

Semiconductor Magazine

Building a Fab – It's All About Tradeoffs

Balancing enormous financial risk with cyclical market demands is like a no-limit poker game

by Katherine Derbyshire

June 2002

A new fab represents an enormous investment of time, money and corporate resources. The success or failure of an entire company can depend on that fab's return on the investment. Maximizing ROI requires a careful balance of many different factors, from the initial fab design to the production ramp phase.

In general, fab planning and construction proceeds from general questions to specific details. To begin, designers must know how many wafers per month the fab will process, what size wafers those will be, and what design rules and process technology will be used. Unfortunately, the optimal combination of capacity, wafer size and process technology depends on the market conditions when the fab begins production – often long after the planning effort has occurred.

WHEN TO BUILD

It takes more than two years to build a fab and ramp production. However, technology transitions happen approximately every 18 months, and semiconductor sales peak every 24-36 months. Neither technology roadmaps nor market forecasts are reliable two years into the future. Yet bringing up production of the wrong technology in the middle of a downturn can have disastrous financial consequences.

The DRAM industry is famous for capacity planning problems. Typically a period of stable or rising DRAM prices and strong demand inspires a burst of new fab construction. As the first new fabs come on line, they drive supplies up and prices down. Fabs that open later in the cycle face stiffer price competition. In 1996-97 and again in 2001, memory selling prices fell below manufacturing costs. Negative returns from memory manufacturing plants contributed to the Asian financial crisis in 1997-98, and the aftereffects of the 2000 building boom have driven Hynix to the brink of bankruptcy and forced smaller suppliers to consolidate.

These enormous financial risks argue for conservative capital investment plans. Yet companies that maintain aggressive capacity investments are able to deliver product to the market more quickly as demand rises. Commodity DRAM suppliers with more capacity are able to exploit market price fluctuations. Capacity helps suppliers of high value-added chips like microprocessors to realize the price premiums that new chips often enjoy.

Conversely, a chip supplier that is unable to meet demand is likely to lose both current and future orders. Reduced orders will in turn limit future capital spending, constraining the company's growth. Available capacity can make the difference between meeting and missing a market window.

Intel owes its continued success in part to its ability to build capacity year in and year out, through good and bad market conditions. By investing even during market downturns, the company almost always has leading-edge capacity ready when market opportunities appear.

Advancing technology makes capacity planning even more complicated. More aggressive design rules deliver higher performance in a smaller package, and usually at a lower unit cost. Higher performance earns a price premium. The first 18 months of a leading-edge fab's life are usually the most profitable and can repay the initial investment very quickly. Fabs using older technology are unable to command this premium and in fact may be at a cost disadvantage compared with more advanced fabs.

Yet technology transitions carry substantial risk as well. New manufacturing tools may not meet their planned delivery dates, may not be ready for the rigors of volume production, or may not deliver the expected process capability. Technology delays can force the fab to miss critical market windows. The closer a company is to the leading edge, the greater the technology risk. For many companies, alliances and joint ventures help mitigate both business and technology risks.

WHAT TO BUILD

By balancing expected market opportunities against business and technology risks, fab planners arrive at the fab conceptual plan. This plan estimates the fab capacity, which in turn determines the number of tools, the amount of wafer storage space, and the size of the cleanroom. These decisions allow the fab owner to develop a budget and begin the design and construction of the fab shell.

Wafer size has a significant impact on both the building floor space and the eventual capacity. According to Peter Hillen, president of U.S. Operations and Worldwide Marketing and Sales at Dongbu Electronics, a planned capacity of 30,000 300 mm wafer starts per month is equivalent to at least 45,000 200 mm wafer starts per month. The 300 mm fab can support much larger product demand and in fact requires higher demand to give an acceptable return on investment.

Fabs using 300 mm wafers require much larger initial investments, particularly in relatively immature automation and wafer tracking systems. While theoretically the larger wafers reduce per die cost, it is not yet clear when those savings will actually materialize for all products. Foundries may find that small lots are difficult to manufacture economically with the larger wafer size.

HOW TO BUILD

After deciding what kind of fab to build, the company must then decide where to put it. In some ways, a fab shell is just another building. In other ways, fabs are unique even among industrial manufacturing plants.

The shell

For example, lithography equipment requires careful analysis and control of ambient vibration. Highways, airports and rail links can all either render the site unsuitable or require more complex seismic isolation. At the same time, both the construction team and the fab employees need highways and other transportation infrastructure, and later ongoing fab operations have their requirements, as well. The seismic analysis must consider conditions created by the fab itself. According to Bart Rogers, project manager for DPR Construction, seismic isolation is often the most time-consuming and expensive part of the construction phase.

Similarly, the fab's power and water consumption place substantial pressure on local supplies. The availability of stable power and abundant water will determine what on-site water purification and power generation facilities are needed. In extreme cases, Rogers said, fabs have been cancelled or relocated because of problems with the local utility infrastructure.

The fab's location also will define the available construction labor pool. Most construction projects are inherently dirty. Unless construction workers have built fabs before, they may not be familiar with clean construction protocols. Final flooring, painting and even landscaping happen earlier in a fab construction project than in other kinds of construction. Materials need to be cleaned before installation and kept

clean during construction. According to Rogers, DPR begins to implement clean protocols from the very beginning. As the fab shell grows and cleanliness becomes more critical, workers become more familiar with the procedures.

Designing the process

Because of the technology and market risks discussed above, designers try to defer process decisions as long as possible. As the fab shell grows, these decisions become more urgent. Lead times for advanced lithography tools in particular can reach a full year during strong growth periods, while shell construction can take as little as six months.

Other processes may have unique facilities requirements. For example, many fabs segregate copper and CMP areas from cleaner processes, affecting wall placement, ventilation and other cleanroom infrastructure. Though leading-edge microprocessors rely on copper interconnects, copper is not yet the metal of choice for all applications. In many cases, aluminum still offers lower cost, better yield and adequate performance.

The layout of the fab as a whole depends on tradeoffs between tool productivity and cycle time. Placing all wet benches or all etch systems close together tends to improve load balancing and tool productivity, says Ron Niv, manager of Tower Semiconductor's Fab 2. This clustering tends to increase cycle time, however, since wafers have to travel long distances between process steps. Placing tools in order according to the process sequence improves cycle time for fabs with a small number of high-volume products, but not necessarily for fabs with many low-volume process variations. Layout decisions in turn control the facility's infrastructure, particularly the placement and design of equipment hookups.

Installing equipment

According to Kirby Hicks, Applied Materials' general manager of Ramp Performance Management, typical fab projects draw a sharp distinction between the construction phase and the tool installation phase. The most cost-effective construction approach is to lay out and install the utility lines in groups, working from one part of the building to the next, and to complete the utility installation before beginning tool hookups. Yet facilities construction is a relatively small fraction of the total fab cost, while equipment and installation represent 70-80 percent of the total. The owner's investment jumps sharply as soon as equipment arrives. To maximize ROI, the owner would like to reduce the time between tool installation and first production.

At the same time, Hicks explained, the fab need not have its full complement of equipment installed in order to begin process development. Often, installed equipment may sit unused for weeks while the process engineers focus their attention elsewhere.

For the project, it may be more cost-effective to gradually bring equipment in and install it as needed over a period of several weeks rather than asking suppliers to meet a more aggressive installation schedule. It may be better to complete tool locations on the critical path to production first, then go back and finish the remaining tool hookups. Such "scattershot" methods increase the cost of the construction phase, but may reduce tool ramp time and allow more cost-efficient timing of tool purchases.

Delayed permits, incomplete tool hookups and similar problems can threaten the schedule and budget of the entire project. For example, an installation team that arrives to find that the fab is not ready for its equipment may either sit around with nothing to do – while billing for time – or return home until the problem is corrected. By balancing the tool installation workload, Hicks said, the fab can avoid these mistakes in the first place and minimize the impact when they do occur. The fab could confirm that the installation is on track before the team ever leaves its home base. Gradual tool installations allow the supplier to keep a smaller installation team in place for a longer period.

PRODUCTION RAMP

The fab can begin process qualification as soon as enough tools are installed to make test chips. Many fabs use SRAM circuits as a test vehicle. However, the SRAM, a standardized array of memory cells, may not be a good model for the random distribution of features typical of a microprocessor or system-on-chip (SOC) circuit. Many potential process problems are layout-dependent. Dishing and erosion in chemical mechanical planarization (CMP), lithography's optical proximity effects, etch loading and many other process behaviors vary with feature density. Via processes depend on the behavior of underlying layers, as well as on the layout – stacked or unstacked – of the vias themselves.

Instead of an SRAM test vehicle, PDF Solutions uses a set of short loop test modules to isolate the different areas of the circuit. A single wafer might include test vehicles for stacked and unstacked vias, dense and sparse arrays, and so forth. The fab can begin processing test vehicles as soon as it can complete the appropriate segment of the total process.

As the process climbs toward production capability, the test vehicle indicates which failures are more common and where the fab should focus its process improvement efforts. This information is more useful in conjunction with the particular designs being manufactured. Dave Joseph, PDF's executive vice president of sales and marketing, explained that one client company experienced 500 failures per billion stacked vias, but only 15 failures per billion unstacked vias. Unstacked vias had a larger yield impact, however, because the fab was manufacturing SOC designs with very few stacked vias.

The yield impact matrix also can suggest design modifications to improve manufacturability. For example, an SOC design with many small memory areas can't tolerate very many memory failures. Building redundant cells into each memory area might drastically increase the total circuit area. Instead, a designer might be able to improve memory yield more efficiently by relaxing the bit cell design rules. Conversely, a chip with a few large memory areas might be able to improve yield with less area penalty by adding redundant cells at a tighter cell design rule.

In fab ramp, as in all areas of fab construction, time is money. Applied Materials estimates that each additional 300 product wafers per week during the ramp phase represents \$675,000 in additional revenue for the fab. The faster a fab reaches full production, the more quickly it can begin to recover its investment, and the more profitable it is likely to be.

Fab construction has been described as a no-limit poker table, where players have to keep increasing their bets or be chased out of the game. Fast-track design, efficient construction and rapid production ramp can help even the odds.