



HARRIS CORPORATION - SEMICONDUCTOR SECTOR

Binu Koshy prepared this case under the supervision of Professor Peter Bell, Richard Ivey School of Business, the University of Western Ontario, solely to provide material for class discussion. The authors do not intend to illustrate either effective or ineffective handling of a managerial situation. The authors may have disguised certain names and other identifying information to protect confidentiality.

INFORMS prohibits any form of reproduction, storage or transmittal of this material without its written permission. This material is not covered under authorization from any reproduction rights organization. To order copies or request permission to reproduce materials, contact INFORMS, 901 Elkridge Landing Road, Suite 400, Linthicum, MD 21090-2909; e-mail informs@informs.org.

Copyright © 1998 by INFORMS.

In late 1988, Harris Corporation acquired the GE Solid State (GESS) semiconductor product lines and manufacturing facilities from General Electric. Prior to the purchase, Harris' semiconductor sector enjoyed a profitable business producing sophisticated niche products. However, shortly after the acquisition, on-time delivery emerged as a crisis issue. The sector had many delinquent orders and customers judged delivery performance inferior compared to other semiconductor vendors.

Throughout 1989, the percentage of ordered line items delivered within one day of the promised delivery date hovered around 75 percent. One survey indicated that 60 percent of Harris' customers wished to replace it as a vendor. Potential sales of \$100 million had been lost in 1989, and sales continued to decline in the following years. What was once a profitable division of the corporation reported a loss of \$75 million in fiscal year 1991. Getting its products delivered on time had become a critical issue for the corporation.

Jon Cornell, sector president pulled no punches: "The sector will not survive unless we solve our delivery problem."

The Company

Harris Corporation was an electronics and electronics systems company based in Melbourne, Florida, with annual sales approaching \$3.5 billion. The corporation operated in several "sectors" of different lines of business. Harris' semiconductor sector had annual sales of \$670 million, with a substantial portion of these sales in components supplied to defense contractors or aerospace companies who were prime government contractors for weapons or space exploration programs.

The acquisition of the GESS semiconductor product lines and manufacturing facilities roughly tripled the size of the semiconductor sector and substantially increased the proportion of its production in competitive, commercial product lines, such as automotive and telecommunications products. Harris' traditional military and space markets were reduced to just a small portion of its semiconductor activity, and Harris now had to provide competitive on-time delivery performance over a much greater product mix.

The newly expanded sector was selling more than 10,000 finished product lines, produced by more than 30 manufacturing facilities in the United States and Asia. The manufacturing databases, control systems, and planning systems in use at GESS and at Harris were different and had proved difficult to integrate. As a result, Harris' management was struggling to cope with data provided in multiple formats on multiple systems and, in some cases, containing serious gaps. Production planning and delivery quotations were decentralized and conducted using a myriad of systems, policies and personnel.

The sector utilized two large MRP systems, one for the former Harris facilities and the second for the former GESS facilities. These MRP systems, although well designed, had serious shortcomings for application to semiconductor manufacturing. In addition, many smaller MRP-like spreadsheet analyses were performed by factory planners. Data on demand, work in process, inventory, and capacity were unreliable and were judged differently by various participants. Working out a plan inevitably involved meetings to negotiate differences, leading to multiple planning iterations. Sector-wide planning cycles were undertaken only once a month and consumed two weeks or more. Quotations and delivery commitments were often little more than judgments made by planners who were forced to work with incomplete information.

Within two years of the acquisition of GESS, the sector had developed a reputation for late delivery.

Semiconductor Manufacturing

Semiconductor manufacturing could be thought of as a two-stage process:

1. *Wafer fab and wafer probe* involved fabricating integrated circuit structures on silicon wafers and testing the circuits, and was conducted in *front-end plants*,
2. *Device assembly and device test* performed in *back-end plants* where these wafers were sliced into individual devices, packaged, and then tested.

The manufacturing process flow for semiconductor products consisted of five serial processes:

- Front-end plants carried out base wafer fabrication (in wafer fab), wafer fabrication (wafer fab), and wafer electrical probe (probe),
- Back-end plants then followed with device assembly (assembly), and device test (brand, test and pack).

Harris had more than thirty manufacturing sites around the world. Each site, or even a single manufacturing facility at the site, typically operated several or all of these processes. Generally, one site integrated the first three processes (the front-end), and a different site, the last two processes (the back-end).

Front-end Plants

Front-end plants consisted of the first three stages in semiconductor manufacturing. During these three stages, round wafers of silicon were processed through wafer fab and wafer probe in small lots of 10 to 50 wafers.

The fab processes (stages 1 and 2) imprinted the wafers with many identical patterns of an integrated circuit. This circuit pattern on each wafer was called a *die*. The probe process (stage 3) electrically tested each die on each completed wafer to identify which of the individual circuits on the wafer were functional. Wafer fab processes had hundreds of serial operations; wafer probe processes generally involved only a few steps.

The subdivision of the overall wafer fabrication process was made to allow for the maintenance of an inventory of semi-processed wafers known as *wafer bank*. The inventory of wafers that had completed the third stage was known as a *die bank*.

Back-end Plants

Back-end plants consisted of the last two stages: assembly, and brand, test, and pack. In the assembly stage, the wafers were sliced up into individual chips which were sealed in plastic or ceramic packages to become *packaged devices*. In the final stage, the packaged devices were tested, labelled, and packed for customer shipment. The assembly and test manufacturing areas processed individual devices in manufacturing lots of 500 to 2000 devices.

One important difference between production of semiconductors and other manufactured/assembled products was that, for semiconductors, one intermediate product was the source for several final products, while in other manufacturing, several intermediate products converged to make up one individual product (for example, an automobile assembly plant). More specifically, a single base wafer was the source product for several types of intermediate wafers; a single die type was the source product for

several types of packaged devices; and a single packaged device was the source product for several finished good types (Figure 1).

Manufacturing of Binned Products

Many semiconductor product families included various quality grades and design revisions; consequently, alternative or substitutable source products were very common. In particular, many intermediate products were the result of *binning*, whereby several quality-graded products emerged from testing a single manufacturing lot of source product. For products featuring binning, the test process was split into two parts, the *initial test* and the *final test*.

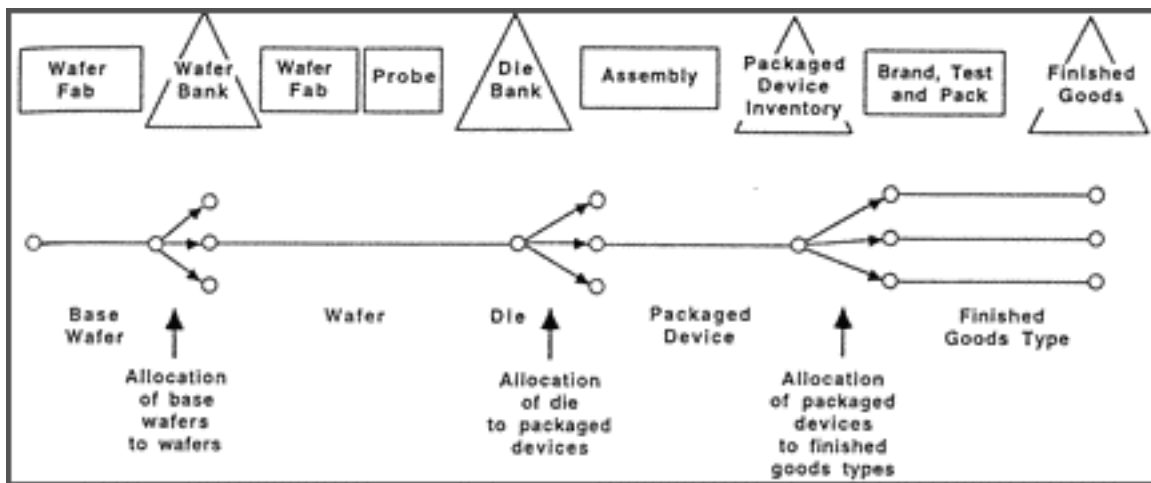


Figure 1: Schematic of Semiconductor Manufacturing Line

The initial test process subjected packaged devices to one or a series of electrical tests to measure the various electrical attributes of interest (i.e., speed of microprocessors, power consumption, resistivity, etc.). Generally, several combinations of attributes emerged from each manufactured lot, each termed a *bin*, where each bin defined a specific range of performance on one or several electrical attributes of performance. For example, a bin definition might be “speed between 20 and 30 megahertz and power consumption less than 100 milliamps.” The inventory of binned packaged devices following completion of the initial test process was known as a *class store*.

The fractions of the source product falling into each *bin* on initial testing were known as the *bin splits*. The bin splits were characteristic of the manufacturing process and were regarded as prespecified for planning purposes.

The final test process was part of “brand, retest, and pack”. When the final test process included a *burn-in* operation (prolonged operation of the devices in an oven), binning tests were often repeated during final testing to determine whether attributes had changed since the initial test. If they had, devices might have to be downgraded (assigned to a lower grade bin) and returned to class stores. However, such a fallout was typically very small

compared to the fallout in the initial test and, for planning purposes, it could be regarded simply as yield loss.

Binning also occurred in wafer probe, where a finished wafer could contain multiple useable electrical grades of die. Analogous to the allocation problem after initial test, more than one graded die type could be suitable as the source product for a particular packaged device type. Moreover, design revisions often resulted in die from several wafer types being suitable for the same packaged device. Figure 2 illustrates the two types of bin splits.

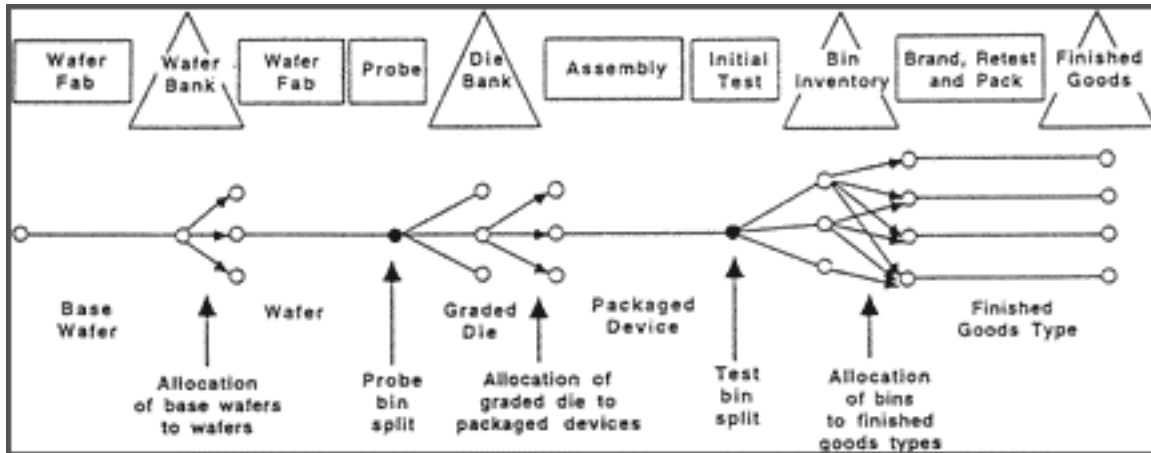


Figure 2: Schematic of Semiconductor Manufacturing showing Bin Splits

The product structure for binned products could be used to model nonbinning products, with appropriate conventions. For nonbinning wafer-probe processes, one graded die type could be thought of as being generated with a 100 percent bin split; and the initial test process flow for such packaged devices included no operations, had 100 percent yield, and 0 days manufacturing cycle time. Using this convention, Harris' manufacturing involved four standard corporate inventory points (wafer bank, die bank, class stores, and finished goods), and six standard manufacturing process flows (base wafer fab, wafer fab, probe, assembly, initial test, and final test). Two of the process flows (base wafer fab and initial test) could be vacant for some product families.

Harris had a number of front-end and back-end sites. It could produce certain die types at more than one front-end site; much less commonly, it could produce certain finished goods types at more than one back-end site (at most two). Typically, it produced low-volume die types at only one front-end fab site, but often produced high volume die types at two or more fab sites, attempting to distribute the manufacturing volume to use capacity efficiently.

The Semiconductor Planning Problem

Several characteristics of semiconductor manufacturing contributed to make planning production quite challenging, and MRP logic difficult to apply.

- In semiconductor manufacturing, a single source product produced several intermediate and final products. This process flow was the exact opposite from other assembly plants where MRP logic was used.
- Semiconductor factories were capital-intensive facilities that operated 24 hours per day, seven days per week. This meant that capacity was determined by the capabilities and performance of the processing equipment in the plant. With the potential market frequently exceeding the factories' abilities, Harris had to limit the market demand it accepted to fit its capacity. Planning methodologies such as MRP had difficulty developing efficient, capacity-feasible factory schedules for manufacturing planning.
- The loading of equipment capacity by semiconductor products was unusually complex. Products were routed through hundreds of steps, requiring weeks to traverse. The routings comprised reentrant process flows in which a product visited a particular equipment type many times for performance of different processing steps, interspersed with steps performed on other types of equipment. For example, wafers following the wafer fab process flow for a 20-layer circuit design had to visit the photolithography workstation 20 times. Fab cycle times ranged up to eight weeks and test cycle times for complex devices could range up to three weeks. This meant that newly released production lots competed for capacity with part-processed job lots, and product releases made over several weeks competed with each other for scarce capacity.
- Traditional forms of capacity analysis generally applied capacity constraints to total factory input or output rates in each planning period, without consideration of planned rates in adjacent periods. When the product mix was dynamic, as it was at Harris, these forms of analysis were inaccurate.

The binning and substitution possibilities in the product structure also frustrated the application of MRP logic to perform requirements planning. The various quality grades and designs made it difficult to establish parameters to enable MRP logic to generate feasible net requirements for factory releases.

Harris' semiconductor sector manufacturing systems department observed that other semiconductor companies tried to apply MRP logic to binning structures in two ways, both of which had problems.

One method was simply to ignore the bin splits and equate requirements for the source product to the sum of requirements for the finished goods. For example, in Figure 3, Bin 1, which corresponds to the electrical requirements of finished good type 1, was suitable for filling demands for either finished good type, while Bin 2 was suitable only for finished

good type 2. If period 1 demand was 25 for finished good type 1 and 80 for type 2, and period 2 demand was 10 for type 1 and 110 for type 2 then this method would set production for the source product at 105 in period 1 and 120 in period 2. If the actual bin split turn out to be 20 percent to Bin 1 and 80 percent to Bin 2, then actual finished good 1 output in period 1 would be 20 percent of 105 or 21, which would be a shortage of four units of finished good 1.

The second method was to select a single bin as the “driver bin” for the requirements planning calculations. Under this method, the planner would choose one of the types of finished goods and divide its demand by the split for the corresponding bin to determine the quantity of source product needed. For example, if Bin 1 was selected as the driver bin and the bin split was 20 and 80 percent (as above) then production of the source product would be computed as 125 (25 of type 1, 100 of type 2) in period 1 and 50 (10 of type 1, 40 of type 2) in period 2, leading to a shortage of 50 units of finished good 2 (20 extra units from period 1 + 40 units from period 2 – period 2 demand of 110 units = shortage of 50 units) in the second period. If bin 2 was the driver bin, production would be computed as 100 in period 1 and 138 in period 2, leading to a shortage of five units of finished good 1 in the first period.

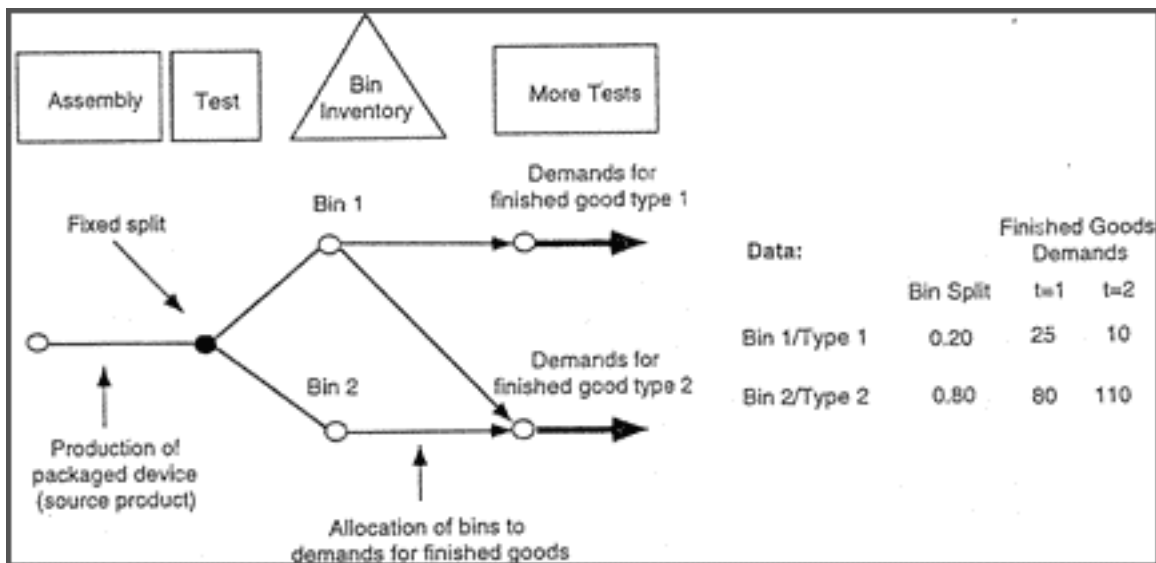


Figure 3: Example of Allocating Bins to Finished Goods

The sheer scale of the planning problem for Harris Corporation was daunting, since it appeared that to solve the delivery problem, Harris needed a planning system that would integrate all aspects of the semiconductor sector. To accomplish this it had to examine:

- How it would determine what orders to take given its limited capacity,
- How it could provide accurate quotations for customers' delivery times,
- How it could forecast demand for each product so that it could be manufactured as effectively and profitable as possible, and

- What factors constrained the production plan, and how the sector could maximize efficiency in these areas.

The on-time delivery crisis forced the sector to establish an extremely aggressive one-year schedule to redesign the sector planning system. Harris assembled a team from the University of California at Berkeley consisting of Robert Leachman, a faculty member at the Engineering Systems Research Center, Dale Raar, a masters graduate, and two doctoral students, Robert Benson and Chihwei Liu, to help with this task..

At the project kickoff, sector president, John Cornell, spoke to the entire sector on company-wide television. In his address he gave his strong support to the project by stating "...[this] is the most important project in the sector, and everyone must do whatever is necessary to expedite the project and to insure its success. (Harris semiconductor) sector will not survive unless we solve our delivery problem."