six-axis devices that combine gyroscopes and accelerometers. Meanwhile, the prices of the chip sets used in these products have declined by about 25 percent per year over the past two years.

Increasing demand. Demand for the first generation of Internet of Things products (fitness bands, smart watches, and smart thermostats, for instance) will increase as component technologies evolve and their costs decline. A similar dynamic occurred with the rise of smartphone usage. Consumer demand for smartphones jumped from about 170 million devices sold per year just four or five years ago to more than a billion devices in 2014. The increase in orders coincided with a steep decline in the price of critical smartphone components.

Emerging standards. Over the past two years, semiconductor players have joined forces with hardware, networking, and software companies, and with a number of industry associations and academic consortiums, to develop formal and informal standards for Internet of Things applications. AT&T, Cisco, GE, IBM, and Intel, for instance, cofounded the Industrial Internet Consortium, whose primary goal is to establish interoperability standards across industrial environments so that data about fleets, machines, and facilities can be accessed and shared more reliably. Other groups have been focused on standardizing the application programming interfaces (APIs) that enable basic commands and data transfer among Internet of Things devices.

Implications for semiconductor players

Analysts have predicted that the installed base for Internet of Things devices will grow from around 10 billion connected devices today to as many as 30 billion devices by 2020—an uptick of about 3 billion new devices per year. Each of these devices will require, at a minimum, a microcontroller to add intelligence to the device, one or more sensors to allow for data collection, one or more chips to allow for connectivity and data transmission, and a memory component. For semiconductor players, this represents a direct growth opportunity that goes beyond almost all other recent innovations—with the exception, perhaps, of the smartphone.

A new class of components will be required to address this opportunity: system on a chip–based devices produced specifically for the Internet of Things, with optimal power and connectivity features and with sensor integration. First-generation chips are already on the way, although it will likely be a few generations before they can deliver all the functionality required. Intel, for instance, is releasing a low-power system on a chip designed for smaller products in automotive and industrial environments that also can be used in fitness bands and other wearable devices. Additionally, sensors based on microelectro-mechanical-systems (MEMS) technology will continue to play a significant role in enabling

Internet of Things applications (see "How big data and connected consumer products could boost the market for MEMS technology," page 31).

It's worth noting that semiconductor players may also be able to profit indirectly from the Internet of Things, since the data generated from billions of connected devices will need to be processed—all those "little" data must be turned into big data—and users will require greater storage capacity, spurring new demand for more servers and more memory. Building on an existing market, semiconductor companies can continue to provide the critical devices and components that are at the heart of these products.

The question, then, is no longer if the Internet of Things can provide substantial growth for semiconductor players; the real consideration is how best to capitalize on the trend. What are the critical challenges or inhibitors? What are the possible enablers for growth and adoption? Based on our research and discussions with semiconductor executives, we have identified potential challenges in two critical areas—technology and ecosystem development.

The technological challenges

Semiconductor players may need to invest heavily to adapt their chip designs and development processes to account for specific Internet of Things system requirements. For instance, because many applications would require devices that are self-sustaining and rely on energy harvesting or long-life batteries, semiconductor companies must address the need for optimal power consumption and outstanding power management in their products. Connectivity load will be another critical concern given that hundreds or even thousands of devices may need to be connected at the same time. The average smart home, for instance, may contain 50 to 100 connected appliances, lights, thermostats, and other devices, each with its own low-power requirements. Existing connectivity solutions such as standard Bluetooth or Wi-Fi will likely not be able to meet smart-home requirements given their power and network limitations.

Manufacturers may also need to emphasize flexible form factor to a greater degree than they currently do. Components must be small enough to be embedded in today's smart watches and smart glasses but also amenable to further shrinking for incorporation into still-unidentified future products. And security and privacy issues absolutely must be addressed. Internet of Things devices will not be used for critical tasks in, say, industrial or medical environments if connectivity protocols have not been established to prevent hacking, loss of intellectual property, or other potential breaches.

Semiconductor players are moving full steam ahead to address some of these challenges. Their efforts in two areas in particular are highly encouraging.

Increased integration. Some semiconductor players are already considering investing in new integration capabilities—specifically, expertise in packaging and in through silicon via, a connectivity technique in electronic engineering, as well as in software development. The emergence of more integrated system-in-package and system-on-a-chip devices is helping to overcome some of the challenges described earlier, in part by addressing power, cost, and size factors. The trend toward multidimensional chip stacking and packaging (2.5-D and 3-D integrated-circuit, or 2.5DIC and 3.0DIC, devices in particular) has resulted in integrated circuits that are one-third smaller than standard chips, with 50 percent lower power consumption and bandwidth that is up to eight times higher—at a cost that can be up to 50 percent lower when compared with

traditional systems on a chip of the same functionality. Monolithic integration of MEMS sensor technologies with complementary metal-oxide semiconductors is considered unlikely for Internet of Things applications. In these instances, the integration of substrates with silicon requires making certain design trade-offs and optimizing both the sensor and the logic circuits. Instead, we expect to see 2.5DIC and 3.0DIC technologies being favored for Internet of Things–specific integrated circuits.

Connectivity standards. The current cellular, Wi-Fi, Bluetooth, and Zigbee specifications and standards are sufficient to enable most Internet of Things applications on the market. Some applications, however, will require low-power, low-data-rate connectivity across a range of more than 20 meters—an area in which cellular technologies and Wi-Fi often fall short. New technologies that target this need are emerging from players such as those in the Bluetooth and Weightless interest groups. The latter is an industry group comprising technology companies that are exploring the use of free wireless spectrum to establish an open communications protocol. Such standardization efforts will enable Internet of Things applications that require broadly distributed sensors operating at low power over low-cost spectrum—for instance, temperature and moisture sensors used in agricultural applications.

The ecosystem challenges

As Joep van Beurden, the chief executive at CSR, notes, only about 10 percent of the financial value to be captured from the Internet of Things trend is likely to be in the "things"; the rest is likely to be in how these things are connected to the Internet (see "Making connections: Joep van Beurden on semiconductors and the Internet of Things," page 10). The semiconductor players that focus primarily on the things themselves should therefore find ways to support the development of a broader ecosystem (beyond silicon) and find their niche as both enablers and creators of value for their customers and their customers' customers. This will mean developing partnerships with players further downstream, such as companies that are building and providing cloud-based products and services.

It will be important for semiconductor companies to remember that different industries are at different levels of maturity and complexity with respect to the Internet of Things—so the roles that components manufacturers can play in application development in certain industries will vary, as will the timing of growth opportunities. The market for home-automation tools, for instance, has established some common APIs, but competing standards remain. A number of application developers have already started generating monitoring products for consumers, and once standardization issues can be addressed, the market may experience significant growth rather quickly. By contrast, the markets for monitoring and control systems in factories and for beacon technologies in retail are much more fragmented and will therefore take longer to develop. In retail, for instance, all the players in the value chain—the stores, the data aggregators, the Internet service providers, and other partners—must sort out their roles and standards of operation before beacon-technology providers can approach them with a clear customer value proposition and business model.

In these instances, semiconductor companies may want to test the waters by forming alliances with hardware companies, systems players, and customers or by finding ways to assist in standards development. In the factory-monitoring-systems

market, for instance, players are attempting to create common standards (through the Industrial Internet Consortium initiative, for example, and the Europe-only Industry 4.0 initiative), even though most of the hardware platforms are still proprietary, as are the data, which reside in legacy systems. Semiconductor players that pursue alliances and standard-setting activities may be able to play an enabling role in defining best practices in Internet of Things privacy, security, and authentication—issues that will be critical in markets such as healthcare and wearables that are dealing with sensitive consumer data.

Given the potential 90 percent distribution of value to players who provide all the technologies "beyond" the silicon, there may never be a compelling enough business case for components manufacturers to develop individual chips and systems for hundreds of thousands of discrete Internet of Things industry applications. We believe semiconductor players should instead design a family of devices that are sufficiently flexible to cater to the needs of multiple industries—that can be used in industrial *and* consumer Internet of Things applications that boast similar characteristics. Our work suggests that these devices will likely fall somewhere along a continuum of application requirements—at one extreme, high-power, high-performance, application-processing Internet of Things devices, such as those embedded in smart watches, and, at the other extreme, low-cost, ultralow-power integrated sensors that support sufficient (but not excessive) functionality and autonomous device operation. To achieve this level of design flexibility and to properly address the opportunity, semiconductor players may need to rethink their approach to product and application development.

• • •

The challenges associated with the Internet of Things are many; semiconductor executives should consider ways to integrate new development models, process capabilities, and go-to-market strategies in their existing operations. Success will require bold moves, boards that are willing to bet on unfamiliar models and activities, and collaboration with those that are developing industry standards. But the semiconductor industry should embrace this era of innovation and reinvention. The opportunities for growth outweigh the challenges, as components manufacturers explore the creation a new class of Internet of Things–enabled semiconductors that can cut across a wider swath of potential customers than existing components can. The sector may be on the cusp of unit growth similar to the surge it experienced with the smartphone—and perhaps an even greater jump.

[1] For more, see *Disruptive technologies: Advances that will transform life, business, and the global economy*, McKinsey Global Institute, May 2013, on mckinsey.com.
[2] Aaron Tilley, "Google acquires smart thermostat maker Nest for $3.2 billion," *Forbes*, January 13, 2014, forbes.com.

Harald Bauer (Harald_H_Bauer@McKinsey.com) is a director in McKinsey's Frankfurt office, **Mark Patel** (Mark_Patel@McKinsey.com) is a principal in the San Francisco office, and **Jan Veira** (Jan_Veira@McKinsey.com) is an associate principal in the Munich office.

Andrew Baker

Making connections: Joep van Beurden on semiconductors and the Internet of Things

The CEO of CSR discusses the progress and growing pains of the Internet of Things market.

Mark Patel and Jan Veira

Semiconductor executives are closely monitoring the development of the Internet of Things—in which physical objects are equipped with sensors and other devices that allow them to share and receive data through a network. Examples of applications in this area include smart watches, fitness bands, and home- and industrial-automation tools. Some are predicting a multitrillion-dollar market opportunity. Joep van Beurden, chief executive officer of CSR, a fabless semiconductor company that produces wireless technologies, agrees but notes that the Internet of Things still hasn't reached its tipping point. "I don't think it has been overhyped by any means. I just think widespread adoption will happen later than we expected," he says. In this edited conversation, Mr. van Beurden discusses growth in the Internet of Things market and the implications of this connectivity trend for semiconductor companies.

McKinsey on Semiconductors: *How would you assess growth in the market for Internet of Things applications relative to the industry's expectations?*

Joep van Beurden: In relative terms, you might say the growth is impressive, but the base is still very small. In absolute size, the market for Internet of Things applications is much smaller than what everyone predicted three or four years ago. Of course, the same sort of thing happened with Bluetooth development in the late 1990s:

every year analysts predicted we would see a significant increase in Bluetooth-enabled devices, and every year it didn't happen—until the early 2000s, when Bluetooth was adopted by leading cell-phone manufacturers and the technology took off. We're all still waiting for that inflection point with the Internet of Things.

McKinsey on Semiconductors: *Which applications do you see as having the greatest promise?*

Joep van Beurden: It's hard to predict. Everyone was initially excited about the promise of wearables but, looking at that market a year and a half later, the uptake has been relatively slow—certainly nowhere near the 100-million-plus device market required to bring the Internet of Things to scale. Many companies are hedging their bets across different industries and shipping reference designs and development kits to a variety of players, small and large. Those players are working on innovative ideas in home automation, medical devices, automotive, and other industries. But it's been a struggle to identify the one Internet of Things application that is going to take off.

McKinsey on Semiconductors: *What is inhibiting growth in the Internet of Things today?*

Joep van Beurden: A lot of analysts have evaluated the potential financial value that Internet of Things applications may create over the next five to ten years—it's a $300 billion or $15 trillion opportunity, depending on whom you listen to. When you drill down, however, you see that about 10 percent of this value is created by the "things," while 90 percent comes from connecting these things to the Internet. The Internet of Things is not just about storing information in the cloud; the data only become interesting when you combine them with sensors and analytics. But a certain degree of alignment must happen for those connections to take place and for the Internet of Things to take off. The industry must adopt common standards and business models, and it must address issues relating to privacy and security.

Getting alignment in all these areas is easier said than done. Consider connectivity efforts in healthcare. Having an Internet of Things–based ecosystem in which medical information is stored in the cloud and accessible by individuals and healthcare professionals from anywhere in the world looks good on paper. But the multiple hospitals and healthcare organizations involved will likely use different protocols for exporting information into the cloud. And not all medical institutions and individuals may be interested in sharing their information. There needs to be alignment on how to collect information and from whom, how to port it to the cloud, how to encrypt it, who will access it and how, and so on.

We are not in that aligned world today. It will happen eventually, because the prize is so large, but it will take time.

McKinsey on Semiconductors: *Semiconductor players are quite far down in the Internet of Things application stack. How important will this network be for semiconductor growth in the coming years?*

Joep van Beurden: Cost and power improvements from semiconductor players will come in time, once a killer application is introduced and achieves the 100-million-plus devices mark. Then highly integrated, cost-efficient devices will be possible. However, it is more important for semiconductor players to recognize that the "things" themselves—the chips they produce—are

Joep van Beurden

Education
Holds an MS in applied physics from the University of Twente

Career highlights
CSR (2007–present)
CEO

NexWave (2004–07)
CEO

Fast facts
Served as chairman of the Global Semiconductor Alliance from 2011 to 2013; previously served as vice chairman of the organization

Began his career as a lecturer in physics and electronics at the University of Zambia

not going to be the game changers. Sure, they may add $30 billion in new revenue through Internet of Things applications, and that would be great, but it will not significantly change the dynamics of the semiconductor industry. We are a low-growth industry, and that is not going to change by selling a few more "things."

McKinsey on Semiconductors: *What role can semiconductor players have in driving Internet of Things adoption? Will this require a change in designers' and manufacturers' business models?*

Joep van Beurden: A critical challenge for semiconductor players will be how to capture more than the 10 percent of value from the things while not stepping too far into uncharted territory—for instance, exploring business models that you have limited capabilities in. It's a fine line. Semiconductors should not become services companies; they need to look instead at where the silicon and the Internet intersect and find ways to enable that connection. For instance, we acquired Reciva, a cloud-based streaming audio aggregator. The company does not offer streaming music or online radio stations; it provides the API layer that allows consumers to get content from streaming music services and online radio stations seamlessly. By enabling the silicon, Reciva allows consumers to access the data in the cloud and do things with them that make the streaming services or online radio stations that are part of its network more valuable.

This sort of enabling model can provide an opportunity for semiconductor players to have their say in standards development. It can also become a nice stepping stone toward larger application markets in industries where the value chain is not as well developed. In retail, for instance, many companies are just now exploring the use of beacon technology—a category of low-power, low-profile transmission devices that can help retailers

provide personalized services to shoppers. The projected market value of the beacons themselves is $60 million a year—a nice figure but not one that will be game changing for my company or others in the semiconductor industry. But because of the information the beacons can provide—what are people buying, and how much?—they will hold a value far greater than $60 million for the retailers that use them. The question is, how do we insert ourselves into that value chain?

McKinsey on Semiconductors: *Should incumbent semiconductor players feel threatened by the Internet of Things? Are they taking enough risks to innovate?*

Joep van Beurden: The short answer is no, they shouldn't feel threatened, and they don't need to take enormous risks. But it is worth noting that over the past ten years we have not grown this industry in any significant way. Every semiconductor CEO is looking for growth in a nongrowth industry. For players in the traditional semiconductor market, the Internet of Things may spark some growth, but it certainly will not change 2 percent industry growth today to the 10 to 15 percent growth we had in the 1980s. In this case, the goal of innovation is not about developing and selling more things; it's an opportunity to rethink the business model—just a little, not in any radical way—and try to create value within the cloud.

McKinsey on Semiconductors: *Are investors and shareholders supportive of the semiconductor industry making bets and experimenting in the Internet of Things?*

Joep van Beurden: It's still too early to tell, but investors and shareholders have generally been supportive; they are certainly not impeding progress to this point. There has also been a lot of interest from the venture-capital industry, although most of it has been focused on pure-play cloud companies. In our industry, there are quite a few hardware Internet of Things players trying to achieve the lowest prices, or power, or what have you for specific Internet of Things applications. Personally, I don't believe that will be the way to create significant additional value; it is a race to the bottom. For me the answer lies in connecting the hardware in a smart way to the cloud, not just in making chips smaller, lower in power, and lower in cost.

Mark Patel (Mark_Patel@McKinsey.com) is a principal in McKinsey's San Francisco office, and **Jan Veira** (Jan_Veira@McKinsey.com) is an associate principal in the Munich office.

Bill Butcher

Semiconductors in China:
Brave new world or same old story?

Will China become home to a world-class semiconductor industry, or will Chinese semiconductor companies continue to pursue global players?

Gordon Orr and Christopher Thomas

Executives of global semiconductor companies have had their eyes on China for many years, primarily as a customer-rich end market and a source of innovation. But now they will need to take an even closer look. Government stakeholders in China have been reconsidering the risk posed by the country's heavy reliance on others for semiconductor components and capabilities, and they are carrying out policy changes that could correct for this dependence. Pair these policy efforts with private-market forces that are slowly but surely strengthening the capabilities of mainland semiconductor companies and multinational chip makers competing in China will likely face a very different operating environment—one with new risks and opportunities.

What's changing?

China is by far the largest consumer of semiconductors; it accounts for about 45 percent of the worldwide demand for chips, used both in China and for exports. But more than 90 percent of its consumption relies on imported integrated circuits. Integrated-circuit companies in China entered the semiconductor market late—some two decades after the rest of the world—and have been playing catch-up ever since in an industry in which success depends on scale and learning efficiencies. The Chinese government made several attempts to build a local semiconductor industry, but none really took hold. Now, however, things are changing on both the business and policy fronts.

Low-cost smartphones designed in China are flooding the market. For instance, Android phones designed in China now represent more than 50 percent of the global market, compared with their negligible presence five years ago. Lenovo's significant deals early in 2014—first acquiring IBM's low-end x86-based server business for $2.3 billion and then buying Motorola from Google for almost $3 billion—further suggest that the customer base for hardware is moving to China. Meanwhile, Beijing and Shenzhen have become innovation hotbeds for wearable devices and other connected consumer electronics. Technology companies in these regions are not trailing others in this area of innovation; they are running neck and neck with other early entrants.

Multinational corporations in every industry—from automotive to industrial controls to enterprise equipment—are increasingly establishing design centers on the mainland to be closer to customers and benefit from local Chinese talent. McKinsey's proprietary research indicates that more than 50 percent of PCs, and between 30 and 40 percent of embedded systems (commonly found in automotive, commercial, consumer, industrial, and medical applications), contain content designed in China, either directly by mainland companies or emerging from the Chinese labs of global players. As the migration of design continues, China could soon influence up to 50 percent of hardware designs globally (including phones, wireless devices, and other consumer electronics).

Fabless semiconductor companies are also emerging in China to serve local customers. For instance, Shanghai-based Spreadtrum Communications, which designs chips for mobile phones, and Shenzhen-based HiSilicon Technologies, a captive supplier to Huawei and one of the largest domestic designers of semiconductors in China, are among the local designers that have shown rapid growth over the past few years.

There has been slower but steady progress among local foundries. For reasons including costs and scale—and, in some cases, export controls—these players traditionally have been reluctant to invest in cutting-edge technologies, always lagging three or four years behind the industry leaders. But the performance gap is shrinking. As global players such as Samsung, Taiwan Semiconductor Manufacturing Company, and Texas Instruments set up shop in China, leading local foundries such as Shanghai Huali Microelectronics Corporation, SMIC, and XMC are poised to benefit from the development of a true technology cluster. At the same time, fewer and fewer chip designs will be moving to technologies that are 20 nanometers and below; following Moore's law is becoming too expensive and is of limited benefit to all but a small set of global semiconductor companies. As a result, low-cost, lagging-edge Chinese technology companies will soon be able to address a larger part of the global market.

A market-based policy effort

The Chinese government is now putting significant funding and effort behind new policies relating to the development of the semiconductor industry. The government's previous attempts to build the industry, dating all the way back to the 1990s, had mixed results because funding plans and incentives were focused more on research and academia than on business. Additionally, investments were fragmented—at one point, the government had invested in 130 fabrication sites across more than 15 provinces, none of which was able to capitalize on the scale and scope of its neighbors' sites, and supporting industries never materialized.

A different type of task force

The Chinese government has convened a task force whose composition and oversight differs markedly from previous groups charged with building a strong domestic semiconductor industry.

The task force includes four important ministries that operate under the State Council of the People's Republic of China. They are the Ministry of Industry and Information Technology, which takes the lead on formulating industrial strategies, policies, and standards; the Ministry of Science and Technology, which drafts policies and plans relating to scientific-research programs and institutions; the Ministry of Finance, which validates the proposed investment plan and assesses it for risk; and the National Development and Reform Commission, which monitors the overall process and reviews the policy draft.

What's different this time, however, is that the task force includes the top 10 to 15 leaders in China's semiconductor industry (convening executives from fabless designers, foundries, and equipment manufacturers) and overarching leadership for the project from Vice Premier Ma Kai, one of the government's highest-ranking officials.

This committee had a direct influence on the State Council during its drafting of the Guideline of the National IC Industry Development Promotion, the high-level policy framework that was shared publicly in June 2014.

The government, realizing that earlier bureaucrat-led investment initiatives failed to bring the desired results, is now aiming to take a market-based investment approach. In this case, decisions about allocating for-profit investment funds will be managed by professionals but will remain aligned with the government's policy objectives. Chinese officials have convened a unique task force charged with setting an aggressive growth strategy (see sidebar, "A different type of task force"). This group helped develop a policy framework that is targeting a compound annual growth rate for the industry of 20 percent between now and 2020, with potential financial support from the government of up to 1 trillion renminbi ($170 billion) over the next five to ten years. Investments will be made by a national investment vehicle (the National Industry Investment Fund) and provincial-level entities. These entities will invest across multiple categories, including project finance and domestic and foreign acquisitions, as well as traditional research and development subsidies and tax credits.

To avoid the fragmentation issues of the past, the government will focus on creating national champions—a small set of leaders in each critical segment of the semiconductor market (including design, manufacturing, tools, and assembly and test) and a few provinces in which there is the potential to develop industry clusters. For instance, SMIC, a leading foundry headquartered in Shanghai, is building a 300-millimeter fab in the Beijing Economic and Technological Development Area. The company signed cooperation agreements with the national and local governments and announced a joint investment of $1.2 billion.

Investors include the Beijing Municipal Commission of Economy and Information Technology, the Institute of Microelectronics of Chinese Academy of Sciences, and the Beijing city government.

The Chinese government has actively pursued consolidation to spur the creation of national champions. For instance, Tsinghua Unigroup, a state-owned enterprise, recently bought two of the top four Chinese fabless companies—in 2013, it acquired Spreadtrum for $1.7 billion and RDA Microelectronics for $0.9 billion—and aims to combine them into a single entity. The new policy framework specifically encourages consolidation within China's assembly-and-test market segment.

Implications for semiconductor players

China released the high-level framework for its new national semiconductor policy in June 2014; the details and the long-term effects of its new approach to developing the industry remain to be seen. Will it lead to a world-class semiconductor industry, or will Chinese semiconductor companies continue to lag behind global players? Three medium-term effects seem likely.

Pressure for localization will increase. China's strong desire for national champions may further tilt the system in favor of local players. According to industry estimates, Chinese original-equipment manufacturers will design more than half of the world's phones in 2015.[1] Under the national-champions model, they may be encouraged to take advantage of domestic suppliers' low-cost strategies and strong local technical support. Additionally, in the wake of global data-privacy and security concerns, there has been even more of a push from the Chinese government for state-owned and private enterprises to purchase from local system suppliers (which, in turn, are more likely to source from local semiconductor vendors).

More partnership opportunities will arise for second-tier players. Many of the Chinese government's previous policies have not offered opportunities for global players to benefit. However, government leaders in China's semiconductor sector are now beginning to realize that the country needs to partner with global technology companies to improve the local talent base and supply chain. As a result, they are more open than ever to win-win engagements between global players and national champions. For their part, top-tier multinational semiconductor companies traditionally have had less incentive to share their intellectual property or transfer technology to China. As such, second-tier players may fare better in this evolving ecosystem since they have less to lose than global giants—and everything to gain. In the winner-takes-all semiconductor markets, these players may benefit from their Chinese partners' deep pockets, becoming better able to match the investments of market leaders.

Chinese companies will become more aggressive in pursuing international mergers and acquisitions. Indeed, it would be quite difficult for Chinese players to build a complete and competitive semiconductor value chain without capitalizing on foreign assets; collaborations between Chinese and global players probably will not be enough to meet the country's objectives. We should expect China to continue to actively seek opportunities to acquire global intellectual property and expertise, usually with the intent of transferring them back home. What's still to be determined, however, is how global governments will react to proposed deals in light of the emerging policy and market changes.

How should multinational players respond?

Most global semiconductor players have invested heavily in their Chinese operations over the years,

In China and elsewhere, government intervention in the semiconductor market has been a mixed bag—some successes, some missed opportunities. But the Chinese government is better positioned than most to make a big policy bet.

but many are still operating below their potential, especially in functions beyond sales and marketing. Considering the emerging policy and business trends we've just discussed, we believe it's a good idea for leaders to inventory their company's current position in China.

This process should start with the most timely and immediate concern—the potential effects of changing Chinese policy. Questions for reflection might include: How will you align your operations with the Chinese government's new plans? Are your relationships in China strong and deep enough to provide you with some warning of potential risk as a result of domestic-policy changes? Do you have an early sense of what those risks might be, and a rapid-response plan to address them? Could you gain advantage by approaching the government with a win-win idea?

For multinational companies operating in China, it is impossible to separate political and regulatory concerns from business—which is why it is also necessary for leaders to take stock of the overall market and the capabilities they bring to the table.

Market-level questions might include the following: Given the different buying factors and supplier-management philosophies of Chinese customers, do you still have a winning product road map? Can you respond to the emerging needs of customers based in China as fast as a local company can? Have you followed your global customers as they set up design centers on the mainland? Which Chinese champions are emerging, and which markets will they attack?

Capabilities-level questions might include the following: How are you leveraging Chinese manufacturing and design talent to win in China—or to win globally? Are your leaders in China as strong and empowered as they are in your home region? Do your global leaders have enough connections in, experiences with, and insights about the Chinese market? How robust is your talent pipeline in China? Can you act as "one company" in the country, or do organizational silos prevent collaboration across the sales, product-development, government-relations, and manufacturing functions?

There is no one right answer to any of these questions; depending on its role and standing in the market, every company faces its own unique challenges in China. Accordingly, we have seen leading semiconductor companies adopt a number of different approaches. Some have taken the initiative to develop R&D capabilities in China, designing chips and applying for patents locally. Others have consolidated all their activities (sales,

marketing, and operations, for instance) under a China CEO who reports directly to the global CEO. One company created an advisory board of senior global executives dedicated entirely to coordinating and pushing the China agenda. Other companies have taken a talent-first approach—for instance, promoting a former China head to a global executive position to add China expertise to the boardroom and soliciting personal commitments from the CEO to visit the country every few months to review status and remove organizational barriers.

• • •

In China and elsewhere across the globe, government intervention in the semiconductor market has been a mixed bag—some successes, some missed opportunities. Still, the Chinese government is better positioned than most to make a big policy bet, with its massive customer and installed-manufacturing base, its deep bench of engineering talent, and its financial resources. It can afford to be patient, confident that macroeconomic forces make its hand incrementally stronger every year.

If the government follows through on its policy intent and steers substantial investment and support toward the domestic semiconductor market over the next decade, it will prompt global players to make their own moves—whether forging new and different partnerships with Chinese players, managing overcapacity in critical segments, or developing complementary or competitive policies of their own.

Whether this policy is ultimately effective or not, its impact will be felt across the industry.

[1] Ian Mansfield, "Chinese phone manufacturers expected to take half the market in 2015," *Cellular News*, March 10, 2014, cellular-news.com.

Gordon Orr (Gordon_Orr@McKinsey.com) is a director in McKinsey's Shanghai office, and **Christopher Thomas** (Christopher_Thomas@McKinsey.com) is an associate principal in the Beijing office.

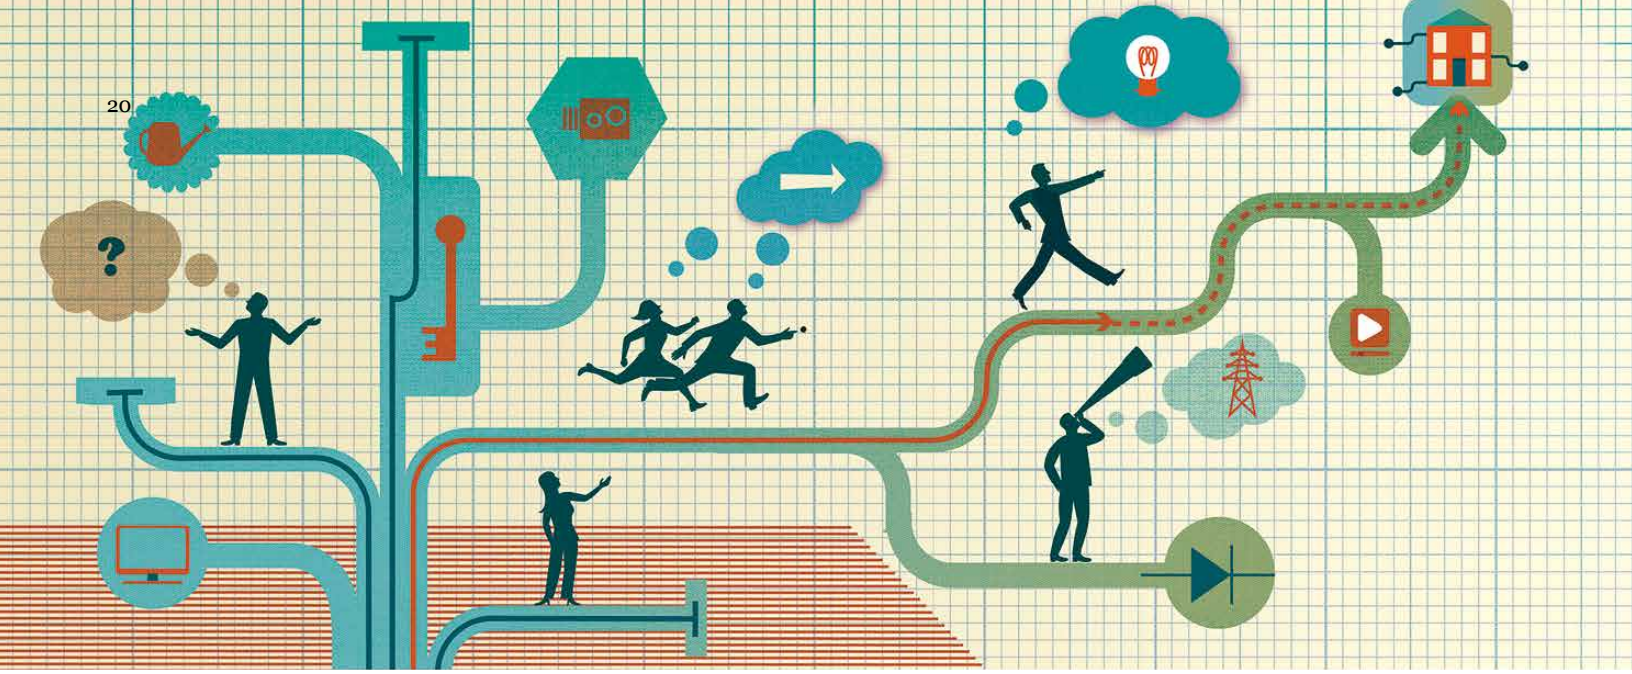

Andrew Baker

Trend spotting: Qualcomm executives consider the next wave of growth in semiconductors

Steven Mollenkopf and Murthy Renduchintala offer their take on the technology, talent, and business strategies required to keep pace in an industry that continues to evolve.

Abhijit Mahindroo, Nick Santhanam, and Bob Sternfels

Mobile technologies have been at the core of semiconductor growth over the past few years, but as growth rates overall have begun to taper off, the industry is asking itself, "What's next? Which technologies and strategic approaches will drive *new* growth?" In this interview, Steven M. Mollenkopf, CEO of Qualcomm, and Murthy Renduchintala, executive vice president of Qualcomm Technologies and copresident of Qualcomm CDMA Technologies, share their perspectives on what's ahead for the semiconductor industry and how Qualcomm and other technology companies can adapt to ever-evolving commercial and technological trends.

McKinsey on Semiconductors: *Growth in semiconductor revenues and profits has slowed as the industry has matured. Is there a growth challenge confronting the industry?*

Steven Mollenkopf: I don't see a fundamental growth challenge. The overall semiconductor industry may be roughly flat, but certain segments are thriving—the mobile industry, for instance. Although annual growth in mobile computing is declining from more than 30 percent to a somewhat lesser number, the mobile industry as a whole is still active and relevant and will remain so. Mobile phones are "talking" to other devices. They are aggregating data and enabling other ecosystems, such as automotive and home, to develop. Almost every piece of electronic equipment can use some technology inherited from the smartphone. Innovation in the area of

smartphone technology is not declining; if anything, we expect the production cadence to multiply, which will create growth.

Murthy Renduchintala: To a large degree, growth now is about achieving technological breadth. Qualcomm, for instance, started off as a modem player, but now we also focus on radio-frequency devices, application processors, power-management offerings, power amplifiers, and connectivity products. We cannot capitalize on new opportunities by looking inward, over-emphasizing profit-and-loss numbers, and being afraid to take on risk. So we approach new ventures with a long-term focus. Failure on the first, second, or even third generation of a product doesn't deter us; we have the scale to conduct new product research and design. Our goal is to lead in the areas of technology that will be sources of growth and new opportunities.

McKinsey on Semiconductors: *A number of semiconductor players have been engaged recently in high-profile mergers and acquisitions. What do you think is prompting this activity?*

Steven Mollenkopf: It's clear that a number of companies are looking at M&A as a means to grow quickly. The Avago Technologies–LSI deal seems to have opened up new possibilities in the industry; that deal is unusual in how the companies have built scale across different product segments. We pursue mergers and acquisitions only when it makes strategic sense for us to do so. In early 2011, for instance, we bought Atheros not just with short-term connectivity growth in mind but also to gain access to a different set of channels for introducing our smartphone technology in new ecosystems. We were able to realize this long-term strategic goal with the deal.

Murthy Renduchintala: I think there are some companies, such as Facebook, Google, and Twitter, from whom large acquisitions are expected. Investors expect these companies to have bold visions, and acquisitions often lend credence to those bold visions. At Qualcomm, however, we are not there yet. We need to be more discriminating in our approach. We have more than 50 R&D organizations across the world, and we are very thoughtful about how and where to acquire more. We cannot radically change our culture—it is the key to our success.

McKinsey on Semiconductors: *One of the important factors in Qualcomm's growth has been the fabless-foundry manufacturing model. What changes can we expect to see in this model going ahead?*

Steven Mollenkopf: One of the most interesting shifts over the past few years has been the importance of increased demand for mobile products, which has been a critical driver of the fabless-foundry model. There are few product categories in which companies can sell more than a billion units a year, and that will not change soon. Technology scale will have a big influence on industry dynamics, and we will see clear consolidation based on who can continue to invest in next-generation technologies. Some of this consolidation is already visible, and it will continue to accelerate.

Murthy Renduchintala: I think we will see greater fidelity between economic cost and benefit. We are already beginning to see "elongation" of industry demand over multiple nodes and technologies. For example, the scope of the mobile-technology portfolio also comprises power amplifiers and radio-frequency devices rather than just bulk CMOS.[1] Interesting collaborations such as the one between GLOBALFOUNDRIES and Samsung across leading-edge technology nodes

suggest an attempt to match supply diversity with demand. There are a number of unknowns, too—for instance, how soon will SMIC be able to reach scale here? What does the future hold for UMC?[2] The strategic landscape could take many different directions, all of which will have an impact on the industry.

McKinsey on Semiconductors: *Another critical factor in Qualcomm's growth has been its management of intellectual property. The broader high-tech industry also seems to be increasingly concerned about how to create and defend its intellectual property. What effect will this focus have on growth and productivity?*

Steven Mollenkopf: The turbulence that you see in intellectual property today comes from the collision of two different types of industries and models. On one hand, you have the cellular industry with its well-understood royalty structure. Standards bodies have been established that make it easier for industry players to collaborate and monetize their intellectual property. The rules have been set. On the other hand, you have companies from outside the industry trying to deliver mobile-computing and wireless products and realizing they also need intellectual property to do that. Eventually, I think players inside and outside the cellular industry will get access to the intellectual property they need, and that model will proliferate.

At Qualcomm, we have invested billions from our revenue base to fuel our research in wireless technologies. These investments have helped the broader industry grow, as well—the industry players using our technology can use our R&D investments, and they would not have been able to do this if we had been unable to monetize our intellectual-property portfolio. Certainly, some players are still using intellectual property as an offensive or defensive weapon, but I think we will eventually evolve to a more stable situation that encourages rather than hinders growth, productivity, and innovation.

McKinsey on Semiconductors: *The major buyers of information technologies used to be the so-called G7 industries—banking and securities, government, insurance, manufacturing, media and communications, retail, and utilities. Now the major technology buyers are consumer-facing Internet companies such as Alibaba, Amazon, Baidu, Facebook, Google, and Tencent, which are developing their own proprietary cloud servers and platforms, often bypassing original-equipment manufacturers. How could this shift affect semiconductor companies such as Qualcomm?*

Steven Mollenkopf: Our business has always been about enabling the success of our ultimate customers. So, for instance, we have spent a great deal of time understanding our telecommunications customers and creating products and services to suit their needs. We have extended that approach to critical cloud players as well.

Murthy Renduchintala: We must be careful not to overstate the effects of such a shift on original-equipment manufacturers. There is a lot of complex hardware and software integrated in a smartphone. It is true that the silicon and firmware must be aligned with the operating system, but a great deal of the fidelity of the smartphone is a function of the hardware and how it is integrated within the device. One of the virtuous benefits of scale is that imperfections in products inevitably get sorted out over multiple generations; over time, these improvements result in more differentiated and innovative products. Original-equipment manufacturers contribute more to this outcome than they usually get credit for.

McKinsey on Semiconductors: *How important is it for Qualcomm to develop a brand?*

Steven Mollenkopf: Marketing and branding certainly matter to us but not in the shape of a sundry flyer or something along the lines of the "Intel Inside" campaign. Our objectives are strengthening our channels and ensuring that our customers understand our business model. I see the need to be targeted and precise in our marketing efforts, but I do not foresee us spending as much on branding as we do on, say, product research.

Murthy Renduchintala: Our branding efforts have always been about maintaining our reputation and technical credibility. Operators check the quality of their networks based on how well a Qualcomm radio works on it. With our technological expertise across a number of cellular-transmission standards, including CDMA and LTE, we are, and should be, perceived as the technology bellwether.

McKinsey on Semiconductors: *Semiconductor start-ups have long been a source of innovation. But recently, there have been fewer new entrants and a decline in venture funding. How does Qualcomm manage its innovation pipeline?*

Steven Mollenkopf: A shrinking innovation pipeline is never good for the semiconductor industry. The hardware platform could benefit from the kind of open-source revolution that enabled the whole "three guys in a garage" era of software innovation. That said, we are interested in people using semiconductors and mobile technology in innovative ways, and the mobile sector is not as significantly starved of venture funding.

Steven M. Mollenkopf

Education
Holds a BS in electrical engineering from Virginia Tech and an MS in electrical engineering from the University of Michigan

Career highlights
Qualcomm (1994–present)
CEO (March 2014–present)

President and chief operating officer (2011–14)

Executive vice president, Qualcomm CDMA Technologies (2008–11)

Fast facts
Is a published IEEE author

Holds patents in areas such as power estimation and measurement, multistandard transmitters, and wireless-communication transceiver technology

Serves as chairman of the Global Semiconductor Alliance and as a member of the board of directors for the Semiconductor Industry Association

Murthy Renduchintala

Education
Holds a BE in electrical engineering and an MBA and PhD in digital communications from the University of Bradford

Career highlights
Qualcomm (2004–present)
Executive vice president, Qualcomm Technologies, and copresident, Qualcomm CDMA Technologies (QCT) (2012–present)

Senior vice president, QCT Engineering (2007–12)

Vice president, QCT Engineering (2004–07)

Skyworks Solutions/ Conexant Systems (2000–04)
Vice president and general manager of Cellular Systems Division

Fast facts
Member of Qualcomm's executive committee

Member of IEEE in both the United States and United Kingdom

McKinsey on Semiconductors: *Government is becoming an active player in a number of industries—finance, healthcare, natural resources. Why not in semiconductors?*

Steven Mollenkopf: I think the semiconductor industry is structurally different from the ones you mention. We do not affect citizens' daily lives as directly as the finance, healthcare, or natural-resources sectors do. Our products could be seen by some as expensive, but they are not prohibitively so. Most important, we are truly a global supply chain, and that makes it incredibly tough for any single government to regulate.

McKinsey on Semiconductors: *Many big companies become victims of their own success and tend to miss the "next big wave." How do you intend to keep challenging yourselves at Qualcomm?*

Steven Mollenkopf: We feel as though we are constantly reinventing ourselves, or at least every five years or so. In the 1990s, we were in the hardware and infrastructure business because that was the way to increase demand for our technology. Then we started making and supplying chips, and that business grew. Then WCDMA came along, and critics predicted that would be the end for us; eventually we took a leadership position in developing products that complied with that wireless standard. When smartphones appeared, critics said, "This is about computing, and Qualcomm knows nothing about it." We performed very well in that domain and exceeded expectations yet again when the LTE standard for wireless access emerged. Essentially, we are *always* challenging ourselves.

McKinsey on Semiconductors: *How has Qualcomm kept its culture intact despite rapid growth?*

Steven Mollenkopf: Culture is a part of every discussion at the company. We have developed a

strong engineering and innovation culture that we have sustained through different products and technologies. We are committed to a culture of continuous innovation and risk taking, and we treat our people exceptionally well—they are our most important assets, after all. There are two things that really motivate our engineers. The first is working on a product they can take pride in; they can point it out to a spouse or a family member and say, "I worked on XYZ, and it is a really cool product." The second is working on a project that is very interesting from a technological perspective. Lucky for us, these two factors are closely linked. As long as we continue to create innovative products and innovative technologies, we will remain in a virtuous cycle that will feed on itself.

Murthy Renduchintala: Our senior-management team really believes that "in your success may lie the seeds of your destruction" if you fail to pay attention to what's going on around you. We never want to get complacent. At the top levels of the company, we spend a lot of time questioning what we're doing, making sure we aren't getting lazy, and always remembering how we got to this point. Just a few bad quarters can be the difference between success and failure, and there are a lot of lessons for us to draw from. That's one of the reasons we invest so much in R&D and believe that you have to think about the business five years out. Also, we demand that our different product and technology groups become stand-alone centers of excellence.

McKinsey on Semiconductors: *Steve, you assumed the CEO role only recently; how has the journey been so far, and what is your vision for Qualcomm?*

Steven Mollenkopf: I have had a great time. I like technology, and I like the team we have. For me personally, it's very rewarding to know that almost everybody has the opportunity to use the tools and technologies we work on. Many technical discussions in the industry today are about topics such as cell phones and high-level operating-system software. In many cases, when I read about new technologies in the newspaper or on blogs, I am personally familiar with the products being discussed and the people who are playing an important role in their creation. I'm at the center of something very exciting, and I'm glad to be a part of it.

Looking five years down the road, I would like us to be the number-one mobile-computing company. To do that, we need strong collaborations. We have always been a relationship-oriented organization, and we have had strong, successful interactions with standards bodies and wireless carriers. We'll continue to do that, focusing on strengthening our ties with operating-system manufacturers, fabs, and other semiconductor players. Delivering a great product today is a much bigger undertaking than it used to be—more than any one company can completely manage on its own. Collaborative activities will be central to our next wave of growth.

[1] Complementary metal-oxide semiconductors.
[2] United Microelectronics Corporation.

Abhijit Mahindroo (Abhijit_Mahindroo@McKinsey.com) is an associate principal in McKinsey's Southern California office, **Nick Santhanam** (Nick_Santhanam@McKinsey.com) is a director in the Silicon Valley office, and **Bob Sternfels** (Robert_Sternfels@McKinsey.com) is a director in the San Francisco office.

Andrew Baker

Executive perspective: Vincent Roche, CEO of Analog Devices, on the next wave in semiconductors

The CEO of a multinational technology firm discusses the state of innovation in the semiconductor industry.

Aaron Aboagye, Abhijit Mahindroo, and Nick Santhanam

The semiconductor industry is at another crossroads: growth is slowing, the cost of innovation is rising, and several disruptive technologies and business models are poised to affect the industry. But the industry has already transformed itself many times over the past 30 years—embracing global models and markets, producing faster (and smaller) connectivity components, and developing new kinds of engineering, marketing, and sales talent. "The companies that will thrive in the future are those that can become bilingual—understanding not just technology but also business," says Vincent Roche, president and CEO of Analog Devices (ADI), a multinational technology firm that produces analog, mixed-signal, and digital-signal processing devices. "The companies that don't sense change, or that sense change but don't respond, or that learn from change but don't adapt quickly or effectively enough, will lose out." Mr. Roche recently sat down with Aaron Aboagye, Abhijit Mahindroo, and Nick Santhanam from McKinsey's global semiconductor practice to discuss where the industry has been, where it is going, and how companies can continue to adapt.

McKinsey on Semiconductors: *You've witnessed significant changes during your tenure at ADI. How do you see the industry and ADI evolving over the next five or ten years?*

Vincent Roche: To my mind, all semiconductor companies face two perennial questions: how to manage increased complexity in products, processes, and business relationships, and how

to react to the pace of innovation. These will continue to be the main themes for many years to come—not just for device companies but also for our customers.

ADI mirrors the broader semiconductor industry in many ways. We are an almost 50-year-old company, and our first 25 years were all about "big iron"—the IBM mainframe era and industrial measurement and control technologies. The second 25 years or so has ushered in the digital communications and consumer eras, and we have capitalized on this to grow from $300 million to $3 billion in annual revenues.

We are in an auspicious period now—a third wave of evolution—where we are combining many products and technologies to do bigger things for our customers while also managing the resulting complexity. It's analogous to the post-Cambrian explosion. The industrial, healthcare, automotive, and energy sectors are realizing big gains due to pervasive sensing, processing, and communication technologies. Meanwhile, we are seeing a tremendous reduction in the number of hardware engineers, especially analog engineers, at our customers' sites, and we are increasingly expected to fill that need by delivering more complete solutions. As a player at one node of the emerging ecosystem, we need a deeper understanding of how to successfully interact and cooperate with all the other nodes to make a difference.

McKinsey on Semiconductors: *Can you comment about the Internet of Things and the swirling attention around big data? What do these trends mean for customers, and by extension, semiconductor players?*

Vincent Roche: Personally, I think the term Internet of Things is overused. That said, many of our customers are realizing that, because of the ubiquity and power of communications technology and the extreme affordability of computing, there is a lot more they can do beyond building basic industrial machines, cars, network gear, or other hardware. They can connect their products to the cloud, capture vast amounts of useful data, and potentially redesign their business models to create new sources of revenue around analytics. These trends will be real game changers, and ADI is embracing the opportunity to highlight our expertise in connecting the physical and digital domains. Our company and other semiconductor players have a real growth opportunity here, because the devices we produce and sell can make it easier for our customers to collect information, perform sophisticated analysis, and do things differently as a result.

McKinsey on Semiconductors: *Many semiconductor players have become successful by leading device-level design in the industry. Now we hear companies talking more and more about software and system-level offerings. What are the implications of this trend?*

Vincent Roche: Compared with other industries, semiconductor players, in aggregate, are investing more in R&D and getting less in return. We are adding more complexity and sophistication to our offerings—for instance, embedding software and algorithms and increasing the functionality of integrated circuits and systems-level offerings. And as I mentioned, there has also been some expectation on the customers' part that we will provide certain hardware-engineering tasks and support capabilities. However, we are still figuring out how to get paid for these basic innovations and the other extras that customers expect for free.

Semiconductors are still the foundation of innovation in the market for information and communications technologies. At the end of

the day, no matter what our customers want to do in the cloud or with sophisticated data analytics, they still need the silicon as a foundation—and not just for incremental innovation but for real breakthroughs. The differentiating features in most automobiles today, for instance, are a result of the innovations that semiconductors enable. For a long time, a chip company was a chip company, but what does it look like now? We have to accept software development and systems engineering as critical domains in our work flow, and we have to organize ourselves around these needs. We have a road map for systems and software development at ADI, and it has become an important part of the innovation conversation happening in the company.

McKinsey on Semiconductors: *Semiconductor companies are facing ever-rising R&D costs and high-risk returns. How can executives successfully manage this challenge and continue to innovate?*

Vincent Roche: At ADI, we believe that "superior innovation makes for superior results." We invest nearly 20 percent of our revenue in R&D, and as long as we're growing, we believe that figure is appropriate. We need to bring more thoughtful risk into the company, pick the places where we play carefully, work closely with customers, and get better at getting rewarded for managing the increasing design complexity. The best, most innovative products we've launched over the

Vincent Roche

Education
Holds a BS in electrical engineering from Limerick University

Career highlights
Analog Devices (1988–present)
CEO (May 2013–present)

Vice president of strategic market segments group (2009–13)

Fast facts
Began his career at ADI in 1988 as a senior marketing engineer in Limerick, Ireland, and has served in various leadership roles ever since

years have been the result of collaborative innovation with our customers—applying the best technologies imagined by our engineers toward solving customers' most critical challenges. This approach has kept us relevant in the marketplace. We are careful to strike a balance between being customer-centric and technology-forward.

As we do this, we do need to be thoughtful about the amount of R&D and ideation we do internally versus externally. I think it was IBM's John Kelly who once said, "The world is now our lab." That is a striking and important statement. In addition to their own efforts, companies need to develop external relationships—with academic institutions, industry bodies, and other companies—to create new technologies. This holds true for early-stage product development as well as latter-stage initiatives.

McKinsey on Semiconductors: *Over the past decade, venture-capital funding and start-ups focused on semiconductors seem to be declining. Is this an issue? What can the industry do to deal with this?*

Vincent Roche: The industry needs to go beyond an incremental approach. We need to tap into new sources for ideas and breakthrough research, and start-ups can help in that regard. Venture-capital firms have been smitten by faster returns on software or Internet ventures, but there are indications that venture-capital investments are beginning to swing back to semiconductors.

Hardware-development cycles are longer and costlier than ever. But if we are truly an industry focused on the long term, thinking about the next couple of decades, we have a responsibility to help manage this problem. By providing funding and advice, and by collaborating in university research, we can help to improve the odds of start-up success.

McKinsey on Semiconductors: *Semiconductor players are increasingly looking toward mergers and acquisitions as a source of growth and competitive advantage. But post-merger integrations are often troubled. In your experience, what is the secret of a successful deal?*

Vincent Roche: The industry is in an acquisition cycle, so as long as capital remains cheap, you'll continue to see companies in our industry pursuing new deals. There really is no magic bullet for succeeding in M&A, but you do need to be clear about why you're pursuing the deal. Are you simply responding to overactive investors? A smart merger or acquisition is one where, in five years' time, you are more relevant to your customers and bringing more capabilities and innovative products to bear. The deal has to make you more competitive for the long term. You should imagine a conversation in which customers and share-holders come back to you five years down the road saying, "The combination of those two capabilities was really beneficial, and we're glad you did it."

Executives also need to be discriminating about their choice of targets. For instance, if innovation is built into the DNA of your company, you will need to pay special attention to culture: Is the target company's culture compatible with your own? How quickly can new technologies and capabilities be integrated? That was an important consider-ation in our recent acquisition of Hittite Microwave. The executives there shared a very similar mind-set and culture with ADI—they were very focused on innovation and developing new technologies and products for customers in markets that were of

strong strategic interest to us. They were close geographically, also in Massachusetts, which helps a great deal. In an opportunity-rich and resource-constrained environment, such as we are in right now, it certainly helps to get more scale, but you need to make sure the conditions are right.

McKinsey on Semiconductors: *With the rise of online giants, semiconductor companies are no longer the natural top choice for electrical engineers and computer-science majors. What will it take to continue attracting top-tier talent to the semiconductor industry?*

Vincent Roche: Just this morning, I met a group of bright, young engineers from Asia, Europe, and the United States, who are just beginning their careers at ADI. We talked about their aspirations, what is happening in the world, in the industry, and with technology. If that group of people is any indication, I am very bullish on the future of our company and the industry. They are very passionate about the work, and they want to innovate and make an impact on the world.

We are attracting the brightest people from great colleges worldwide. Although the big Internet companies have been grabbing the headlines and a lot of the engineering talent, I believe the pendulum is swinging back.

Analog, in particular, is a specialized craft. We train our people in foundational, core skills by exposing them to the greatest minds in our field. They go through multiple cycles of learning and stay with us a long time. I really think we are in great shape on the people front.

McKinsey on Semiconductors: *What would you like your legacy to be at ADI?*

Vincent Roche: We are a 50-year-old enterprise, and it is my goal to develop and position ADI to thrive for the next 50 years. We will need to continually innovate, relative to our past and relative to the market. We will need to increase our fluency in both the technology and commercial domains.

I am a student of evolutionary theory, and I think great companies have the same attributes as great societies, cultures, and nations—they sense, they learn, and they adapt. They don't just focus on competitors; they find mutually beneficial opportunities to cooperate.

Additionally, I have been privileged to work with two industry legends—Ray Stata and Jerry Fishman—who built our business. I stand on the shoulders of giants, and I hope to leave a legacy that they would also be proud of.

Aaron Aboagye (Aaron_Aboagye@McKinsey.com) is a principal in McKinsey's New Jersey office, **Abhijit Mahindroo** (Abhijit_Mahindroo@McKinsey.com) is an associate principal in the Southern California office, and **Nick Santhanam** (Nick_Santhanam@McKinsey.com) is a director in the Silicon Valley office.

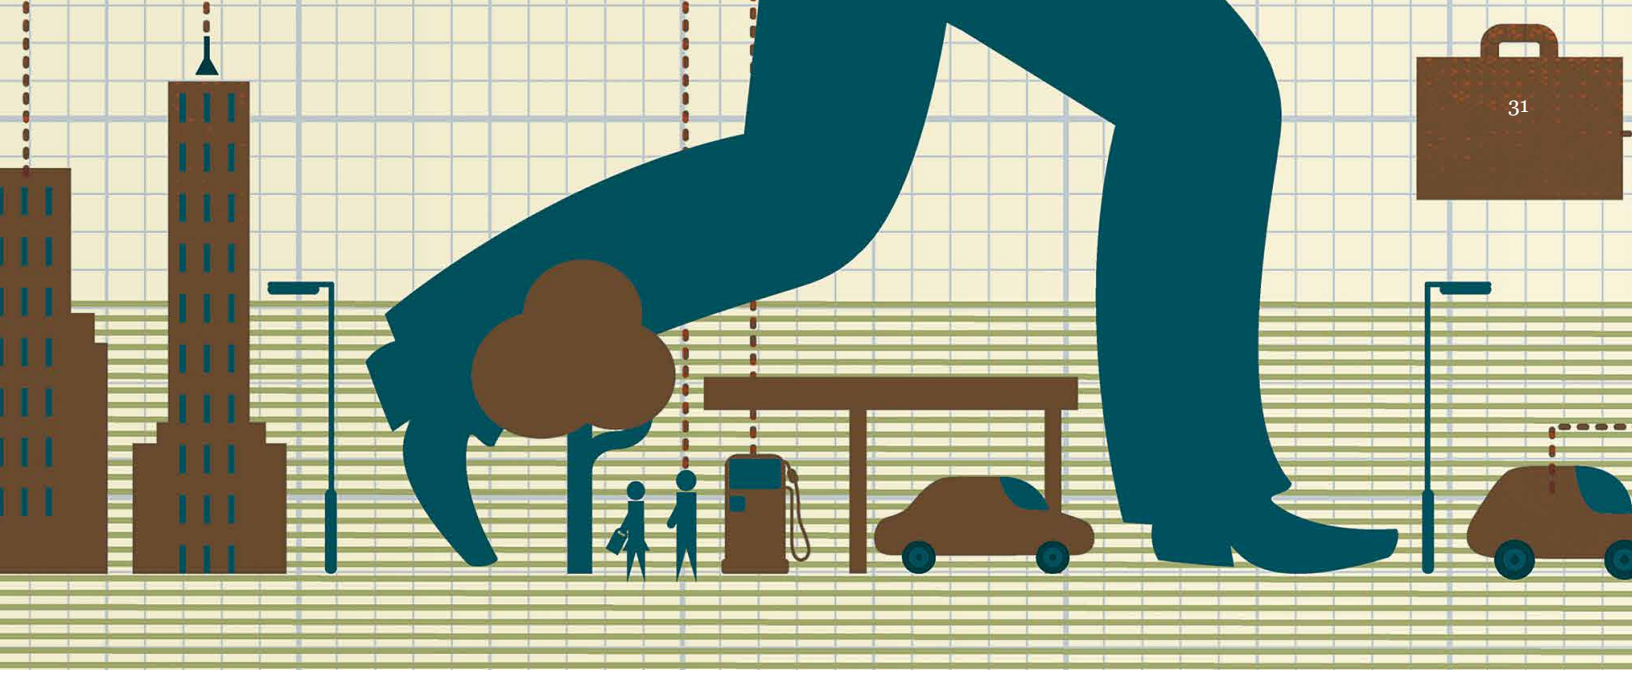

Andrew Baker

How big data and connected consumer products could boost the market for MEMS technology

MEMS technology continues to thrive in familiar markets, but with the advent of the Internet of Things, a significant new opportunity is emerging for industry players.

Harald Bauer, Sebastian Schink, and Florian Thalmayr

The overall market for microelectromechanical-systems (MEMS) technology—a category of devices that includes, for instance, inertial measurement units, gyroscopes, accelerometers, and pressure sensors—is projected to grow from about $11 billion in 2012 to approximately $23 billion by 2018 (exhibit).[1] While mobile phones, automobiles, and healthcare will continue to make up a large share of the MEMS applications market, there is another, potentially higher-margin use of the technology emerging within the next five years: as a critical enabler of the Internet of Things.

For more than a decade, as information has become increasingly digitized and computing power more robust, researchers, governments, and businesspeople have talked (in varying terms) about the emergence of smart, global, "object to object" communications. We define the Internet of Things as a universe of uniquely identifiable objects that are connected to a common network (public or proprietary) through which information about them can be exchanged (actively or passively) and analyzed.

In this universe, sensor-enabled equipment could facilitate the monitoring of production activities in a chemical plant. Sensor-enabled equipment in an automotive plant could allow managers to better predict which machines need maintenance, thereby decreasing production downtime. In the consumer world, connected objects might include

Exhibit

The MEMS market is projected to grow.

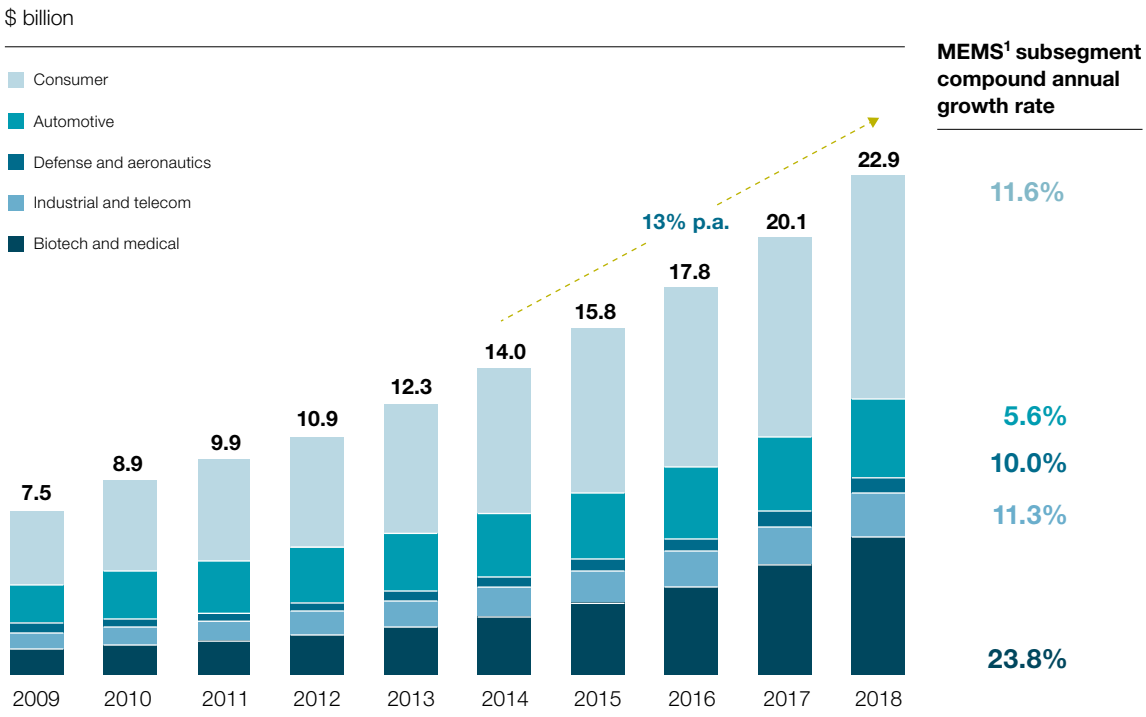

[1]Microelectromechanical systems.

Source: iSuppli; Yole Développement; McKinsey analysis

cars, fitness bands, washing machines, or home security systems.

MEMS sensors, already a proven technology in the mobile-telephony market, will provide the critical backbone for the Internet of Things simply because they enable the generation and collection of all the "small data" required to accumulate the big data that feeds the network and can then be analyzed to inform a range of business activities. MEMS-based infrared sensors and vibration and temperature sensors will help industrial-plant managers pinpoint root causes of production problems—for instance, detecting increased vibrations ahead of a machine's failure or diagnosing cracks in equipment. And building-management professionals are exploring the use of MEMS-based sensors to monitor and control ventilation systems and energy usage, room by room. Real-time climate and lighting conditions, as well as meteorological information, can be captured through these sensors and fed into a control system to allow for predictive heating and cooling—a smart building.

At the point where MEMS technology and the Internet of Things intersect, there is an opportunity for semiconductor companies and other industry

players to innovate, increase revenues, and reach new customers. In this article, we consider how MEMS technology is evolving and how fabs, integrated-device manufacturers (IDMs), and foundries can manage emerging trends—specifically, anticipating a shift toward technology acquisitions, player consolidation, and operational changes.

The technology advances

Over the next several years, the global MEMS market is expected to expand at a compound annual growth rate of about 12 percent, compared with only 3 to 4 percent growth in the semiconductor industry overall, according to the technology consultancy Yole Développement. Low-cost, low-power, small-footprint sensors built from MEMS technology can already be found in many mobile phones, cameras, and tablets on the market, as well as in every automobile and in some healthcare devices. These sensors enable the automatic rotation and adjustment of images on iPhone screens and support navigation functionality. Increasingly, MEMS-based microphones are replacing the condenser microphones embedded in cell phones, headsets, and laptops.

But the demand for MEMS-based sensors will increase exponentially not only because of the ubiquity of smartphones but also because of the rising popularity of connected consumer lifestyle products. Examples of the latter include glasses, watches, wristbands, and other wearables that allow individuals to monitor their heart rate, activity level, calories consumed, or sleep patterns, as well as smart appliances that allow consumers to optimize their home energy consumption or ensure home safety through the use of remote controls.

MEMS production is inherently scalable, enabling continuous improvements in chip performance, size, and cost (manufacturers' costs can drop 10 percent or more each year). A good example of such scalability is MEMS-based bulk-acoustic-

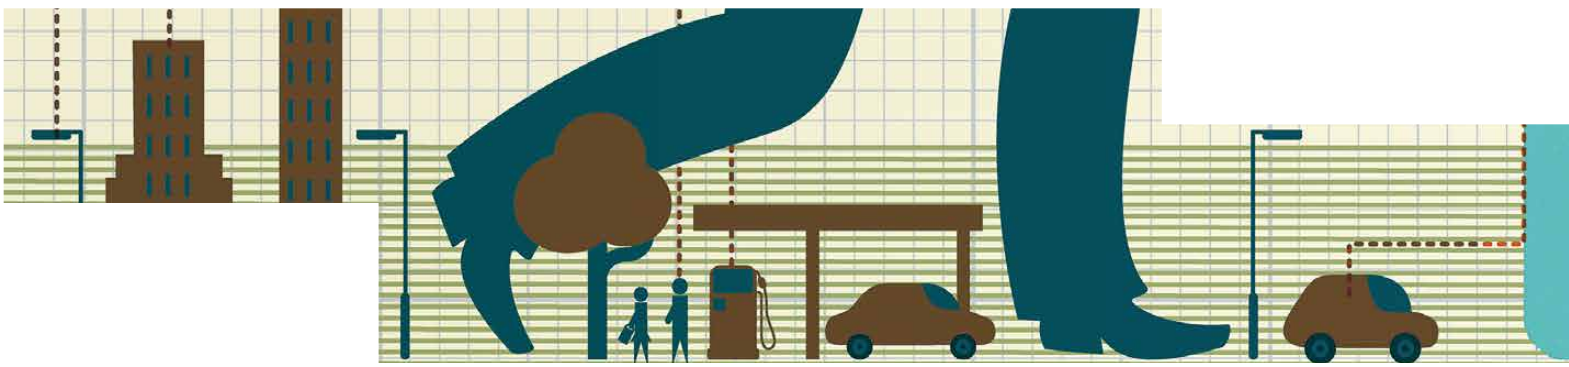

wave (BAW) filters and duplexers for radio-frequency front ends in mobile devices. If one looks at the bill of materials associated with today's high-end 4G smartphones, one would see that the cost contribution of BAW filters and duplexers is comparable to what it might have been when these devices were embedded in early-generation mobile products such as a Samsung CDMA watch phone in 2001. However, the number of duplexers in today's smartphones (for instance, the iPhone 6) has quadrupled while the size of a single duplexer has been reduced by a factor of 30. Performance increases have helped mobile users realize significantly improved cell-phone reception and battery life.

The innovation cycle for MEMS technologies is decreasing as developers improve on previous MEMS releases. As a result, many newer types of MEMS-based technologies are making their way into products. In the market for mobile handsets, for instance, MEMS-shutter-based display technologies could replace LCD screens. And MEMS-based micromirror technologies are gaining favor in the burgeoning market for projectors and head-mounted display products (think Google Glass and Oculus Rift).

Some chip makers, specifically in the mobile and sensor areas, are exploring the economic and operational benefits of developing integrated MEMS modules, which constitute a sensor or timing component, a logic component, and connectivity capabilities. The MEMS chip maker Sand 9 and Intel recently demonstrated an integrated transceiver for cellular phones, including frequency reference. Sand 9 is providing MEMS-based timing devices that are 50 percent smaller than conventional timing devices. Thus they can be copackaged with Intel's transceiver chip through overmolding.[2] Such MEMS-based stacks can provide enhanced functionality at lower cost, with a smaller footprint (which is critical for use in more compact, connected consumer electronics, such as smartphones and fitness bands) and with relatively low power consumption.

How will the market respond?

The trend toward MEMS integration suggests that a substantial part of the market will be served by a few big players that can offer "many in one" chips. Companies that produce single-device chips (with only an accelerometer or only a gyroscope, for instance) will still be able to thrive but mostly in niche areas outside of consumer electronics—for instance, providing chips for certain automotive and military applications, both of which represent the traditional customer base for components manufacturers. We expect the MEMS market to follow a typical "hogs' cycle" over the next few years, with pronounced periods of overcapacity followed by periods of price erosion, investment cutbacks, and shortages, followed by a wave of consolidation.

Non-MEMS players may attempt to acquire MEMS suppliers to integrate their devices into their own silicon and improve their competitive positions. That was the case when ROHM targeted Kionix for acquisition in 2009. The Japanese wireless-communications company wanted to add sensor technologies to its portfolio, and Kionix was one of the market leaders at the time for MEMS accelerometers. Rather than take the time to build MEMS capabilities in-house, ROHM acquired Kionix's proven platform, design expertise, and manufacturing capacity. In part because of this acquired technology, ROHM was later able to partner with the German

tech firm EnOcean to establish EnOcean's energy-harvesting wireless technology in the Japanese market.[3]

In-house development of integrated MEMS can be difficult, because it is time consuming and because many of the device and process technologies associated with MEMS are already patented. So even the leaders in the MEMS market will likely turn to acquisitions to build the combination devices that will further strengthen their portfolios, although incorporating rivals' technologies into their stacks may prove to be difficult given the specific production processes associated with various MEMS devices. One company that produces resonators, for instance, may rely on a process whose core element is deposition of a piezoelectric layer, while another company producing a device for the same application may rely on a process whose core elements are creating a tiny air gap and controlling oxidation levels.

Partnerships with or strategic investments in MEMS providers may prove to be the most effective way for semiconductor companies to quickly gain access to emerging MEMS devices. The components maker Alps Electric, for instance, was able to tap into Qualtré's MEMS expertise by making a $3 million strategic investment in the company in June 2013. The companies are aiming to codevelop and bring to market three-axis inertial sensors based on Qualtré's BAW technology and supported by Alps Electric's manufacturing capabilities and global sales resources.[4]

The stakes for MEMS players

These trends are likely to have direct implications for individual players within the MEMS market. Among them, fabless design companies will probably have an advantage because of their role as small but critical contributors to innovation in the MEMS market and their ability to react quickly to swings in demand (the latter due to their low overhead). The fabless chip designer InvenSense is now one of the fastest-growing MEMS companies, with a cost structure that is 15 to 20 percent lower than its competitors among IDMs, according to Yole Développement.

For their part, IDMs should not need to replace much of their existing equipment to capitalize on the growing demand for MEMS devices, because many single-function MEMS devices do not require bleeding-edge resolution or lithography. Older fabrication plants should work fine. However, the IDMs that want to explore MEMS integration (producing stacked devices with many functions) will need to take the time and find the resources required to harmonize their manufacturing processes, since many of them will have multiple process-technology platforms in place as a result of legacy acquisitions. These semiconductor players may also need to adapt their sales and marketing processes to reflect their new offerings—MEMS modules rather than components.

Meanwhile, the symbiotic relationship between foundries and fabless companies is expected to evolve to accommodate MEMS. For certain manufacturing technologies, such as complementary metal-oxide semiconductor, or CMOS (a technology for constructing integrated circuits), foundries have been able to offer standard design libraries to their fab clients, creating flexibility for customers that want to, say, design a single chip that supports two different operating voltages. One family of recipes can be used to make thousands of different products. Previously, most foundries supported proprietary MEMS development but did

not offer a common MEMS design library. However, more and more foundries, such as X-FAB Semiconductor Foundries and GLOBALFOUNDRIES, are starting to own and offer "open" processing modules for the production of MEMS devices—providing, for example, a module comprising silicon-on-insulator wafers with pre-etched cavities, poly through silicon via for interconnects, and hermetic sealing. In this way, foundries are able to help remove barriers to entry for small fabless players, which could fuel innovation within the start-up scene.

• • •

The mobile market has already lit a fire under MEMS chip producers; innovation cycles are getting shorter and costs are coming down. That flame will grow as consumers find value in interconnected objects, and as industrial clients see the advantages in adopting new Internet of Things applications to complement their existing capabilities. Producers and suppliers of MEMS-based technologies must recognize the effects that the move toward integrated MEMS devices will have on their operations and in the marketplace. Those that adapt can seize significant opportunities for growth; those that don't risk falling far behind. ○

[1] R. Colin Johnson, "MEMS market to top $22 billion by 2018," *EE Times*, November 8, 2013, eetimes.com.
[2] R. Colin Johnson, "Sand 9 MEMS cracks cellphone market," *EE Times*, September 3, 2013, eetimes.com, and "Integrated MEMS Oscillator for Cellular Transceivers," presentation at 2014 IEEE International Frequency Control Symposium, Taipei, May 19–22, 2014, ifcs2014.org.
[3] "EnOcean and ROHM announce strategic partnership at electronica 2012," November 12, 2012, enocean.com.
[4] "Alps Electric and Qualtré, Inc. announce partnership for next generation inertial sensors; strengthen relationship with strategic investment," June 7, 2013, qualtre.com.

Harald Bauer (Harald_H_Bauer@McKinsey.com) is a director in McKinsey's Frankfurt office, and **Sebastian Schink** (Sebastian_Schink@McKinsey.com) and **Florian Thalmayr** (Florian_Thalmayr@McKinsey.com) are consultants in the Munich office.

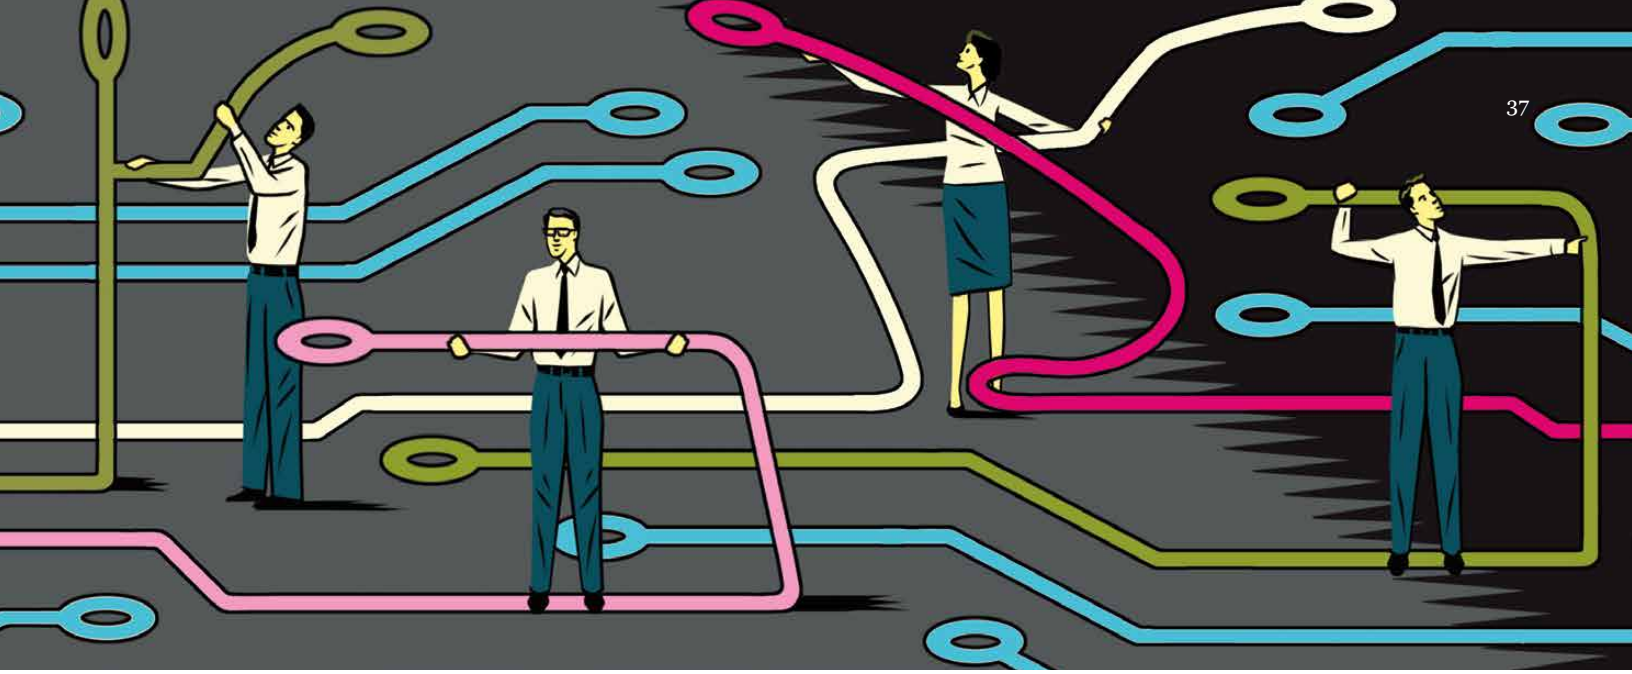

Bill Butcher

How semiconductor companies can get better at managing software development

They may want to consider adopting one of four basic organizational structures.

Gang Liang, Christopher Thomas, and Bill Wiseman

Software isn't just for purchasing or outsourcing anymore. Increasingly, semiconductor companies are exploring in-house software development as a way to reduce costs, improve time to market, and differentiate themselves from competitors. Organizations that have traditionally been focused only on hardware (silicon wafers and circuits) are hiring more software engineers to support the development of the integrated circuits that are at the heart of most consumer electronics, medical devices, automobiles, appliances, and pieces of heavy equipment in use today. In the late 1990s, it was common for chip makers to invest in one software engineer for every ten hardware engineers; today the ratio is closer to 1:1 or, in some cases, 1:1.5.

Chip makers are pursuing software development primarily to meet customers' growing demands for more sophistication in and more support for the components they buy. Indeed, the leaders in various sectors of the semiconductor market—Intel, MediaTek, and Qualcomm, among others—now routinely earmark significant portions of their R&D budgets for software development, and some are already providing end-to-end software-based products for customers. MediaTek, for instance, provides both software and reference designs with its chip sets so that customers can create their own branded mobile phones with basic features. Its end-to-end offering has helped customers cut their own time to market and R&D costs and has allowed them to focus more

on strategies for branding and differentiating their products.[1]

But while some market leaders have been at this for a while, other players are only just starting to take a closer look at how they build, use, and manage software (see sidebar, "What are they building?"). They face a number of challenges: software resources at hardware-oriented companies tend to be limited, and engineering talent can be scarce and hard to acquire and retain. Additionally, selling silicon is the main business, so efforts to divert scarce resources toward software development may meet internal resistance.

Semiconductor companies that choose to pursue a rigorous software-development program will need to have the right organizational structure in place—one that enables companies to motivate talent, control the R&D budget, launch products more quickly, and meet customers' expectations. Some semiconductor companies have established a central software organization to support all of their business units, while others are struggling to keep up with "rogue" development efforts happening within various business units, each of which has its own software team. Our work points to four potential organizational structures that software-minded semiconductor companies may want to consider adopting to get the most from their existing development efforts and to make it easier to pursue new software R&D initiatives: completely decentralized, completely centralized, hybrid, and leveraged. There are advantages and potential pitfalls associated with each, and the appropriate structure will differ for every company based on its existing talent, resources, and overall business objectives.

Four ways to organize

We have seen 1,000-person companies make the transformation from one organizational structure to another within 12 months, but a years-long effort is much more typical, particularly if the company is starting from scratch or having to make hard decisions about which groups to merge. In either scenario, a change in metrics and mind-set will be required. Executives will need to develop mechanisms for tracking the productivity gains from their software R&D, and they will need to foster engagement and commitment to software development across the company.

Completely decentralized. Semiconductor companies with a number of different business units that have little or no business or technical crossover likely would find it easiest to pursue a completely decentralized software organization: each of the business units funds and manages its own software group, and the unit's general manager retains the autonomy to deploy software resources where needed. In the 1990s, Intel's architecture business unit boasted a dedicated software organization that created homegrown development tools to support its x86 systems. Even today, the software group works closely with a number of third-party software vendors—Oracle and SAP among them—to optimize those applications for every generation of its central processing units. Intel also had separate software groups dedicated to its NOR Flash and i960 businesses. The NOR Flash software team built up a strong capability in device drivers, and the i960 software team focused on enabling Intel silicon to work well with third-party software and applications. There was almost no overlap in customer bases or operations among those business units.

The completely decentralized model works well so long as the units' businesses and technologies remain independent—which, in today's semiconductor environment, is quite rare. If, for instance, units are combined or new businesses emerge and need the same fundamental software and technologies being developed and managed in other groups, it makes less sense (operationally and financially) to duplicate efforts. One large integrated-chip manufacturer, for instance, had created separate handset and tablet business units as each of those technologies emerged in the marketplace. There were separate software-development groups for each unit, but the company eventually realized that development teams in both used the same system on a chip, which meant the company was wasting its resources and needlessly creating conflict and competition between two groups of engineers.

Completely centralized. For semiconductor companies whose business units rely on all the same technologies, a completely centralized software organization will be most efficient and effective. Under this organizational structure, software-development and technological expertise radiates from a central group—one that reports to the C-suite—removing potential redundancies and significantly reducing resource and development costs. That is the case for one large US components vendor. It provides integrated circuits to manufacturers of a range of consumer electronics, including laptop computers, mobile phones, tablets, and other devices. The underlying graphical processing unit and other technologies embedded in its chips are standard, so one version optimized for Android is suitable for all its customers. Having one centralized software group allows the company to better manage all its licensees and reduce development costs.

What are they building?

At the start of the shift toward in-house software development, many semiconductor companies were focused primarily on developing their own firmware—software embedded in their integrated circuits that would dictate how the chips would function. Over the past few years, some have started working directly with operating-system vendors to make sure their device drivers will work seamlessly and their processors will perform optimally in those environments; still others began to release software tools (compilers, debuggers, tuners, and the like), plus common libraries and middleware, so third parties could create optimized applications for their company's chips. Most recently, semiconductor companies have started to create end-to-end, embedded software products for original-equipment and original-device manufacturers.

A completely centralized model also confers upon semiconductor companies other benefits, including a consistent approach to R&D planning, a standard set of software-development and management tools, a common software-development process and methodology, and comprehensive rules and standards for assuring quality and appropriately managing source code. By establishing this level of consistency across all business units, chip makers can reduce their R&D costs and accelerate growth in new and critical businesses that may not otherwise have the funding or technical capabilities to pursue software development as a complement to their existing work. This centralized structure also may facilitate offshore expansion or development outsourcing—the strategic moves favored by many semiconductor companies these days—by making it easier for them to manage global engineering resources or maintain relationships with vendors.

There are a few drawbacks, however. For instance, the funding model for this approach can be complicated. In many companies that use a completely centralized model, the business units pay a "tax" based on their needs, financial strength, and other criteria. This can be a headache for the finance team, which has to calculate the difference between projected needs and actual demand for each business unit—each of which would obviously want to pay as little of this tax as possible. Additionally, under the centralized model, the business units would have less control over software development as a resource. Often, the objectives of the centralized software group and the business unit will not be completely aligned; the units may have unique requirements that a centralized organization simply may not be aware of. For example, video playback on a mobile device and videoconferences over the Internet both use the same H.264 video codec (or compression software), but the implementation of H.264 in each case is quite different. A centralized software group could easily deliver the common video codec but likely would not have the technological expertise to support its implementation on both mobile and Internet platforms.

It is critical for companies that adopt a centralized model to pay attention to process, metrics, and collaboration—for instance, convening a small team that represents the interests of each of the business units and the centralized software group. The team would meet regularly to analyze software priorities and rank them according to business units' needs and the impact of certain projects on the company overall. It is also good practice to establish service-level agreements between the centralized software group and business units to help clarify roles and responsibilities—and to preserve some level of control for the general managers involved.

Hybrid structure. This organizational approach combines the financial and technological efficiencies provided by a centralized software group with the greater flexibility and controls that a decentralized structure may offer the business units. At first glance, it seems to present the best of both worlds. In reality, there are significant funding and operational challenges to address. Under the hybrid model, the technologies and software capabilities that are common to all business units become the property of a centralized group, while the technologies and software that are unique to a particular business unit are maintained and developed separately. One leading

maker of mobile chip sets, for instance, has a centralized software group that is charged with enabling and optimizing the Android operating system for use with the company's generic system-on-a-chip architecture. But each of the company's business units "owns" a version of this technology that it uses in ways that are specific to its group.

Besides just holding on to the common software, the centralized group should also establish best practices for its use and encourage sharing among all the other software teams within the company. To that end, a joint committee should be convened to manage common software-development priorities, and service-level agreements should be drawn up. But as with the completely centralized model, a charge-back process must be established; the use of common technologies would be subject to a tax based on revenues, profits, or other criteria, and each business unit's software organization would be required to fund its unique development initiatives separately.

Leveraged structure. Many semiconductor companies have a core business and a number of units that are derivative of the core. For instance, one chip manufacturer's core business is in microcontrollers and microprocessors, primarily for the automotive market, but increasingly its technologies are also being used in medical and consumer applications. For it and companies like it that are exploring market expansion, a leveraged software organization may make the most sense. Under this structure, the software group would report to the core business unit rather than to a centralized corporate team. As with the hybrid model, the software organization would own the completed software components and resources but would deliver them to the rest of the company. For instance, the software team in the chip manufacturer's consumer-products business unit could take technologies developed by the software team in the company's automotive unit and modify them to suit its and the market's needs. As with the centralized model, the core business's software group would need to establish best practices in software development and encourage sharing across the organization, but the other business units would have to fund their own unique development initiatives.

Which model?

To determine which of these structures is best, companies need to consider their existing software capabilities—that is, the type of software R&D they are currently undertaking (if any), their overarching objectives relating to software, and the funding and other resources at their disposal. They should also consider their competitors' software capabilities.

It is unlikely that many companies would pursue a completely decentralized model; this type of structure just isn't the norm in today's semiconductor environment. But the companies that already have lots of software R&D experience, or that have a core business unit with several businesses feeding off of it, will want to explore hybrid or leveraged models. The individual business units would immediately benefit from software technologies that are already in hand (managed by the centralized software organization), but they would retain the flexibility to create unit-specific products based on their unique technical and business needs. Such companies could realize less duplication of effort and waste. By contrast, the companies that have limited software R&D experience may want to set up a centralized software organization focused on just one business or a few business units at first—starting narrow

to ensure that success is within reach but establishing best practices that can be rolled out more broadly as software-development initiatives gain momentum.

These decisions won't necessarily be permanent; as semiconductor companies move from a single-minded focus on developing silicon components to a broader focus on delivering end-to-end offerings built around their integrated circuits, their software organizations will need to change as well. In the transition from one model to another, executives may need to introduce key performance indicators and other metrics to help the software organization (however it is structured) quantify the impact of its development efforts and to help project leaders set and meet personal targets. Because of the global scarcity of technical talent, semiconductor executives also may need to adjust some of their human-resources practices—for instance, providing attractive, high-profile assignments in which software experts actively participate in product design and planning, or letting software engineers lead higher-level strategy discussions. Most important, executives who are bringing software R&D in house will need to become steeped in basic software terminology and concepts. They don't have to be experts, but gaining at least a rudimentary understanding of what the software can and can't do may help them achieve their business objectives in the long run.

• • •

The software-development function in most semiconductor companies typically flies under the radar—until growth slows and executives with cost cutting in mind notice the large cadres of engineers they've acquired over the years or until a new business opportunity emerges and executives notice how few engineers they have on staff. We believe executives need to be more proactive; they need to recognize the complexity and collaboration associated with software development and react accordingly.

[1] *Android Authority*, "MediaTek is riding high, how far can it go?," blog entry by Simon Hill, April 25, 2014, androidauthority.com.

The authors would like to thank Harald Bauer and Ondrej Burkacky for their contributions to this article.

Gang Liang (Gang_Liang@McKinsey.com) is a senior expert in McKinsey's Boston office, **Christopher Thomas** (Christopher_Thomas@McKinsey.com) is an associate principal in the Beijing office, and **Bill Wiseman** (Bill_Wiseman@McKinsey.com) is a director in the Taipei office.

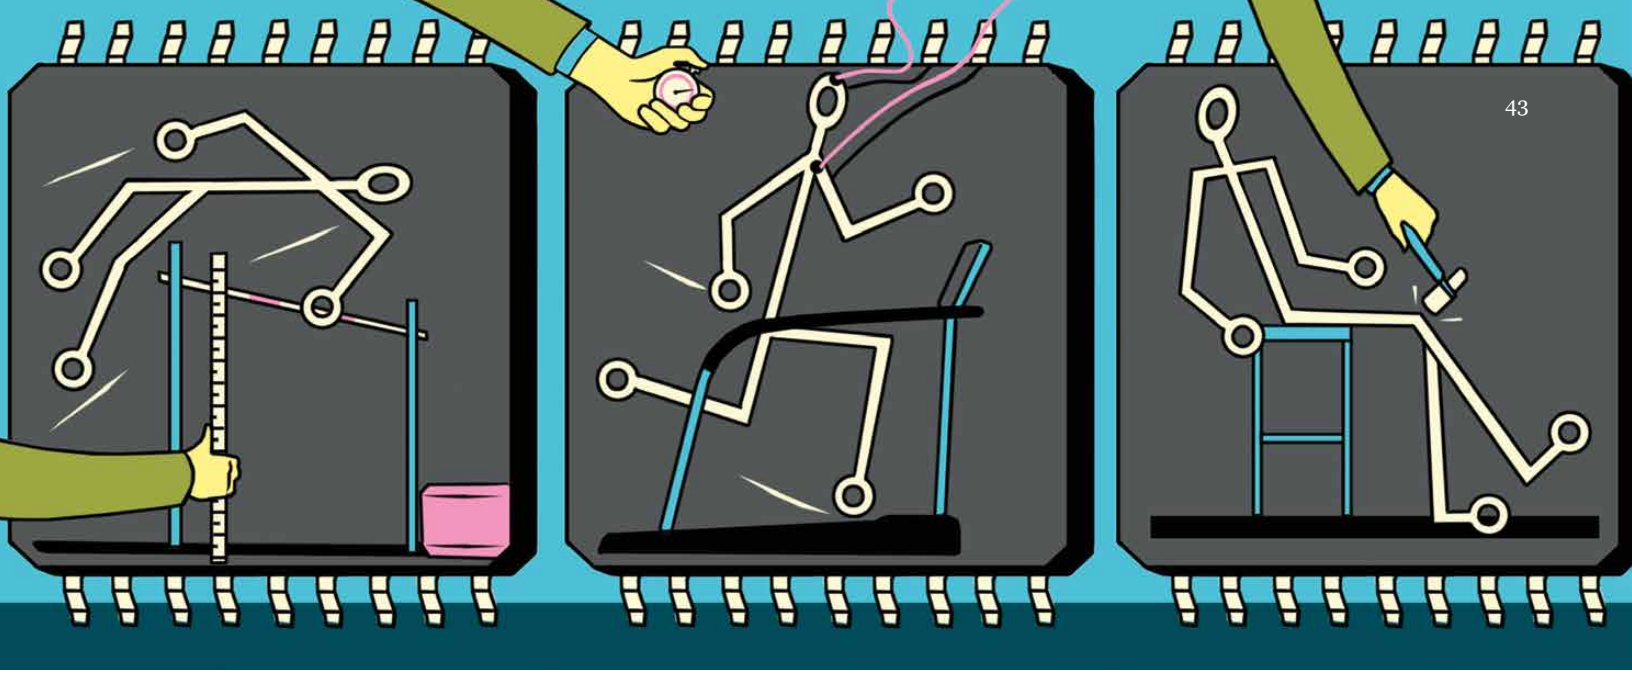

Bill Butcher

Standing up to the semiconductor verification challenge

Companies should seek faster, more cost-effective ways to test the quality of complex system-on-a-chip devices.

Aaron Aboagye, Mark Patel, and Nitin Vig

The tail is wagging the dog in most system-on-a-chip (SOC) development efforts.

Design verification, the end-stage process of ensuring that everything on an integrated circuit works as planned, consumed more than 55 percent of the total time spent on a typical SOC design project in 2012, up from 49 percent in 2007, according to the Wilson Research Group's 2012 Functional Verification Study. This increase in time spent is a direct reflection of newer, more complex generations of semiconductors that have many more transistors and many more functions, all of which must be carefully vetted. Flaws that are found late in the production process, or not at all, can create poor customer experiences that can damage chip designers' reputations. So most companies have accepted the risk-versus-efficiency trade-off and are relying on conservative, resource-intensive approaches to design validation and testing. As a result, however, they are often forgoing potential profits from the timely release of their SOC devices. What's more, the demand for ever-increasing complexity in today's circuitry is not likely to slow down, so the percentage of total project time that must be spent on SOC verification will likely continue to increase—unless semiconductor companies rethink their approach.

An obvious first step is to assess and adopt testing technologies that can help streamline the

verification process. But some of these tools and techniques can be expensive to implement, particularly for smaller semiconductor players. So companies may also want to consider ways to simplify the verification tasks associated with a particular chip or family of chips and examine the steps they can take to improve the infrastructure they have set up to support verification efforts—for example, creating a centralized verification organization, with a dedicated senior leader, to oversee the testing process. Indeed, by enhancing their capabilities in three areas—technology, process, and organization—and by taking time to develop an overarching verification plan rather than tackling the verification process chip by chip, companies may significantly reduce both their time to market and the cost of development associated with their SOCs.

In this article, we consider the use of data analytics in verification-project planning and discuss ways to simplify integrated-circuit testing. We also look at the advantages of different verification technologies for different players, as well as ways to establish an organizational infrastructure that facilitates efficient SOC testing. These technology, process, and organizational levers can be used to complement—and jump-start—companies' traditional approaches to verification. Our study of productivity measures associated with more than 1,400 integrated-circuit projects suggests that, by streamlining the verification process, chip companies may be able to increase their productivity by at least 10 percent—for instance, by closing their design-specification stages faster, producing more SOC devices, and moving them to market more quickly. Such an approach could provide even small semiconductor players with a means to differentiate themselves from larger rivals.

The verification process

The traditional approach to verification begins after chip design is complete. The chip design is simulated via a "test bench," which consists of software code written in the hardware-description languages that a designer uses to test each functional block of an integrated circuit. The test bench instantiates the chip, supplies stimulus signals, and measures and evaluates the resulting responses from the chip. This process enables the test bench to determine whether the block meets predetermined specifications.

When dealing with a complex system on a chip, different designers typically will work on individual functional blocks within the chip, often following different schedules. Complications can arise because the designer working on functional block A often requires input from blocks B and C to verify his or her design. As SOCs have become more complex and the number of transistors on them continues to climb, managing these interrelationships has become increasingly unwieldy. And the task will not get any easier with the ongoing trend toward miniaturization.

Resolving the verification challenge

Our experience, industry research, and expert interviews suggest that semiconductor players often skimp on the time spent in verification planning. Put simply, design teams don't know what they don't know about their approach to verification and are therefore missing significant opportunities to improve aspects of this critical quality-control process. Teams will often focus most of their verification resources on executing the project and staying on schedule, leaving less time for up-front critical thinking (exhibit). As a result, they may not recognize the chances they have to save costs and create production

efficiencies by, for instance, reusing certain pieces of intellectual property or simplifying chip architectures. In the face of rising production costs and complexity, semiconductor players should take more, not less, time for planning, using readily available project and process data to set their verification objectives. Armed with this information, design teams can find ways to introduce simplicity into their verification processes and tailor their use of new technologies to specific verification situations rather than assuming one size fits all. The data can also inform companies' attempts to build a robust verification organization.

Use data analytics in planning discussions

Verification may happen late in chip development, but conversations about quality control should occur quite early and often. Companies should convene project-launch planning discussions, bringing together members of the verification team, leaders on individual projects, design engineers, and senior managers. The goal is to get an accurate read on the resources required to carry out verification tasks for every SOC in development, what the testing schedule should look like, and the potential risks associated with certain SOCs or families of chips. Input from members of the verification team in these conversations will be crucial; they will have the institutional knowledge required to make qualified estimates. Advanced analytics can play an important role here, as it now does in decision-making processes across most industries. Using historical project and process information, semiconductor players can develop a comprehensive verification-process database that, over time, will allow senior managers to see where and when critical pain points are likely to emerge in the typical verification process and react accordingly. The data and planning discussions can also help companies determine how much verification is enough—often, the decision about when a device is "finished" is partially based on

Exhibit

Verification efforts can significantly outpace projections without solid up-front planning.

System-verification tasks, % of effort spent

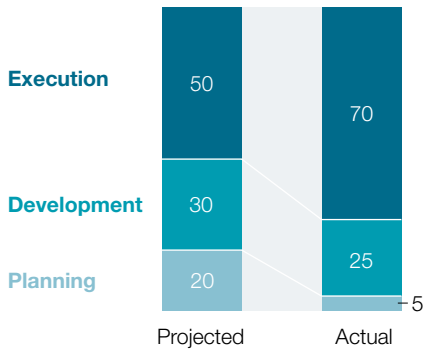

how much time is left before it is slated to launch rather than how much time is actually required to ensure reliable functioning of the chip. As a result, bugs are found late, and mask layers need to be regenerated.

By contrast, we have seen verification teams use the data at hand to prioritize various test scenarios associated with particular SOCs. In this way, they can quantify not only the number of tests required to ensure the operability of their intellectual property but also the number required to prove that a component or module does *not* work. One semiconductor player, for instance, used analytics in feasibility discussions about a new integrated-circuit concept. Members of the verification team met with engineers and senior managers to outline their projections of the validation effort and the resources that would be required to bring the device to market. They prioritized the testing scenarios that would need to take place before the device could be deemed done. As issues emerged, the group was able to go back to its plan and, based on the data, recalibrate activities and objectives associated with the development of that integrated circuit. Over time, the team built up a rigorous database of project-verification information; the accuracy of its work-plan projections improved significantly for each successive project.

Simplify the elements to verify

Based on existing project data or user feedback, there may be ways to streamline chip designs or head off performance issues long before verification tasks come into play. In the design-exploration phase, for instance, engineers can consider ways to reuse current intellectual property rather than introduce new intellectual property that might complicate eventual verification efforts. To build the business case for minimizing changes, engineers could review the verification efforts associated with earlier system-on-a-chip tape-outs—the phase in which designers share the photo mask of a circuit for fabrication. Design teams could then categorize SOC projects according to how much (if any) intellectual property was reused and, for each chip or family of chips, compare the verification efforts that were required at the end of development.

When a chip design *does* require new intellectual property, teams can identify and develop the required electronic system–level or C++ verification models early on to ensure that downstream verification is feasible and would not introduce unexpected issues. Before chip design even begins, verification and engineering experts should test the logic embedded in structures that are complex (such as first-in/first-out data structures) or time sensitive (such as arbitration controllers). If testing challenges emerge, the engineers can simplify the designs at the outset, before teams have sunk significant time and resources into the development process. Of course, teams will not have unlimited time and manpower to perform this kind of up-front testing. They may decide to use this approach only when new and critical features are being implemented or only in the development of the most complicated SOCs—such priorities may be determined during early planning discussions, using the data at hand.

Assess the latest verification technology

Verification teams have always had access to a wide range of technologies for creating high-level simulations and prototypes of circuits. There are simulators for testing register transfer–level designs and logic gates. There are hardware-acceleration techniques (also known as emulation techniques) to speed up the verification of large

designs. But newer tools have emerged that allow for robust, mixed-signal simulation so that digital and analog design components can be verified together. And next-generation emulators can operate at higher speeds and handle even larger designs.

The technology has improved; still, none of these approaches, on its own, is a panacea for companies' inability to find design flaws early and release bug-free products. Simulation, the least expensive approach, can be too slow for large designs. Emulation is faster but more costly. Prototyping can provide immediate test results but may be prohibitively expensive for some companies—particularly smaller semiconductor players.

To take advantage of new technologies but keep costs in check, companies can use basic emulation techniques instead of prototyping, and they can exploit cosimulation tools that simultaneously model hardware and software functions to verify hardware and relevant portions of software code. Small semiconductor players may also want to explore the use of cloud-based servers and computing infrastructures, which are provided these days by a number of electronic-design-automation vendors. Third-party IT resources may be particularly useful during tape-out periods, which typically constitute crunch time for smaller project teams: they need the extra computing power for managing tape-out tasks but not during normal work periods, so many consider it a waste to build out large computing infrastructures that would be underused much of the time. To ensure that design data would not be compromised, semiconductor players would need to work closely with third-party platform and service providers to establish rules and protocols for creating secure cloud-based environments.

Larger companies with deeper pockets and pools of talent should attempt to push the technology envelope further—for instance, using mixed-signal simulation as well as virtual-testing platforms in their verification processes. Mixed-signal simulations can generate relatively accurate, cost-effective results, given the faster simulation speeds now possible. Engineering teams may still need to prototype new devices or portions of a system on a chip, but even in those instances, they can use mixed-signal simulation to improve the accuracy of their findings.

In fact, semiconductor companies of any size could realize great cost savings and productivity benefits by making virtual platforms an integral part of their SOC planning and design cycles. The common platform, which would be used for setting goals and for overseeing progress toward those objectives, could help mitigate the need for rework as a chip moves along the production track. Software and hardware engineers could collaborate from the outset on SOC design stages, which would have a favorable impact on verification stages downstream and could allow semiconductor players to close the specification phase of SOC development much faster. Having a virtual platform allowed the engineering team at one semiconductor company to accelerate its software development and have it ready for ramping up and debugging the underlying silicon when the hardware became available. The company had just transitioned from being a hardware provider to a being software-and-services provider, and the virtual platform allowed it to ensure that customer use cases in system- and application-level scenarios were factored into the verification process, which in turn allowed it to reduce the number of iterations required and hence the overall cost and time for development.

Establish the right organization

When it comes to organizing their verification efforts, organizations should have a centralized way to manage verification methodology and architecture development. Activities in these areas should be under the direction of a senior verification leader, aided by a small team, who collaborates with various stakeholders in the organization. He or she should delineate verification standards—giving teams clear targets and well-defined outcomes while affording them the freedom to use the approach that works best to meet those standards. In this way, the leader can encourage teams to think strategically and make decisions based on a common, company-wide understanding of objectives rather than project teams' sometimes insular understanding of what needs to be accomplished.

Additionally, companies may want to create a centralized team for system-verification tasks but maintain a decentralized one for the module-verification tasks that are part of intellectual-property design and development. Consider the development of a wireless system on a chip: the teams responsible for designing the individual radio-frequency and baseband modules would also be responsible for verifying those parts of the chip, but a central systems team should take charge of verifying the transfer of data among these and other modules. The module-level team would require engineers skilled in intellectual-property function and protocols for intellectual-property verification. The system-level team, by contrast, would need engineers familiar with chip architecture, applications, and customer use cases for system verification.

• • •

Future system-on-a-chip advances could be at risk if semiconductor companies fail to address the current verification crunch. Lacking an intervention, SOC projects could end up devouring so many resources that only a few major players can still afford to play the game. But while the SOC verification challenge is real, it may also provide opportunities for semiconductor companies to differentiate themselves competitively in the marketplace. They can use these ideas to reduce both their costs and time to market while ensuring high levels of product quality. In the fast-moving semiconductor industry, that combination could be unassailable.

Aaron Aboagye (Aaron_Aboagye@McKinsey.com) is a principal in McKinsey's New Jersey office, **Mark Patel** (Mark_Patel@McKinsey.com) is a principal in the San Francisco office, and **Nitin Vig** (Nitin_Vig@McKinsey.com) is a consultant in the Chicago office.

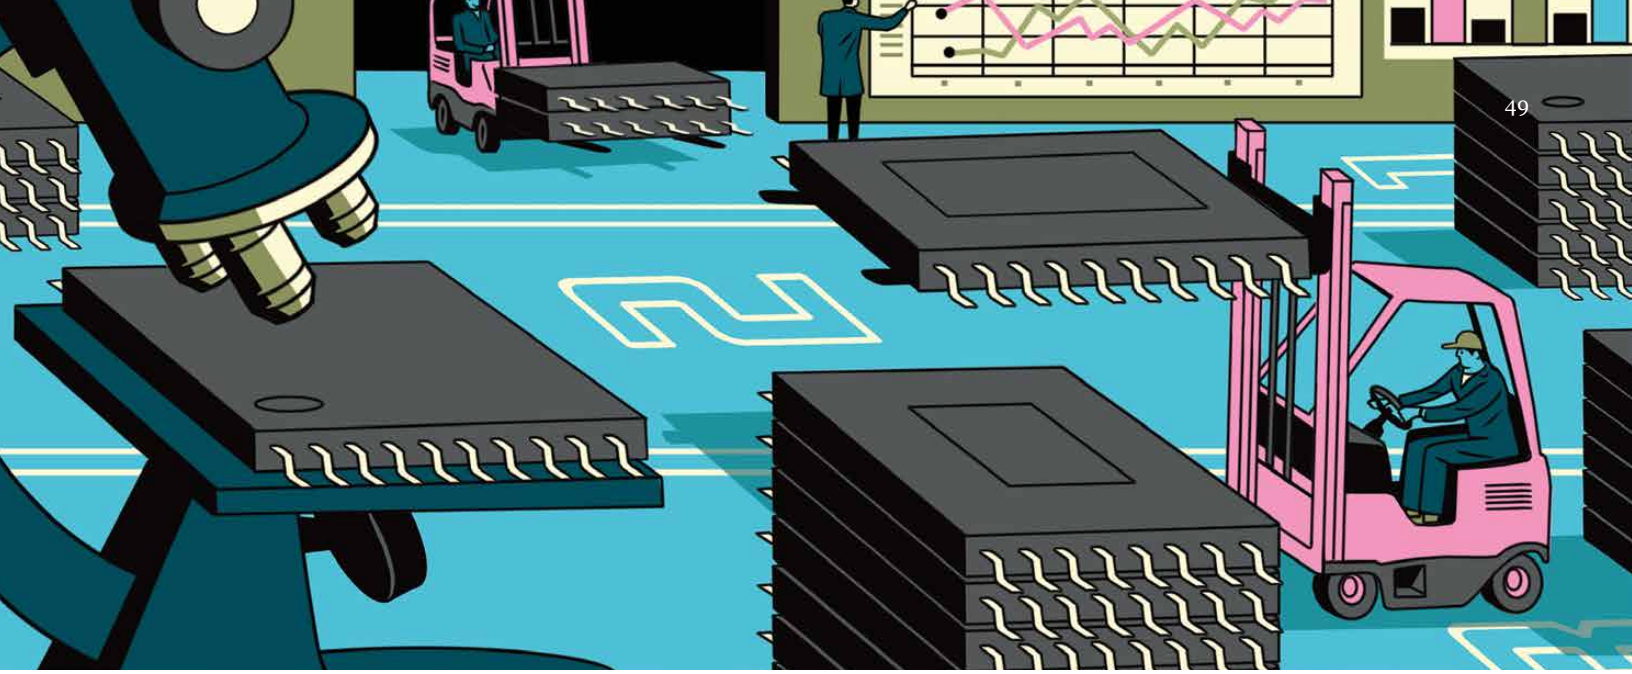

Bill Butcher

By the numbers: R&D productivity in the semiconductor industry

Four insights on the people, places, and processes that could help companies optimize output.

Aaron Aboagye, Dorian Pyle, and Alexander Silbey

Most integrated-chip-development projects are late to market, with more than half of them falling more than ten weeks behind their planned delivery dates.[1] Why is this so? Our analysis of more than 2,000 projects at more than 75 companies suggests that semiconductor executives and project teams routinely overestimate how productive they are and underestimate the complexity associated with their R&D efforts. As a result, they end up falling short on staff and other resources required to complete existing projects on time and to develop and launch new R&D initiatives.

"Productivity" generally refers to a ratio of output generated versus labor and other resources expended. Measuring the amount of resources used in semiconductor development is relatively straightforward. Measuring the quantity of output produced, however, is not. Output can vary tremendously within a single R&D organization—one team might develop 22-nanometer/5-gigahertz microprocessors, and another might develop 0.25-micron analog sensors, along with a number of other devices. This variability has traditionally made it difficult for semiconductor executives to get a clear, consistent read on their development efforts and find opportunities to improve.

What's more, most semiconductor R&D teams tend to rely on gut-feel estimates of complexity,

using qualitative up-front estimates to assign subjective labels to activities—for instance, designating a certain impending change as a "minor modification" or a "derivative release." Their estimates often do not properly account for all the nonlinear activities involved in product development, the increased complexity (even in seemingly simple updates), and interdependent project-team relationships.

The advent of big data and advanced analytics is making it easier to address the variability and complexity associated with semiconductor R&D. We have worked with semiconductor project teams to implement a "complexity index" in their R&D organizations—using historical project and process data to compile absolute measures of projects' technical characteristics, technical difficulty, and total development effort, and normalizing the differences among projects. As a result, managers can more accurately benchmark projects across the company and against industry peers. Armed with data, they can better assess risk and can reprioritize resources and projects accordingly—thereby significantly increasing their odds of on-time delivery.

Indeed, our quantitative look at R&D productivity in semiconductor companies has revealed four critical insights relating to the people, places, and processes required to optimize output.

Team productivity is strongly (and negatively) correlated with team size

Academics have long asserted that productivity is a function of team size, noting that output decreases as larger teams are mobilized. Our analysis supports that assertion. We considered R&D organizations in two different integrated-circuit markets: three organizations designing integrated circuits for the automotive sector and three organizations producing them for the wireless sector (Exhibit 1). In each case, the R&D organizations' productivity decreased as project-team size increased. The lesson? Companies can accelerate an R&D project by throwing more bodies at it, but each additional person tends to have diminishing effects. Put simply, every project has a natural limit beyond which adding more people does not increase throughput.

Each development site added reduces R&D productivity

As semiconductors incorporate more features, and thus more complexity, into their designs, it can be difficult for R&D organizations to assemble large enough teams on one site to handle new process steps. The company may decide to expand the project to multiple sites, simply to get to critical mass. However, semiconductor executives often don't have the tools and metrics that would allow them to consider the long-term effects of this decision—which can be quite significant—on productivity and schedules. Our research suggests that when companies expand teams from one site to three, productivity can drop by about 20 percent (Exhibit 2). The management practices and team dynamics that may have been effective in lower-complexity, single-site projects no longer work when far-flung team members are charged with managing increasingly intricate development tasks.

By using advanced analytics, semiconductor executives and R&D project-team leaders can explicitly account for a potential multisite penalty before deciding whether to expand. A Pareto analysis,[2] for instance, could help them quantify a project's complexity, balancing the costs associated with implementing certain process steps against potential returns on those investments. Using these data, company leaders could

Exhibit 1

Productivity on semiconductor teams usually falls as the size of the teams increases.

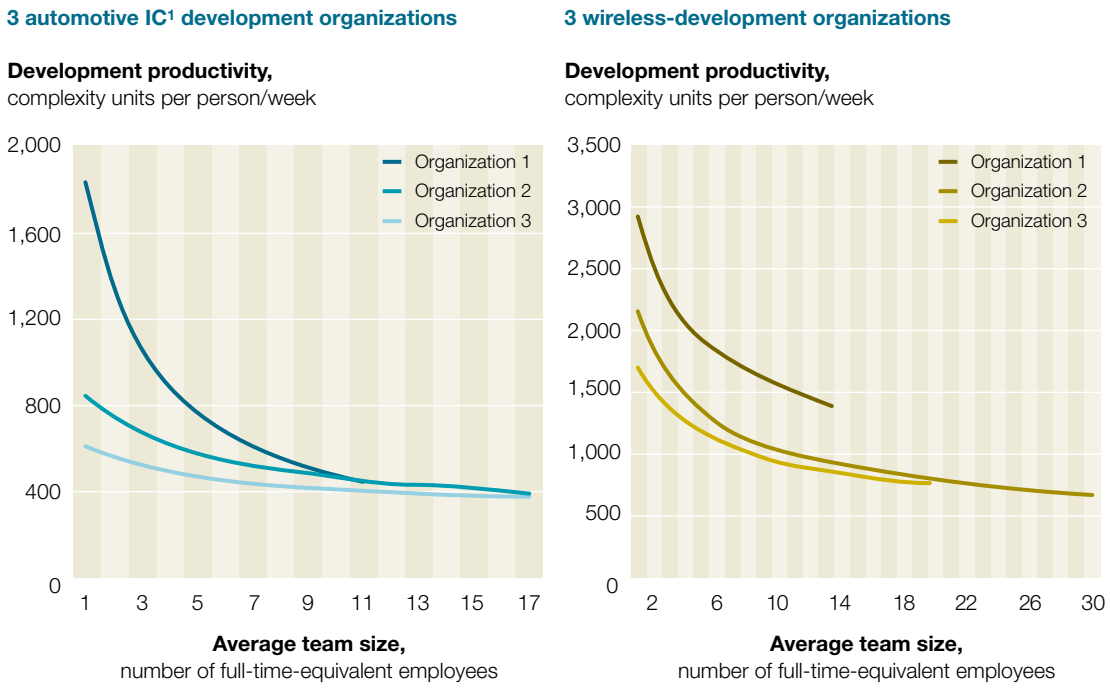

[1] Integrated circuit.

target the minimum complexity needed to satisfy market requirements. In turn, they could reconsider project-team composition—and likely assemble smaller teams in fewer sites. One semiconductor company was able to increase its productivity by 30 percent by downsizing from more than six sites to only three; functions and tasks were consolidated and partitioned among high-functioning units at the three core sites.

Don't make assumptions regarding the 'build or reuse' question

R&D organizations will often attempt to reduce cycle time and development costs by building a robust portfolio of standardized technology blocks with open interfaces and validated functionality. In this way, they can minimize the number of different design versions required and quickly turn these building blocks into a final product. But sometimes project teams need to modify these blocks because they don't have quite the right feature set or performance specs. The question then becomes, how much time and effort will these modifications take? In our interviews with several hundred design managers, most believed that reusing 50 percent of the design would save 50 percent of the development effort—a reasonable assertion. But our analysis of more

than 35,000 intellectual-property blocks suggests something very different. The relationship between reuse and effort is not linear. Instead, effort actually grows with modest amounts of reuse and then tapers off rapidly with high amounts of reuse (Exhibit 3). Furthermore, the assumption that a little reuse is better than none at all is not supported by our data: the numbers show that, no matter the type of circuit being developed, there is often little benefit when less than 40 or 50 percent of schematics are reused.

Consider the effects of time spent in all development phases, not just in design and verification

At most semiconductor companies, executives and R&D project teams spend heavily on design tools, engineering skills, and research methodologies associated with the middle and later stages of component development, when design and verification teams are fully ramped up. This focus is necessary for companies to stay competitive, but it shouldn't come at the expense of other parts

Exhibit 2

Development teams that span multiple sites can be up to 20 percent less productive.

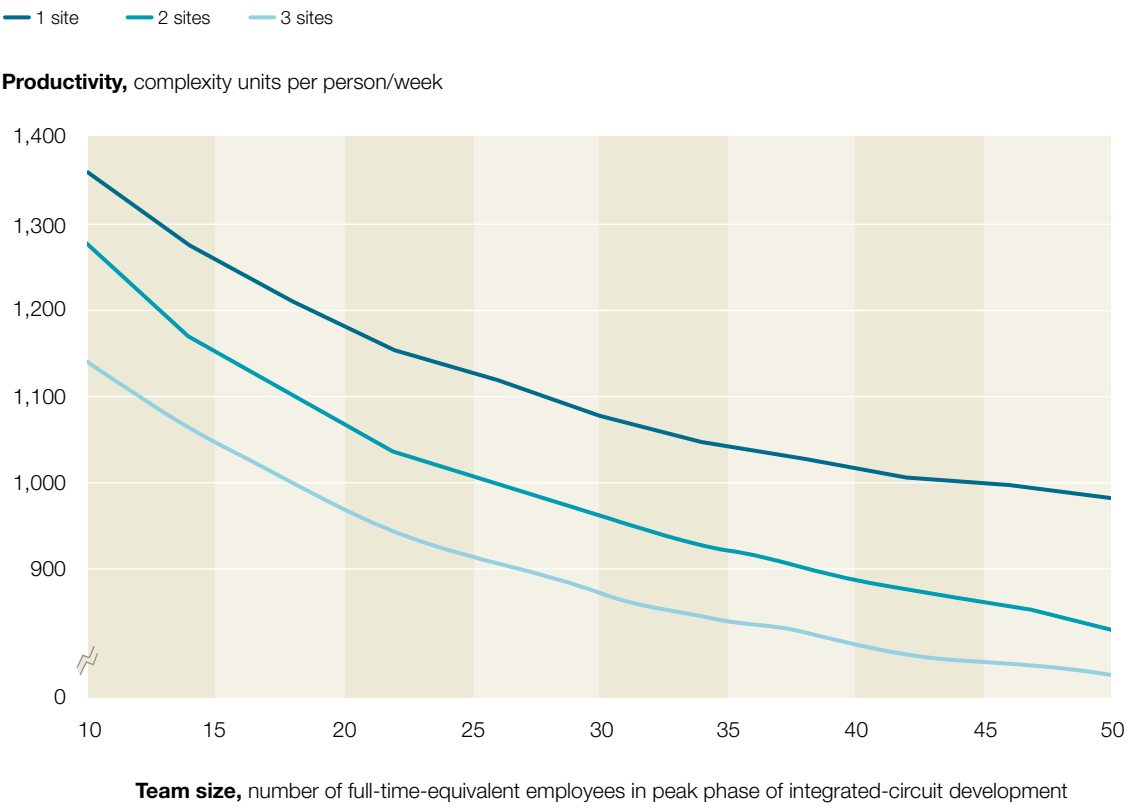

of the cycle, which our research suggests can have an enormous effect on time to market. Semiconductor players may be missing out on opportunities to cut weeks, or even months, from predevelopment phases of production. Based on our research on more than 2,000 integrated-circuit projects at more than 75 companies, for instance, the bottom quartile of companies is taking an average of 40 weeks for specification tasks while the top quartile is taking only 10 (Exhibit 4).

One R&D organization's time to market lagged behind its peers by more than six months; as a result, the company's market share and revenues were slipping. A closer benchmarking analysis demonstrated that the biggest contributor to the delivery gap was the number of projects the R&D organization had started with "fuzzy" front-end development. These projects tended to spend three calendar quarters on the drawing board before execution began, while peers' projects took less than one quarter to make that leap. As a result of this exercise, the R&D group implemented a project-introduction process that facilitated early interaction among design engineers, the marketing team, and lead customers. With the launch of this new process, the R&D group was able to

Exhibit 3

Project teams' expectations about their ability to reuse existing intellectual property are often overly optimistic.

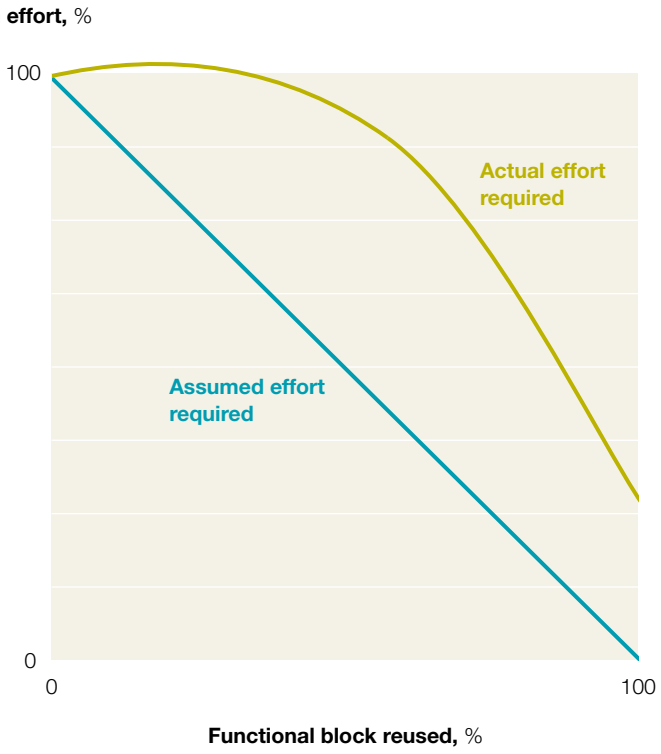

Exhibit 4

Project teams often miss opportunities to optimize processes in specification and post-tape-out phases.

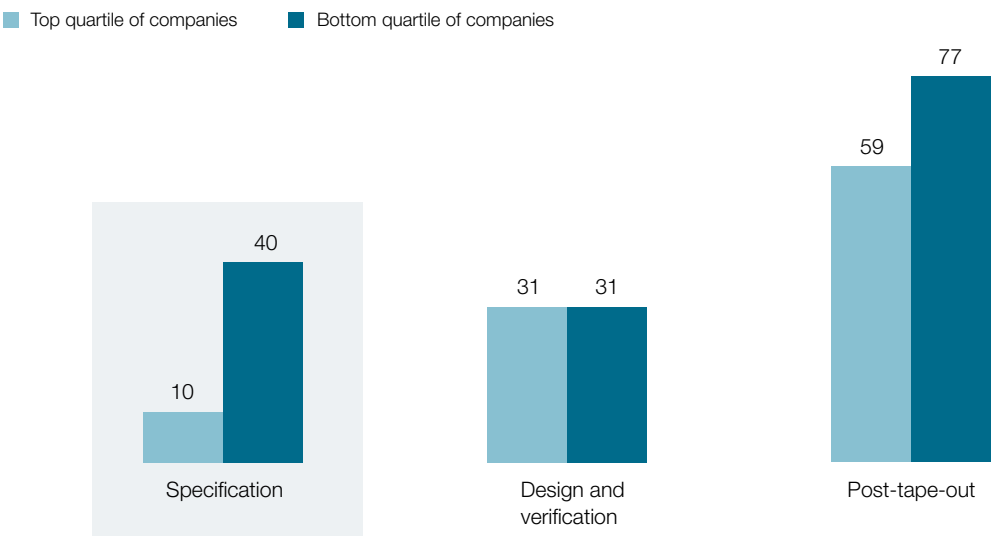

sharpen its front-end development capabilities, improve its time to market on most projects, and regain its foothold in a competitive market.

• • •

These findings point to the need for lean R&D organizations, where project teams are co-located, limited to only the optimal number of team members required, and kept staffed according to plan for the entire life cycle of the project. They also highlight the importance of using data to rationalize investments and strategic decisions; given the variability in output at most semiconductor companies, gut-feel approaches are simply not rigorous enough. Semiconductor R&D project teams must necessarily be focused on innovation and creating next-generation product features. Using advanced analytics, however, these teams can address cost and viability factors related to their innovations. They can present realistic estimates about what they can launch and when, which can give them an advantage when competing for scarce development dollars.

[1] From McKinsey analysis of more than 2,000 integrated-circuit-development projects.
[2] A Pareto analysis is a decision-making technique for determining which project inputs and other factors are having the greatest effect on the project's outcome, whether positive or negative. It is based on the Pareto Principle, which states that for many events, about 80 percent of the effects come from 20 percent of the causes.

The authors would like to thank Ondrej Burkacky for his contributions to this article.

Aaron Aboagye (Aaron_Aboagye@McKinsey.com) is a principal in McKinsey's New Jersey office, **Dorian Pyle** (Dorian_Pyle@McKinsey.com) is a consultant in the Silicon Valley office, and **Alexander Silbey** (Alexander_Silbey@McKinsey.com) is a consultant in the Chicago office.

Advanced-packaging technologies: The implications for first movers and fast followers

Adoption of 3-D technologies appears inevitable, creating both opportunities and risks.

Seunghyuk Choi, Christopher Thomas, and Florian Weig

The commercial reality for most integrated-circuit (IC) manufacturers is that node migrations and changes in wafer sizes are slowing down even as capital expenditures are increasing. One way for manufacturers to preserve their edge on their circuits' small sizes, low costs, and high performance is to incorporate newer chip-packaging options such as 2.5-D integrated circuits (2.5DICs) and 3-D integrated circuits (3.0DICs) into their production processes. These advanced-packaging technologies, many of which are still in their infancy, promise greater chip connectivity and lower power consumption compared with traditional packaging configurations.

Given these advantages, their adoption seems inevitable. According to our research, the number of integrated circuits containing 2.5DIC and 3.0DIC technologies is expected to grow tenfold—from about 60 million units in 2012 to well over 500 million in 2016 (Exhibit 1). Meanwhile, advanced packaging has become a technology priority for the Chinese semiconductor industry, according to the high-level policy framework released by the State Council of the People's Republic of China in June 2014. The council aims to have advanced packaging account for about 30 percent of all packaging revenues earned by Chinese vendors by 2015.

But there is still a lot of uncertainty in the market about 2.5DIC and 3.0DIC technologies—for instance, when and how exactly to adopt these newer packaging configurations, who will dominate among the players, and the role China will play. There are significant risks and investments (of time and money) associated with being an early adopter—the first movers will need to help reduce multiple technology standards to only a few, for instance, and will need to reconsider their roles within the manufacturing value chain. Companies in all semiconductor sectors (for instance, memory suppliers, logic producers, foundries, and packaging subcontractors) must explore strategic alliances and partnerships to ensure that a viable ecosystem for advanced packaging develops. For IC manufacturers, foundries, and others, there is also the potential to gain defensible leads in pricing and volume against rivals. Semiconductor players are therefore facing critical decisions when it comes to advanced packaging, choices that will be more or less complex depending on whether they aim to be first movers or fast followers.

Technology and market overview

Before making any strategy or process changes, semiconductor players must consider where the advanced-packaging market has been and where it is going.

The process. For IC manufacturers and foundries, end-stage packaging represents the smallest and least profitable component of the semiconductor manufacturing process (see sidebar, "What is advanced packaging?"). The entire packaging process engenders a series of front-end, middle, and back-end activities that are carried out after the integrated circuit has been designed but before chip testing begins. Critical packaging activities from start to finish include drilling (etching, lithography, and insulation), copper filling of the insulated hole to enable connectivity, grinding the surface of the wafer to expose the copper pillar (also called reveal), bumping the pillar to soften the surface, chip stacking, and chip testing.

IC manufacturers tend to manage many of the front-end activities in this process, but most of the midstage and back-end activities are performed by foundries that specialize in outsourced assembly and testing (OSAT). Compared with the integrated-device-manufacturing (IDM) market, the OSAT market is much more fragmented; the combined sales of the four companies that lead this segment account for only 45 percent of the entire OSAT market. OSAT players have lower profit margins (about 20 percent versus 40 percent for IDMs) and higher material and labor costs, and they primarily compete on operational efficiency rather than innovation.

2.0DIC technology. Existing 2-D integrated-circuit (2.0DIC) flip-chip and wafer-level packaging technologies have shown solid growth over the past five years and are used in a number of mainstream applications—predominantly in high-end smartphones (the iPhone and Samsung Galaxy, for instance) and tablets, which must meet stringent size and power-management requirements. Flip-chip packaging involves applying soldered bumps on the top side of a fabricated wafer; the integrated circuit can then be flipped and aligned with grooves on an external circuit to enable the necessary connections. This form of packaging occupies less space in products and offers higher input/output rates, because the whole surface area of the chip can be used for interconnection instead of just the outside edge, as is the case with traditional

wire-bonding methods. In wafer-level packaging, the integrated circuit is packaged while it is still part of the silicon—meaning the package is the same size as the die, and the manufacturing process is streamlined, because conductivity layers and solder bumps are applied to the integrated circuit before dicing occurs.

2.5DIC and 3.0DIC technologies. Emerging 2.5DIC and 3.0DIC technologies promise to extend flip-chip and wafer-level capabilities, enabling multiple dies to be stacked vertically together through the use of interposers and through silicon via (TSV) technology. The TSV stacking technology allows for a greater amount of functionality to be packed into the chip without having to increase its size, and the interposer layer (which essentially performs a routing function) serves to shorten critical electrical paths through the integrated circuit, creating faster input and output. So according to our estimates, an application processor and memory chip encased using advanced-packaging technologies would be about 30 or 40 percent smaller and about two or three

Exhibit 1

2.5DIC and 3.0DIC technologies are growing.

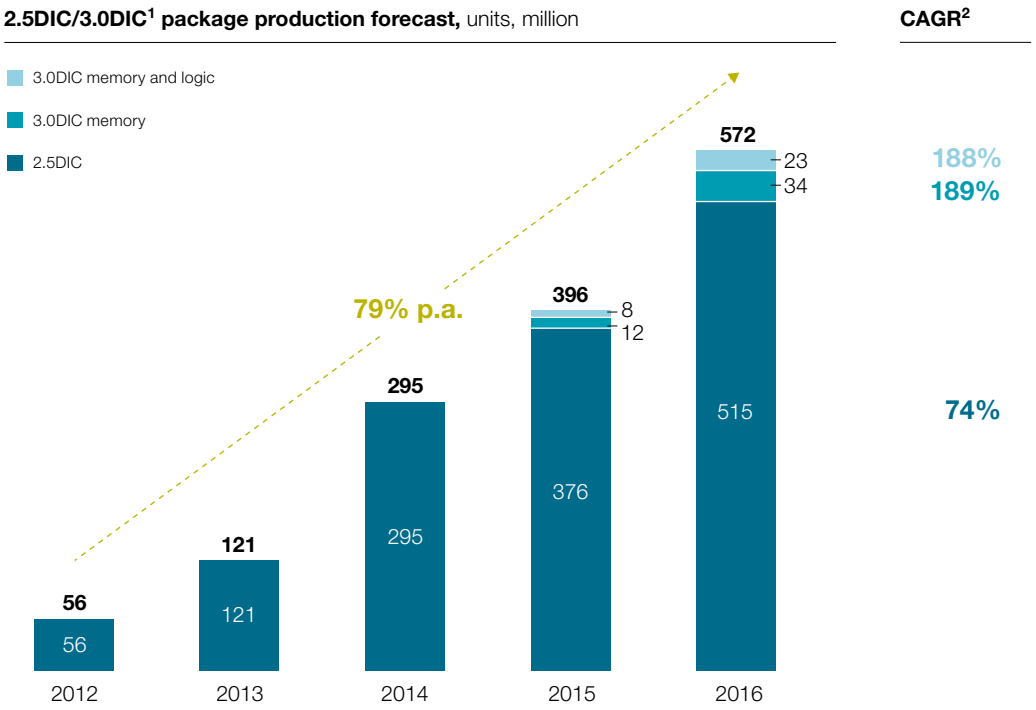

[1]2.5-D integrated circuits/3-D integrated circuits.
[2]Compound annual growth rate. Figures have been rounded up.

Source: Gartner; New Venture Research; McKinsey analysis

What is advanced packaging?

During the final stages of semiconductor development, a tiny block of materials (the silicon wafer, logic, and memory) is wrapped in a supporting case that prevents physical damage and corrosion and allows the chip to be connected to a circuit board. Typical packaging configurations have included the leadless chip carriers and pin-grid arrays of the 1980s, the system-in-package and package-on-package setups of the 2000s, and, most recently, 2-D integrated-circuit technologies such as wafer-level, flip-chip, and through silicon via setups (exhibit).

Exhibit

Integrated-circuit packaging has evolved since the 1970s.

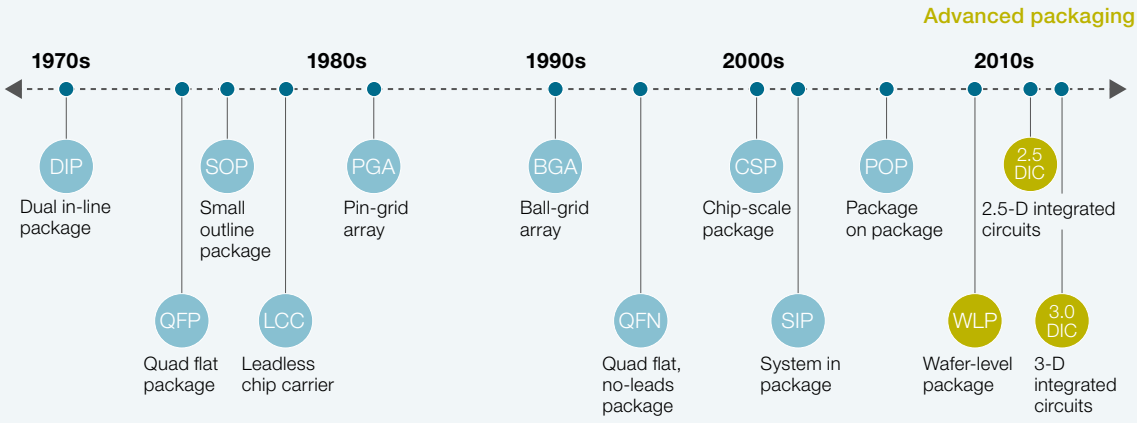

Source: IC Insights; Yole Développement; McKinsey analysis

times faster than a chip packaged using older technologies and may create power savings of up to 40 percent or more. Demand for 2.5DIC and 3.0DIC technologies is dependent upon a range of factors, of course, including a thriving market for low-end smartphones, tablets, wearable devices, and other connected consumer goods, as well as an ecosystem in which multiple semiconductor companies (not just a few big players) are committed to upgrading to newer packaging technologies.

How will the market unfold?

The sophistication of 2.5DIC and 3.0DIC technologies, and the economics for the IC manufacturers and OSAT players that produce them, means that IDMs and foundries will still need to handle the front-end work, while OSAT players will remain best suited to handle the back-end processes, such as via reveal, bumping, stacking, and testing. The latter activities rely on interposer manufacturing, a cost-sensitive process with low technical

Exhibit 2

Who owns the gray area?

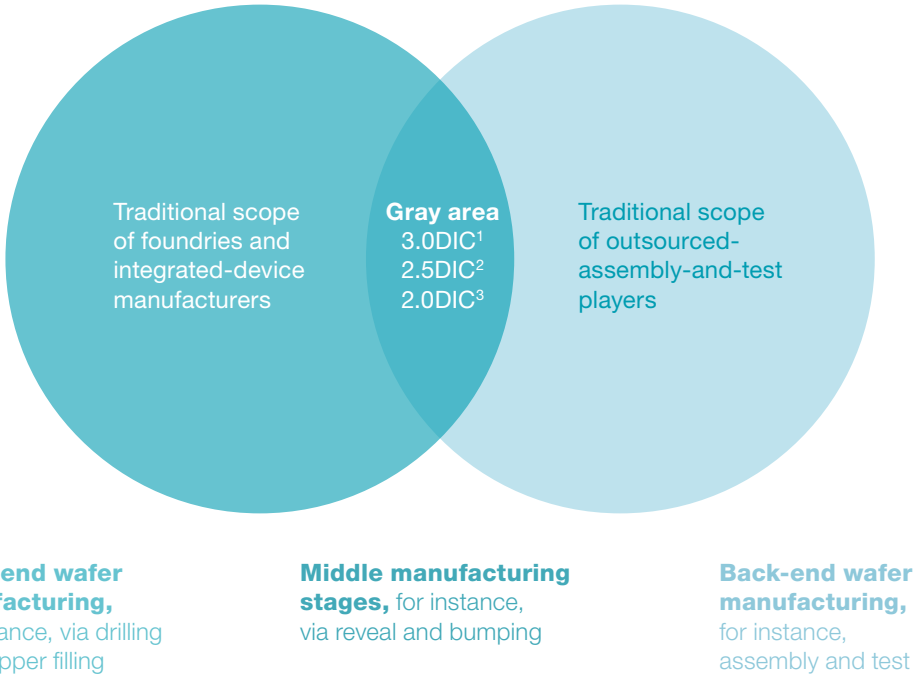

[1]3-D integrated circuits.
[2]2.5-D integrated circuits.
[3]2-D integrated circuits.

Source: IC Insights; Yole Développement; McKinsey analysis

requirements. But there is a gray area emerging midstream, as Exhibit 2 shows, and IC manufacturers may need to reconsider their role in this stage of production, exploring the trade-offs between taking on higher process and implementation costs and gaining improved performance and competitive advantage through early adoption of 2.5DIC and 3.0DIC technologies.

Indeed, the market probably will not move monolithically; different segments likely will make the transition based on the relative benefits of investment and the level of competition. The IDMs and foundries that produce high-end application processors, higher-end image sensors, enterprise memory devices, graphical processing units, and central processing units will likely be among the first to make a move. In fact, some leading-edge graphical processing units and high-end memory products are already in the early-adoption phase. But those that traffic in integrated circuits for lower-end products, such as basebands

for low-end to midrange handsets, will likely transition much later (Exhibit 3). Early adopters would likely include companies such as Intel, Samsung, and Taiwan Semiconductor Manufacturing Company—those with enough scale to drive up volume, bring down costs, and reduce the risk enough so that others will follow suit. Fast followers may then find it easier to make the transition but may also be limited to collaborations with first movers as their only means for capturing cost and performance advantages from advanced-packaging technologies.

For their part, some OSAT foundries are also preparing for a ramp-up to 2.5DIC and 3.0DIC technologies by collaborating with larger foundries to serve fabless players. For instance, Amkor Technology, whose client base includes

Exhibit 3

3.0DIC adoption scenarios vary.

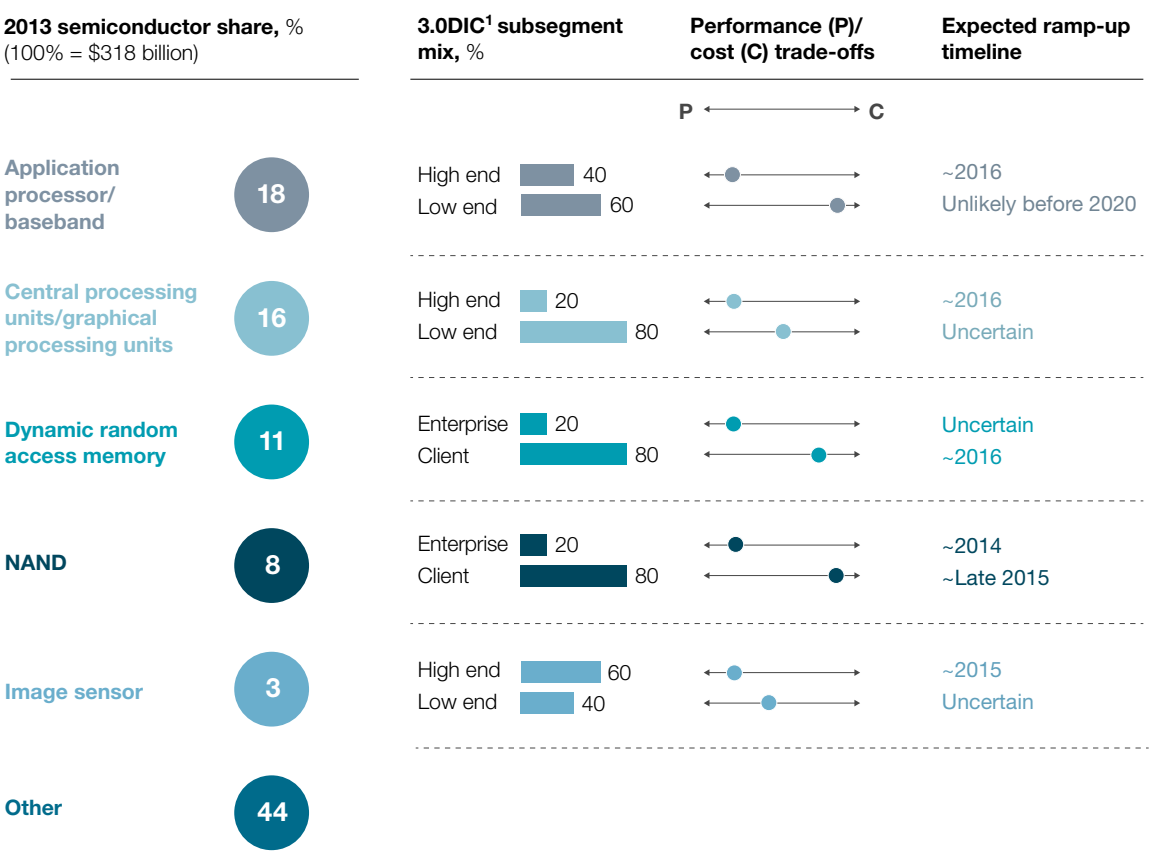

[1] 3-D integrated circuits.

Source: iSuppli; McKinsey analysis

most of the major fabless players across the globe, has been closely working with Xilinx on qualifications relating to TSV technology.

Overall, we believe two adoption scenarios could unfold. First is a slow and steady transition where semiconductor companies would gradually move from flip-chip and 2.0DIC technologies toward incorporating 2.5DIC and 3.0DIC technologies into their chips; the latter technologies would account for between 20 and 30 percent of the advanced-packaging market by 2022, but with only a few big players adopting them and implementation costs that would still be 50 percent higher than 2.0DIC costs. Second is a hard right turn in the industry, where 2.5DIC and 3.0DIC technologies would account for more than 50 percent of the advanced-packaging market by 2022, and multiple industry players would have adopted 3.0DIC technologies and collaborated to strengthen the advanced-packaging ecosystem. Implementation costs would be only 20 to 30 percent higher than those associated with 2.0DIC. A slow and steady transition is more likely, given that production costs are not dropping fast enough and potential end markets for devices that would contain 2.5DIC and 3.0DIC chips (wearables, for instance) have garnered early buzz but have been slow to develop.

Implications for first movers and fast followers

What will it mean to be a first mover in 2.5DIC and 3.0DIC packaging technologies? The early adopters will need to invest significantly in the ecosystem—hiring new engineers, for instance, or spending the time and money to establish partnerships. They will also need to find cost-effective ways to upgrade their equipment to handle newer TSV-based technologies and processes. In some cases, existing 2.0DIC machinery can be expanded to meet newer capacity requirements. But IC manufacturers and foundries may also need to purchase and install new equipment for, say, TSV etching or copper filling. We estimate that in preparation for the shift to advanced-packaging technologies, industry players may invest between $200 million and $300 million on such equipment in 2016. IC manufacturers and foundries could also address this need by entering into partnerships with equipment manufacturers to codevelop bonding, plating, and reveal capabilities that they may not have.

First movers will also need to shape the industry's discussions about packaging standards. Currently, for instance, there is no standard method for temporary bonding and debonding of integrated circuits; different plants use either laser, heat, or

The early adopters will need to invest significantly in the ecosystem and shape the industry's discussions about packaging standards.

mechanical processes to do the same job, thereby missing an opportunity to not only save costs but also minimize quality issues. First movers should consider working with other players in the advanced-packaging industry to establish common process recipes, equipment specifications, logic-to-memory interfaces, and so on. Several such partnerships and initiatives are under way. The semiconductor industry association JEDEC Solid State Technology Association (formerly the Joint Electron Device Engineering Council) for several years has been working toward a standard for the use of 3.0DIC packaging technologies in IC manufacturing. In addition, GLOBALFOUNDRIES has developed the Global Alliance for Advanced Assembly Solutions to accelerate innovation in semiconductor connection, assembly, and packaging technologies; alliance members include Amkor Technology, ASE Group, and STATS ChipPAC in the assembly-and-test area.

For their part, fast followers can mitigate risks and minimize investments as first movers take the lead. As 2.5DIC and 3.0DIC technologies take off, however, fast followers will likely want to get back into the fray. They will need to closely monitor the first movers' activities, participate in discussions regarding standardization, and keep the lines of communication open with customers to gauge their needs in advanced packaging. They may also want to track potential M&A partners—for instance, TSV equipment makers.

• • •

Collaboration among OSAT players, IDMs, foundries, and others in the semiconductor market will be critical for building a reliable advanced-packaging ecosystem—one that recognizes the importance of scale, second-source providers of packaging services (to preserve customer choice), and strategic alliances among memory suppliers, logic IDMs, foundries, and subcontractors. It will be an important factor in allowing companies to optimize their returns on advanced-packaging technologies and ensure continued innovation.

The authors would like to acknowledge Chris Lim and Bill Wiseman for their contributions to this article.

Seunghyuk Choi (Seunghyuk_Choi@McKinsey.com) is an associate principal in McKinsey's Seoul office, **Christopher Thomas** (Christopher_Thomas@McKinsey.com) is an associate principal in the Beijing office, and **Florian Weig** (Florian_Weig@McKinsey.com) is a director in the Munich office.

Fab transformation: Four markers of excellence in wafer production

To succeed with their lean initiatives, managers should focus on improving plant uptime, equipment utilization, process variability, and product quality.

Cosimo Corsini, Tommaso Debenedetti, and Florian Weig

Achieving excellence in complex semiconductor manufacturing environments is difficult. The chip-fabrication process involves numerous, nonlinear steps and stages, and it requires advanced technologies that must be deployed in ultraclean environments that are expensive to set up and maintain. Demands from clients can change quickly—a car manufacturer may want to switch to a different type of chip in its vehicles, integrating that component into its production lines within two months. Plants must be able to continually adjust their production recipes, schedules, and priorities to accommodate these new requests.

To crack the code of excellence, some semiconductor companies have tried implementing lean principles, with varying levels of success. The core concepts behind lean programs are well established and fairly straightforward. But as many semiconductor players have learned, the application of lean principles on the shop floor is much more complicated than it seems. To successfully modify some of the most advanced and difficult production processes in the world of manufacturing, managers and equipment operators must show full dedication to the change effort, but that focus can be hard to maintain when there is pressure to improve performance immediately

and when both managers and employees are skeptical about proposed process changes.

We have found that refocusing lean efforts on four critical dimensions of plant operations—uptime, utilization, process variability, and product quality—can provide a jumpstart. To increase the odds of maintaining process improvements over the long term, managers should also establish a culture that relies on data analysis, problem solving, and cross-functional collaboration. One large manufacturer of eight-inch silicon wafers that adopted this approach was able to increase its output by more than 25 percent and decrease its cycle time by 20 percent. The team at this fab did not need to invest in more equipment or increase its head count to achieve these goals. Instead, fab managers systematically reviewed plant processes and behaviors and, in response to their findings, adopted new, lean practices. In this way, they were able to improve equipment reliability and uptime, work-flow management, plant agility, and product quality. The company sought to re-create itself as a lean organization—and its investment in this pursuit showed significant returns within 18 months.

Let's take a closer look at what we'll call Fab X, the challenges it faced, and the actions it took to improve operations—actions other semiconductor companies may be able to emulate.

Facing production challenges

Fab X was seeking to increase its moves per day—the number of times a wafer advances from one step in the manufacturing process to the next—but, for a variety of reasons, activity was stalled below target. Plant leaders had publicly stated their desire to adopt a lean approach and improve the company's efficiency and effectiveness, but that philosophy was not reflected on the shop floor.

A close assessment of operations at the semiconductor plant revealed that there was no shared understanding among managers across the plant of where bottlenecks were occurring, and there was too little time spent conducting timely, detailed analyses of overall equipment effectiveness, given the work in progress. The lead team was more likely to try to find a temporary fix for a faulty machine so it could meet a weekly production quota, rather than task a team to explore root causes of the problem and get rid of it once and for all. Every project was "urgent"; too many work streams and activities were being launched at the same time, with limited or no time allotted to appropriately assess outcomes. Managers did not prioritize projects, nor did they monitor quality in any systematic way. So cycle times increased while volumes decreased.

Meanwhile, a detailed look at the organization overall revealed there was little communication among senior managers situated in a shop that was hierarchical in nature. The senior leaders were technicians with deep knowledge about product and equipment specs, and they valued that form of organizational capital above all else. They failed to recognize the importance of gathering input on the production process from all levels of the plant and across functions and were missing the signs that employees were confused about the performance feedback they were being given and the direction in which the plant was going. Managers did not see the long-term advantages of creating an inclusive work environment that would engage employees and establish a culture

of continuous improvement. The resulting low morale contributed to decreased productivity.

Focusing on four markers of excellence

Fab X's experience was not unique; these are the perennial problems for the industry. Oversight of complex enterprises requires a very high level of expertise in production and line management, equipment maintenance, process and production engineering, and quality control. But Fab X was able to turn around its fortunes by optimizing its performance in the four critical areas of plant operations mentioned earlier.

Uptime. All fabs tend to experience two main production delays—when machinery goes offline for scheduled repairs, and when it shuts down unexpectedly. To address the former, Fab X introduced a new scheme for planning equipment maintenance based on advanced analytics. After the production of a certain number of wafers, cleaning must take place. The information managers were using to determine the optimal time for this changeover had been incomplete—different units collected and recorded the information using different methods. Fab X now uses sensors and tags embedded in its equipment to collect data that can then be run through various simulations—asking, for instance, what will the impact be if we take down a high-temperature furnace on nights and weekends or at certain hours? The plant is also relying more heavily on tried-and-true lean production methods such as the single-minute-exchange-of-die process, which emphasizes quick change of parts used at various stages in the manufacturing process—altering the sequence of part replacements, for instance, or automating various replacement steps. In the case of unscheduled outages, Fab X recognized it could not necessarily plan for every shutdown possibility, but managers did implement structured problem-solving sessions focused on figuring out exactly what went wrong. Previously, senior managers would have spent the time justifying among themselves what happened rather than trying to fix it. By contrast, their daylong discussions of root causes—which involve fab managers and representatives from across all functions—have allowed the fab to realize an almost 70 percent reduction in equipment downtime (both scheduled and unscheduled).

Utilization. Another production-cost challenge for fab managers is minimizing standby, or the time a machine tool is available for use but not actually in operation. Tools that are perpetually in standby mode can cost the plant thousands of dollars per minute. At most plants, managers may try to address production shortfalls by investing in more tools, even though the existing ones are being underutilized, or firing and then hiring new line staff, hoping they will do things differently. Fab X was able to increase the utilization of tools in its plant by combining quantitative and qualitative research to redesign work flows, redeploy existing staff, and standardize certain shop-floor activities. To determine the right number of people needed to operate each piece of equipment in each of its production bays, for instance, managers shadowed shop-floor operators, recorded their observations, and discussed their findings with shift leaders and operators. The critical part of this process was collecting feedback from the operators and convening team discussions to foster continuous improvement. In these discussions, fab managers learned that handoffs between operators on a

given tool and between operators handling different parts of the fabrication process were a big time sink. So they considered the optimal times required for shift changes, breaks, and other shop-floor activities that were indirectly related to production. Based on these data, managers standardized their transfer activities and created schedules that allowed them to allocate the right resources at just the right times. Through its efforts to calculate staffing needs from the bottom up and reallocate operators more effectively, Fab X was able to reduce its standby times by 70 percent.

Variability. Fab managers must maintain a careful balance in work flow. One small bottleneck in the wafer-production process can throw off lead times and performance across the entire plant. To better manage the work in progress, leaders at Fab X assessed equipment utilization rates and cycle times, and identified several machines that, given their history of outages, alarms, and operator issues, had the potential to become huge bottlenecks as demand increased. Just as they had in their uptime analysis, the fab managers convened root-cause discussions, pulling in representatives from different functions—for instance, production, engineering, maintenance, and quality control—to assess the critical reasons for variability among some of the machines and to develop a plan for boosting overall equipment effectiveness. As a result of their collaboration and analysis, the Fab X team revised the dispatching rules associated with the challenged equipment—for instance, requiring the system to deliver a wafer faster or immediately—and took other steps to increase capacity. Through these efforts, the plant was able to minimize bottlenecks, improve its overall work flow, and reduce its overall cycle-time variability by up to 15 percent.

Quality. Often fabrication plants seeking to increase production and reduce cycle time believe that they will need to make small sacrifices in quality to do so. This is false; process improvements do not need to come at the expense of quality. Lean principles applied to improve manufacturing operations will indirectly affect the quality of the semiconductors being produced. To diminish the effects of chronic quality issues, managers at Fab X focused on identifying core process and product flaws. A critical point for shop-floor personnel was to identify errors where they are generated and not at the end of production, when other components have already been added to a cracked wafer. To do so, fab managers compared the process steps during which errors typically happened with the process steps during which errors were actually found and were able to differentiate between the early leaks in error detection, the chronic process issues that led to product flaws, and those errors that could have been avoided through root-cause analysis. (The data were drawn from the process information the plant routinely collected as part of its operations.) As a result of these findings, the plant has increased its yield and, over time, has gradually decreased its waste.

Developing lean teams and capabilities

Fab managers cannot realize the same sort of improvements in uptime, utilization, variability, and quality that Fab X did without having the right team and infrastructure to implement and support a shop-floor transformation. They must create an environment that emphasizes data analysis, problem solving, cross-functional collaboration, and execution.

Fab X introduced new tools and technologies—for instance, data-visualization tools and software

applications that would assist in the daily tracking of key performance indicators, procurement decisions, and other process parameters. The plant also reorganized its leadership structure to include a core "lean team" whose primary activity in the fab was to oversee efficiency efforts. That team, many members of which were steeped in technical rather than "soft" skills, underwent a series of workshops focused on developing competencies in coaching, planning, conflict management, and delivering and receiving feedback, among other things. The sessions involved role playing and one-on-one interactions. Additionally, another 100 employees, at different levels of the company, were trained as change agents for lean transformation, so not all the change was top down. This focus on improving the health and sustainability of the organization is ongoing, so it is still too soon to quantify the overall effect of the company's lean transformation, but Fab X has been identified within the industry as a best-practice plant.

Interviews with employees and operators at Fab X before managers there undertook a lean transformation suggested that they understood the need for change—the lag in performance was apparent—but different constituents within the plant held different beliefs about why the change needed to happen. And while all agreed that cross-functional collaboration was crucial, none felt that top management had made this a priority in its day-to-day operations. In post-transformation discussions with employees, the same respondents reported a shift away from competition among functions and shifts; clearer "rules of the road," with a reduction in the number of key performance indicators to just several crucial ones; less focus on firefighting and more feedback sessions involving people from all levels of the fab; and a robust, data-oriented approach to monitoring results and modifying processes. "The distance has closed between us and senior leaders," one operator noted.

Fabs that want to achieve lean transformation can similarly use surveys, interviews, and feedback sessions to build awareness among employees about the need for performance improvement, to educate them about lean principles and approaches, and to ensure that there is sufficient appetite and willingness to embrace this sort of change. This is not an easy or a short exercise; without a change in organizational mind-set, it can be difficult to sustain a lean program over the long term.

• • •

The production processes and activities associated with semiconductor fabrication are highly volatile and very complex, and applying lean principles in these environments can be difficult. But as Fab X learned, significant performance improvements are possible when companies train their lean efforts on four main areas—uptime, utilization, variability, and quality—and develop a corporate infrastructure that supports this focus. The fabs that do can reduce downtime and waste, increase cycle time, and improve the quality of their products over the long term.

The authors would like to thank Guido Frisiani for his contributions to this article.

Cosimo Corsini (Cosimo_Corsini@McKinsey.com) is a principal in McKinsey's Milan office, where **Tommaso Debenedetti** (Tommaso_Debenedetti@McKinsey.com) is an associate principal; **Florian Weig** (Florian_Weig@McKinsey.com) is a director in the Munich office.

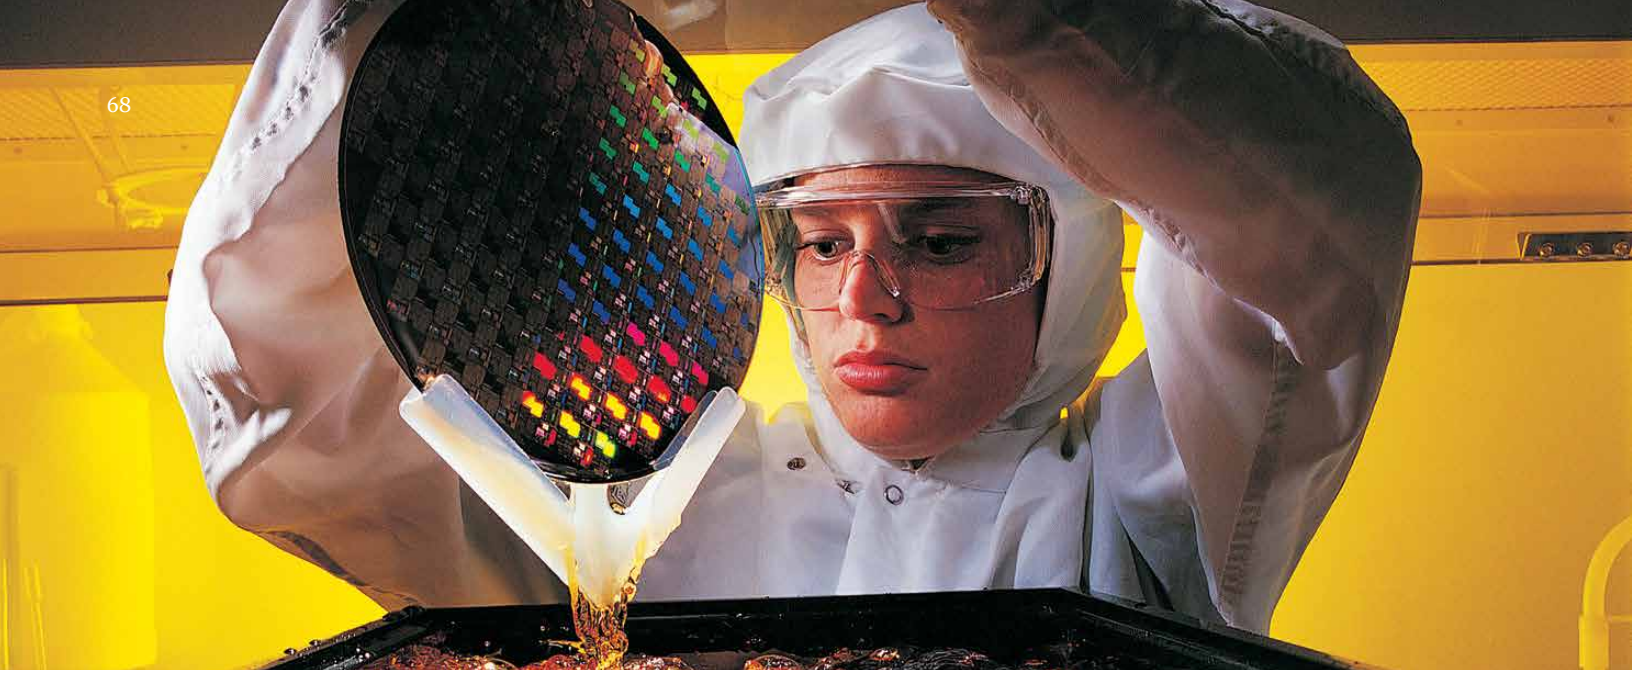

Fab diagnostics:
A data-driven approach to reining in the cost of indirect materials

Companies that use a set of core analytics to assess consumption patterns can gain better control of production expenses.

Harold Janin, Mark Patel, and Florian Weig

Indirect materials—the substrates, chemicals, slurries, specialty gases, pads, films, spare parts, and other critical ingredients used to make integrated circuits—typically account for more than 30 percent of the cost of front-end semiconductor fabrication.[1] But fab managers have had a tough time getting a handle on these expenses, in part because of the limited control they have over materials pricing and because they are more likely to examine projects, supplies, and production activities in isolation rather than considering their impact across a fab's entire portfolio.

To deal with these and other cost issues, semiconductor executives should adopt an analytics-based approach to materials cost management. By systematically assessing the data the fab collects on processes and materials, managers can better understand spending by supplier and by fab. As a result, they can emphasize cost-management efforts that may have the greatest impact, and they can undertake discussions with suppliers more confidently.

In this article, we introduce several data-centric methods that managers and engineers can use to identify cost-saving opportunities and reset priorities. Based on our experiences, these tools can help managers achieve cost savings of more than 15 percent—far better than the single-digit average savings typical even in mature, 200-millimeter fabs. The tools can provide a straight-

forward, repeatable reading on the resource situation at any fab. But implementing them successfully requires support from all the departments that are using the respective chemicals or other materials. All have a vested interest in ensuring that the fab can reduce costs year after year to keep pace with the price erosion the industry is experiencing.

Savings stumbling blocks

Why don't more fabs achieve better results from their cost-management programs? There are two main factors.

A narrow view of consumption. Resource-management efforts have tended to be ad hoc, in part because of the relentless pace of product development and the number of nodes in play. In this climate, managers evaluate costs by project, and some waste is considered part and parcel of the production process. Additionally, the decentralization of production activities often leads to a lack of coordination among semiconductor-module teams, sales teams, procurement specialists, and other units within a fab. This can result in a poor understanding of the types of chemicals required and which suppliers to target. For example, in one company, a module-engineering team in the lithography department was trying to optimize the mix of ingredients necessary for a single resist (a thin layer of polymer used to transfer a circuit pattern to a semiconductor substrate), to minimize cost overruns. The team was unaware, however, that the number of products using this particular resist was expected to fall in the near future—and that focusing on this recipe would have little impact on costs.

Limited control over resource pricing. The typical semiconductor manufacturing process involves a wide variety of chemicals; a fab may stock and use more than 50 chemicals within the lithography stage alone. But there are only a handful of established chemicals suppliers, and plants are reluctant to switch to new ones given the long lead times required to qualify them—in some cases, it can take up to a year. As a result, incumbent suppliers are shielded from price pressures, and fab managers have less opportunity to explore potentially more advantageous relationships with existing or alternate vendors.

Applying the diagnostics

There are a number of analytic tools and techniques that engineering teams and semiconductor executives can use to better manage production resources, but we believe two are particularly effective for ensuring that no cost-containment measures are left on the table: the heat-map analysis and the mass-balance analysis. The former is a prioritization tool; it gives fab managers a high-level overview of the line items associated with semiconductor production, and it allows them to spot the gaps in their management of certain chemicals and other inputs. The latter offers a deep dive into the consumption patterns revealed by the heat-map analysis, giving fab managers the information they need to make smarter, more cost-effective resource and operations decisions. When combined with other methodologies—among them, spending analyses and time-to-failure and complexity assessments—heat-map and mass-balance assessments can provide the backbone for a strong, systematic cost-management program.

The heat-map analysis. The first step in any effort to reduce costs is to know which materials are in greatest demand or have seen the most significant changes in usage over a time period being considered. Heat maps are effective for creating this level of transparency. A module-engineering

When combined with other methodologies—among them, spending analyses and time-to-failure and complexity assessments—heat-map and mass-balance assessments can provide the backbone for a strong, systematic cost-management program.

team can inventory and record all items used across the fab using the vast amount of routinely collected product and process data—albeit usually in uncoordinated fashion. With input from procurement managers, the team can then categorize and rate indirect materials and maintenance items along several dimensions relating to consumption and pricing. In this way, the team can spot meaningful gaps in their cost-control programs.

At one large fab, for instance, managers assumed they had created a comprehensive program for reducing their consumption of indirect materials, simply because of the breadth of their efforts: there were more than 100 cost-cutting initiatives going on throughout the company, most of them focused on optimizing existing product mixes. This is justifiably a common focus; we have seen many cases in which too much of an expensive chemical is incorrectly prescribed for a production process. But fab managers had not fully explored other cost-cutting opportunities focused on different cost-containment parameters—for instance, emphasizing waste reduction and considering the possibility of reducing the amount of certain chemicals used in setup and rework activities (process steps that happen in support of core chip development). By undertaking the mapping exercise, fab managers saw the gaps in their approach and inconsistencies across sites; different fabs were using different amounts of chemicals, even for the same tech nodes. Through their analysis, they were able to reprioritize their cost-cutting initiatives and, for some chemicals, the fab was able to realize savings of up to 50 percent.

The mass-balance analysis. This tool enables fab managers to drill down into the findings presented by the heat map and further delineate chemicals consumption. The goal is to create a snapshot of actual consumption patterns associated with particular ingredients compared with projected usage. Using these data, module-engineering teams and procurement managers can examine individual causes of waste.

The results of this analysis can be eye opening. One company's mass-balance analysis revealed a flawed batching process. The chemicals bath the company employed during the clean-tech stage

could accommodate up to 100 wafers at a time. Through the mass-balance assessment, however, the company recognized it was processing many fewer than that, wasting up to 40 percent of materials used in this step. Fab managers conducted workshops to generate ideas and determine how to address the challenge. By altering its batching steps and tool configurations, the company was able to improve its load factor, reduce waste, and cut costs.

• • •

There will be inevitable roadblocks to implementation: resistance to change from module engineers, a shortage of time and talent within the modules to carry out new projects, and insufficient management capacity to lead the qualification process when adding suppliers. The fabs that adopt this approach may also require new technology systems for collecting data, as well as analysts and engineers who can perform regressions and other forms of data mining. They may also need to bolster capabilities in portfolio management; the analytical approach we are suggesting may turn up more cost-containment projects than fabs will have the time and resources to execute, and managers will need to focus on the projects with the biggest impact.

We cannot understate these challenges, but they should not stop fab managers from exploring analytics-based cost-reduction programs. Even small reductions and improvements will help put fabs in a better long-term cost and operations position.

[1] Front-end fabrication refers to the process of forming transistors directly in the silicon wafer.

Harold Janin (Harold_Janin@McKinsey.com) is a consultant in McKinsey's Calgary office, **Mark Patel** (Mark_Patel@McKinsey.com) is a principal in the San Francisco office, and **Florian Weig** (Florian_Weig@McKinsey.com) is a director in the Munich office.

Bill Butcher

Beyond the core: Identifying new segments for growth through value-chain partners

A systematic process for assessing supplier and customer capabilities and relationships can help semiconductor companies identify adjacent markets and promising opportunities.

O Sung Kwon, Mark Patel, and Nicholas Sergeant

It is becoming increasingly difficult for semiconductor players across all sectors—whether in manufacturing, capital equipment, or chip design—to find business opportunities beyond their core customers and products. The current economics of the industry simply do not support companies' efforts to dabble in new areas. The necessary investments in the core carry a high price tag. The costs of creating new platforms, engineering for the next node dimension, or developing new integrated-circuit (IC) designs are now reaching the billions. And most of the new market segments targeted by chip and equipment manufacturers will inevitably have significant barriers to entry—not the least of which are unfamiliar operating models and sales channels, and aggressive incumbents.

Given these challenges, some companies are looking for growth opportunities closer to home and finding unexpected resources in their existing networks. They are partnering with suppliers and customers who participate in adjacent market segments or acquiring technology from elsewhere within the value chain (Exhibit 1). Intel did just that with its 2010 acquisition of McAfee, a deal that many in the mainstream and technology trade press described as a strategic move by the chip maker to grow "outside of the PC and computer server

markets" and gain a toehold in smartphones and consumer electronics.[1] There have also been several large, high-profile deals in the market for NAND flash-memory technology among vendors looking to broaden their customer base (Exhibit 2).

Based on our work over the years with a number of global semiconductor companies and our research on mergers and acquisitions in high tech, we have identified a process for assessing potential opportunities for growth through value-chain relationships, partner capabilities, and adjacent applications. In this article, we focus on ways to create opportunities among suppliers. It is typically easier for semiconductor companies to look upstream first because they will inherently understand their suppliers' businesses better. By contrast, an evaluation of customer-focused opportunities will likely require more time and resources.

There are three steps semiconductor companies can take to spot growth opportunities: identify the complementary market segments that could provide growth, identify the suppliers and technologies that could provide access to those segments and determine the right mechanism by which to grow—for example, will it be through partnership or an acquisition?

Companies may perform the first step routinely as part of their strategic-planning processes; however, they are less likely to tackle the second step with the diligence required to understand which new market segments a target or partner company can help usher them into. It is relatively straightforward to understand the products and services a target can offer you as a customer. But what capabilities are behind those offerings? What patents and technologies does the supplier hold that can generate growth for your company in adjacent markets? What sort of competitive advantage does the supplier hold in markets that might be of interest to your company? The research and conversations associated with the first two steps can help nudge semiconductor executives outside their comfort zones, while the third step will likely raise larger strategic and resource questions as companies begin to view their value-chain partners in a very different light.

Identify complementary market segments

Semiconductor companies can kick off the search by casting a wide net, as they would at the beginning of any strategic initiative. This means taking an exhaustive inventory of all the technologies and capabilities that their suppliers provide and considering how those assets could be applied in other market segments. For instance, an equipment manufacturer may discover that the technologies it uses in its wafer-inspection machinery can also be applied in the nondestructive testing schemes deployed in the aerospace or automotive industries. Or a company may recognize that the technologies it uses to produce the ingots that are at the core of its integrated circuits can also be used to produce solar-photovoltaic cells. The company's inventorying process could generate a healthy list of potential growth areas. But before making a move, semiconductor players should also consider how similar the potential target-market segments are to their current businesses. For example, are product-development processes and standards in the photovoltaic-cell market similar to those used in standard IC development? If so, the company's likelihood of success in that market will be higher because of the company's

Exhibit 1

Some companies are finding growth opportunities in the supply chain.

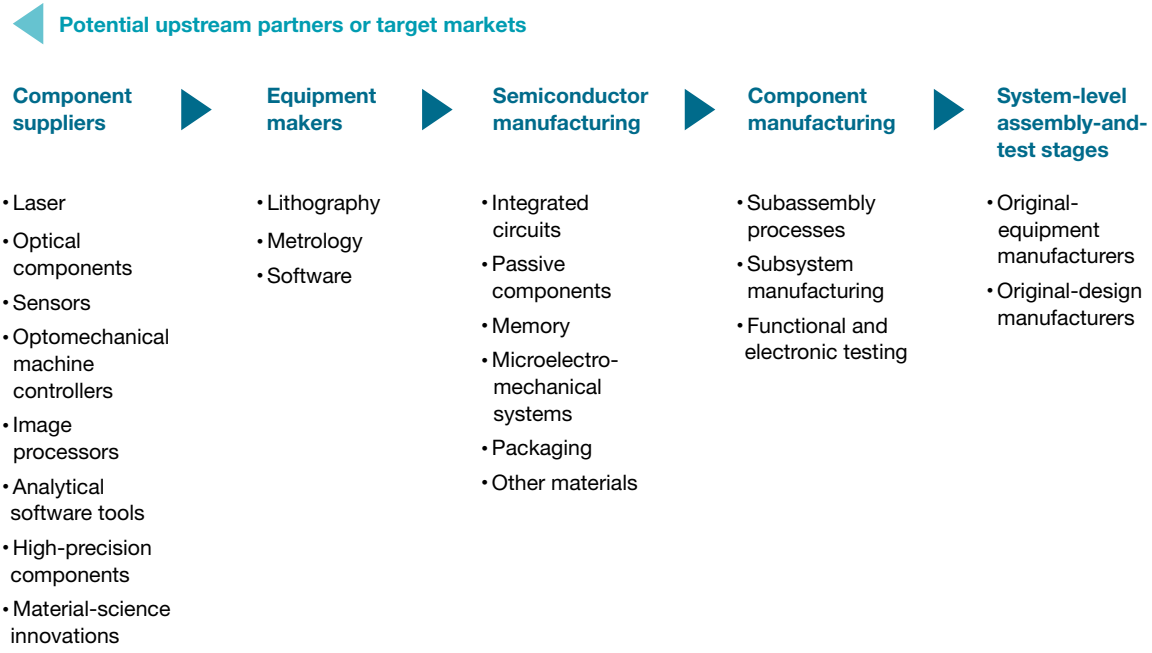

familiarity with critical aspects of the new market. The inventorying process can help the company focus on potential growth areas that are beyond its core but do not stray too far from its existing competencies and experience.

Find the right supplier

Once it has a market segment in mind, the semiconductor player must then determine which supplier, or set of suppliers, can provide a differentiating platform or infrastructure. There are a number of factors the semiconductor player should consider when determining who to target or partner with—among them, the company's level of spending with the supplier, the supplier's competitive position and revenue growth, and the compatibility and strength of the supplier's intellectual-property (IP) portfolio (for instance, the number of patents it holds inside and outside of the target market segment) compared with the manufacturer's own IP (see sidebar, "Using recursive- growth analysis to move beyond core markets").

A look at the supplier's business model and sales channels could also reveal potential alignment. Let's consider a semiconductor company whose business model is centered on the initial sale of manufacturing equipment. In its scan of suppliers, the company sees several potential partners that are focused on the long-term servicing and sale of replacement components. Partnering with or acquiring these suppliers would require the semiconductor company to develop new capabilities in its sales force and supply chain—so these players would likely not be an immediate priority (unless the semiconductor player is ready to make large investments in

sales and logistics), and the choices and opportunities could be narrowed further.

A chief consideration when picking a supplier as a partner (or target) should be minimizing the changes required to succeed in the new segment. Semiconductor executives must also consider the regulatory environments of the market segments they are targeting and their potential risk exposure—the costs of noncompliance could be very different depending on prevailing industry mandates. For instance, a company's proposed use of a product or technology in a new healthcare segment might not just benefit from but actually require partnership with a core supplier that understands the US Food and Drug Administration approvals that might be needed to bring the idea to market.

To perform a comprehensive assessment of the opportunities, the company will need input from across all functions, including business-unit representatives and executives in finance, IT, procurement, and strategy.

Determine the most effective mechanism for growth

Once the company has identified a priority supplier, or set of suppliers, it needs to decide how it will gain access to that player's technology or platform—that is, will it be through acquisition or partnership? There are a range of factors to consider. An acquisition may be a good choice if the semiconductor player can immediately capitalize on the technology acquired from a target company, and in cases where exclusivity

Exhibit 2

A number of companies have made acquisitions in the NAND flash-memory market.

Acquirer	Target	Date	Value	Details
SanDisk	**Fusion-io**	July 2014	$1.1 billion	Important elements included flash-storage systems, enterprise solid-state drives (SSDs), and flash software
SK Hynix	**Softeq**	June 2014	N/A	SK Hynix acquired the NAND firmware arm of Softeq to strengthen its flash-controller solution
Toshiba	**OCZ Storage Solutions**	Jan 2014	$35 million	Toshiba acquired OCZ assets (including intellectual property relating to the Indilinx controller) after OCZ filed for bankruptcy
SanDisk	**SMART Storage Systems**	Aug 2013	$307 million	Acquisition included serial ATA and serial attached SCSI enterprise SSDs
IBM	**Texas Memory Systems**	Oct 2012	N/A	Deal involved enterprise SSDs and multilevel-cell flash technology

continued on page 78

Using recursive-growth analysis to move beyond core markets

Advanced analytics and big data can be powerful tools for identifying growth opportunities beyond the core. Using them can make it easier for semiconductor players to identify potential acquisition targets by assessing information relating to company patents, revenues, products, and other critical identifiers. A thorough analysis of public and proprietary patent databases, for instance, can provide insights not only into the quantity and quality of patents owned by a particular company but also about the shifts in the patent landscape over time. Semiconductor players may be able to evaluate how a potential target's patent portfolio stacks up against other players' patent collections, the respective players' areas of focus, and the applicability of their intellectual property across various market segments.

One particular methodology developed by McKinsey, recursive-growth analysis, takes the user through several "degrees of separation" in markets and products to identify less-obvious expansion possibilities—ones that are far from the company's core segments but that still optimize and build upon core capabilities. This methodology draws on a database of growth activities pursued by more than 200,000 companies across 2,000 industries worldwide, as well as some 600,000 financial records that have been mined to estimate and visualize growth and profitability in particular areas. Using recursive-growth analysis, the semiconductor player identifies peer and target companies or business units and an initial set of markets in which peers may be active but the company is not. The process is repeated, allowing the semiconductor player to rank and further explore opportunities uncovered during each pass. The exhibit provides an example of just such an exercise performed by a company with expertise in semiconductor-testing equipment looking to move beyond its core market.

Exhibit

Recursive-growth analysis can identify opportunities beyond core markets.

The recursive-growth tool generates a wheel of activities and market segments radiating from a company's core capabilities (darker segments) in the center to unexplored opportunities at the edges (lighter segments).

and control of IP are important to build or maintain competitive advantage; the semiconductor player would gain immediate ownership of the supplier's technologies, sales channels, and personnel.

By contrast, partnering may be the right choice if there would be antitrust issues associated with an acquisition, if the cost of acquisition would be too high, or if only a portion of the supplier's business is attractive to the semiconductor company but the supplier is unwilling to carve it out from its core operations. Partnering rather than acquiring may be also be desirable if the venture would be risky—for instance, where the market is still nascent or where the technology is not proven.

Let's consider two examples that highlight the different mechanisms:

ASML acquires Cymer. Dutch-based ASML, over the past ten years, has emerged as a leading provider in the rapidly evolving market for photolithography systems and equipment. Among its competitors, only Nikon has retained a market share above 10 percent over that period (Exhibit 3). Photolithography tools typically cost upward of $40 million per unit. Next-generation extreme-ultraviolet-light tools are expected to cost between $80 million and $120 million per unit. Only a few device and equipment manufacturers will be able to sustain these high product-development costs long term.

Exhibit 3

ASML has increased its share of the market for photolithography systems and equipment.

Core players in wafer-fabrication lithography equipment,
% of market[1]

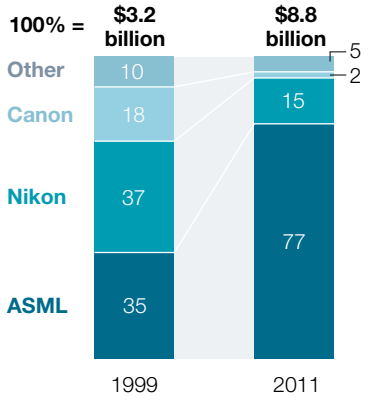

[1]Figures may not sum to 100%, because of rounding.

Source: Strategy Analytics, 2012

Exhibit 4

Industry investments ensure that innovation continues in extreme-ultraviolet-light technology.

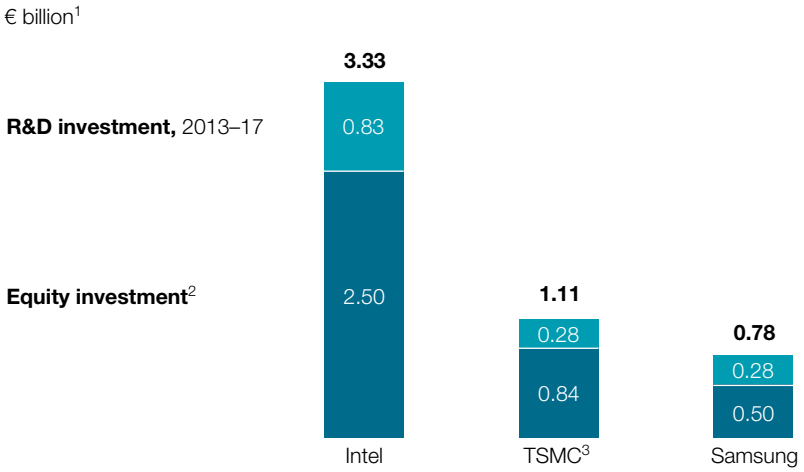

[1]Figures have been rounded up.
[2]One-time payment.
[3]Taiwan Semiconductor Manufacturing Company.

Source: *The McClean Report*, IC Insights, 2014, icinsights.com; *Solid State Technology*, 2014

In this environment, ASML targeted laser maker Cymer in a cash-and-stock deal valued at about $4 billion. In announcing the news, ASML suggested the deal would help to accelerate the development and commercialization of extreme-ultraviolet-light sources, a critical technology for enabling the continued downscaling of transistor-node size.[2] Cymer and Japan-based Gigaphoton shared the market for deep-ultraviolet-excimer lasers for photolithography, and both had been developing extreme-ultraviolet sources for several years using laser-produced plasma. But, according to company officials, the light-source technology that Cymer owned was central to ASML's growth plans. The company also gained access to a number of potential new partners, namely Cymer businesses that supply optics to LCD and organic-light-emitting-diode manufacturers.

Manufacturers partner with ASML. In July 2012 ASML announced a customer coinvestment program to enable minority investments in ASML to support and accelerate the company's research and development of new technologies for extreme-ultraviolet lithography and the fabrication of 450-millimeter wafers. Intel was the first to sign on, and Samsung and TSMC joined in August 2012, with their combined investments totaling €5.2 billion (Exhibit 4). The partnerships reflect a realization among industry players that productivity improvements and growth in semiconductor manufacturing have traditionally

come from increasing the diameter of silicon wafers—and, more recently, from the accelerating change in enhanced lithography technologies such as immersion lithography, deep-ultraviolet-light sources, and the extreme-ultraviolet-light sources described earlier. But there are high development costs associated with these still-nascent technologies. When announcing the program, ASML noted that the collaboration spreads financial, R&D, and implementation risks among multiple parties and enhances the company's ability to improve shareholder and customer value.

• • •

The environment for technology companies has been difficult the past few years. The industry has always been cyclical, but it is absolutely possible that the current slow-growth environment is now the new normal. The companies that are best able to ferret out the opportunities in their supply chains and find seemingly elusive pockets of growth will have the advantage.

[1] Ashlee Vance, "With McAfee deal, Intel looks for edge," *New York Times*, August 19, 2010, nytimes.com.
[2] Roberta Cowan, "Chip gear maker ASML buys Cymer for $2.5 billion," *Reuters*, October 17, 2012, reuters.com.

The authors would like to thank JehanZeb Noor for his contributions to this article.

O Sung Kwon (O_Sung_Kwon@McKinsey.com) is an associate principal in McKinsey's Southern California office, **Mark Patel** (Mark_Patel@McKinsey.com) is a principal in the San Francisco office, and **Nicholas Sergeant** (Nicholas_Sergeant@McKinsey.com) is a consultant in the Silicon Valley office.

November 2014
Designed by Global Editorial Services

Printed in the United States of America.