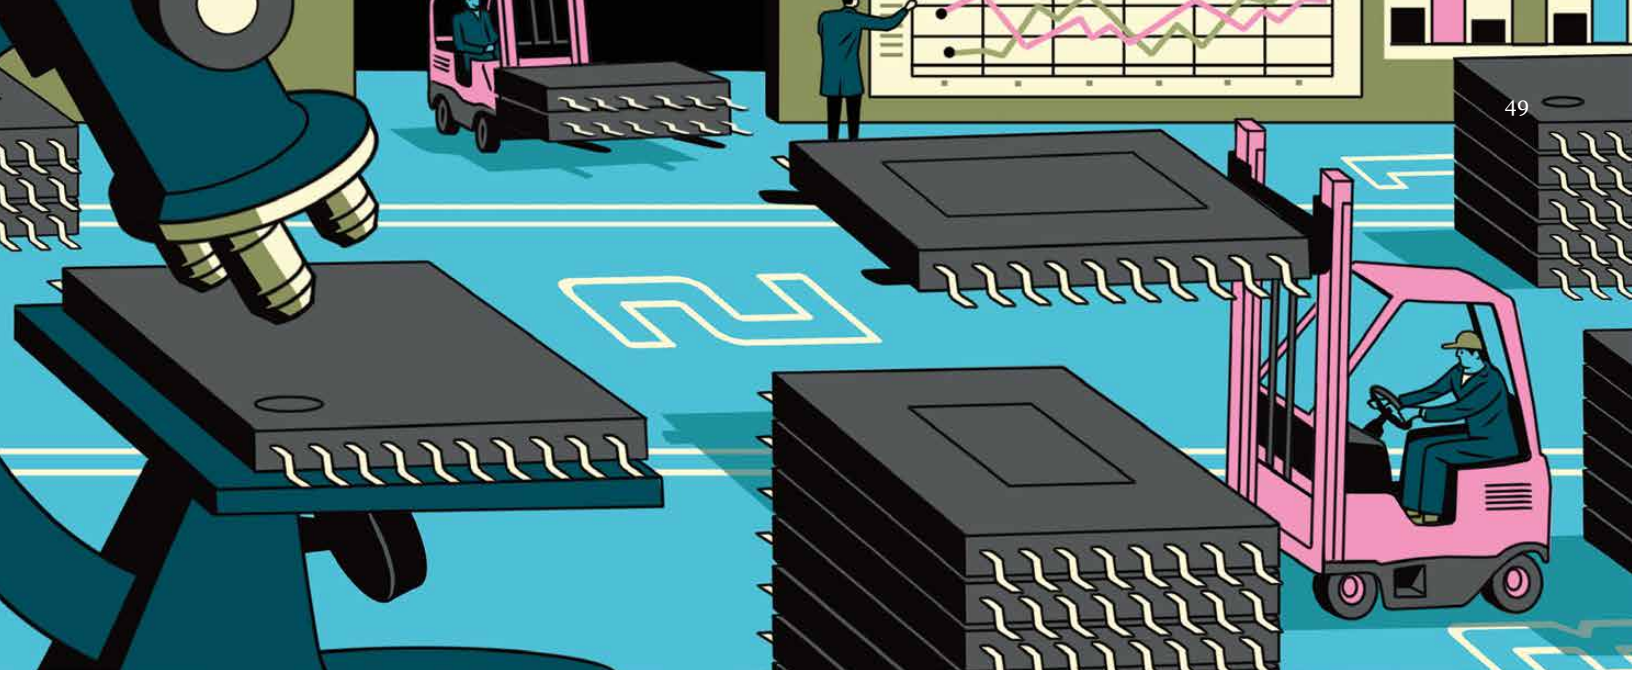

Bill Butcher

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.

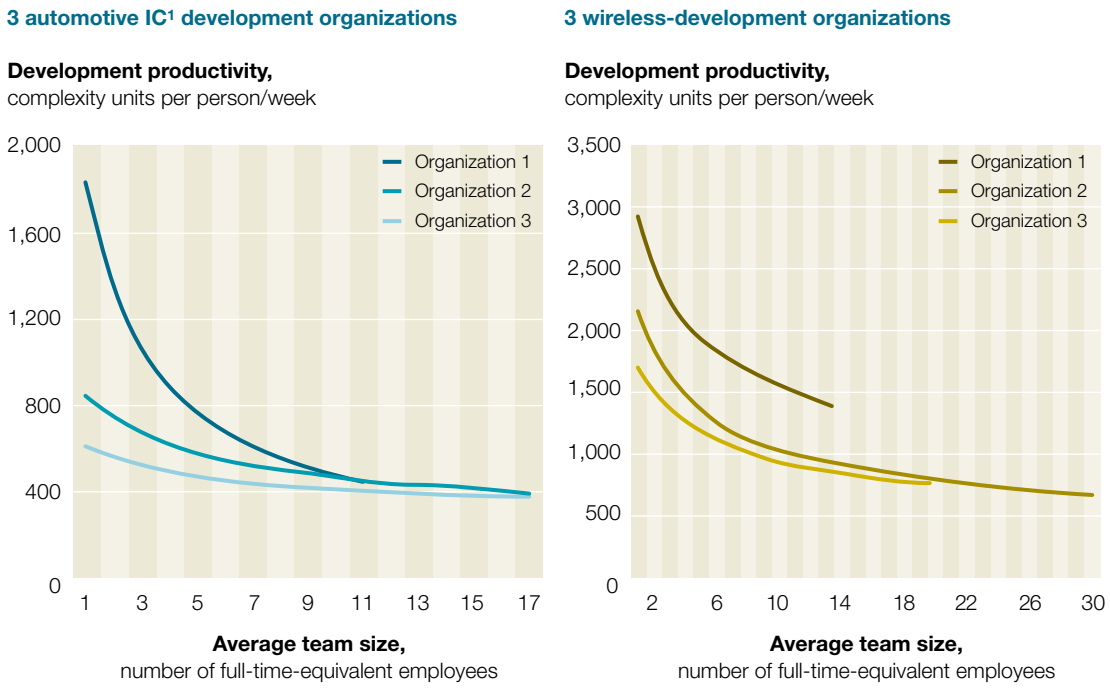

[1] Integrated circuit.

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.

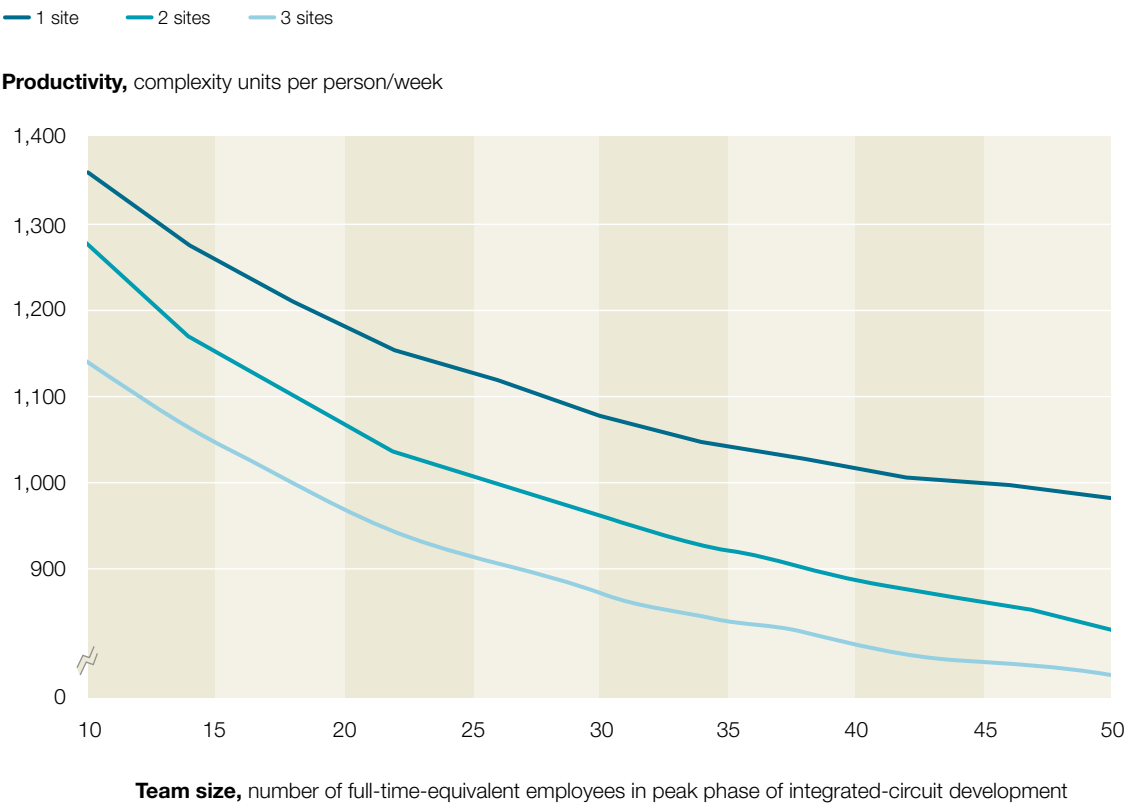

of the cycle, which our research suggests can have an enormous effect on time to market. Semiconductor players may be missing out on opportunities to cut weeks, or even months, from predevelopment phases of production. Based on our research on more than 2,000 integrated-circuit projects at more than 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.

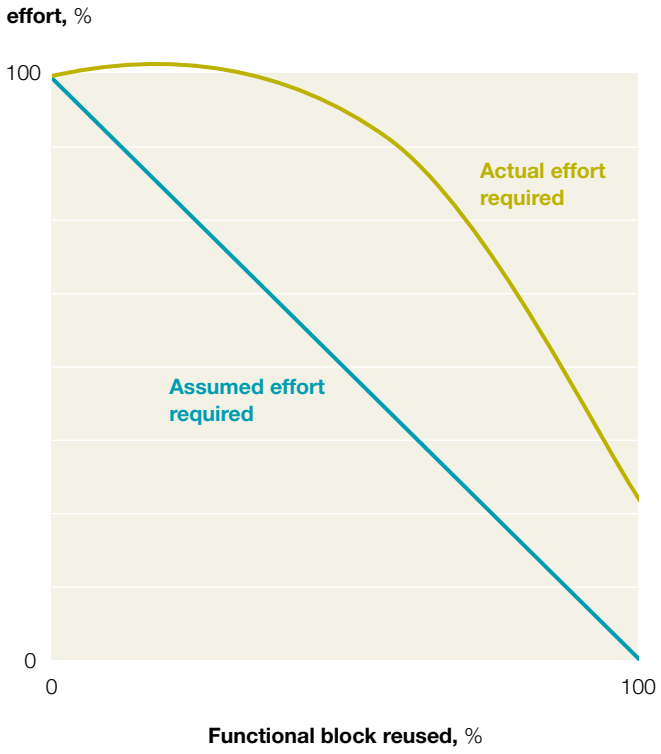

Exhibit 4

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

Time spent on phase, weeks

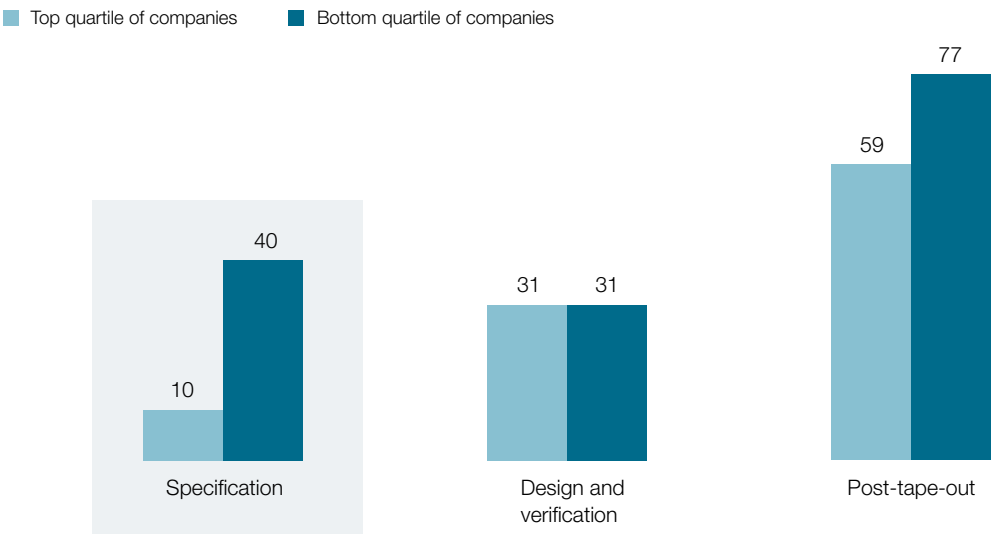

sharpen its front-end development capabilities, improve its time to market on most projects, and regain its foothold in a competitive market.

• • •

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.