Static or Dynamic — The Selection Process for a Memory System

National Semiconductor
Memory Systems Division
J.C. Blackie†
November 1977

INTRODUCTION

Today's memory system designer is faced with a bewildering number of technology choices for his memory devices. He is offered overlapping performances from different technologies, different cell structures, and different packages. In addition, he is told of new developments daily that will obsolete all of the information he has already gathered. He is placed in the position where he must choose a technology, not on today's price, but on a projected price based on estimated minimum costs from vendors who are not yet in production with the devices they are offering.

There are many charts available showing the relationship between cost per bit, and learning experience, and the relationship between die size, and cost, on devices sharing the same learning experience; yet the driving reasons behind the costs are not always discussed. In addition, the device costs given do not reflect any additional expense that the user must incur in additional test costs. This additional testing can require significant capital expenditure and additional processing cost if the memory system is to achieve the reliability levels possible from semiconductor memories.

In an attempt to identify these costs more clearly, this paper will outline the decision processes and qualification testing to be followed by a design team choosing devices to be used in a line of standard memory systems that are to be competitive with existing core systems.

The systems were required to be single-card products utilizing power available in the user's system. They cover memory ranges from 16k x 18 to 128k x 18, in increments of 16k. The fact that the systems are to compete directly with well-established core modules, places a high emphasis on low cost and reliability. In addition, it dictates that static memory devices have to be seriously considered for potential applications where dynamic memories can not meet existing interface restrictions. In order to be cost effective with core systems, a volume selling price goal of 0.3¢/bit for the lowest density system, dropping to 0.18¢/bit at the highest density, was required. An analysis of the manufacturing costs of a wide range of semiconductor memory systems currently in production provided a set of non-storage cost constants. These constants define price goals for the memory devices. The systems analyzed had complexities ranging from cards with transparent error correction and power-down modes for battery back-up operation, to those offering high-storage densities with minimal overhead logic.

Figure 1 shows total non-storage costs in cents per bit of memory system versus memory system size. Although this data was compiled from dynamic memories, static figures were obtained by deducting the cost of all elements whose sole purpose is to support the need to refresh the system. Using the numbers shown in the table, one is able to establish maximum prices for memory devices that allow the design of semiconductor systems that are cost effective against core modules. Using a composite of predicted pricing from many vendors, the lowest cost per bit anticipated in mid-1978 for dynamic memory is 0.06¢/bit for both a 4k x 1 configuration and a 16k x 1 configuration. The lowest cost per bit anticipated for static devices is 0.15¢/bit for a 1k x 1 configuration or a 4k x 1 configuration. In each type, pricing per bit of the higher density configurations are projected to be on a significant downward trend while the lower densities are flat. Another fact that emerges is that static devices can be cost effective with core at low densities, but do not yield the lowest cost per bit solution in applications that can accept dynamic devices. Static devices greatest application appears to lie in areas where:

a. System performance is severely impacted by cycle stealing for refresh of dynamic memory.

b. Board space limitations do not allow space for additional logic required for dynamic interface

c. Memory requirements are small and can share a board with other system components.

As the system design goal is to replace core systems, there are many situations where the requirement to interrupt system usage for refresh cycles is impractical, thus there is a requirement for static versions of all modules.

Figure 1 shows that even with devices currently in production (1k x 1 static and 4k x 1 dynamic) semiconductor memory can be cost effective up to 32k x 18; the choice becomes marginal at 128k x 18. However, with the advent of the higher density devices in 1978 (16k x 1 dynamic), semiconductor systems establish a clear price advantage over the entire range (16k — 128k). Having established that cost effective devices are available, the designer is now faced with the task of choosing which device and which vendor should be selected. Since the goal is core replacement, the designer can restrict his choices to MOS devices that provide access times in the 150 ns to 200 ns range. These speeds allow him to provide systems to match or exceed typical core performance, even if a transparent ECC scheme is used to enhance reliability.

MEMORY SYSTEM CAPACITY	*TYPICAL NON-STORAGE COST TO BUILD/BIT	
	DYNAMIC	STATIC
16k x 18	0.046¢	0.032¢
32k x 18	0.04¢	0.027¢
64k x 18	0.039¢	0.027¢
128k x 18	0.039¢	**

*Includes PC board, peripheral circuitry and assembly

**Single card product not available

FIGURE 1

†Refer to Introduction. This paper presented at COMPCON 1977

Because it is uncertain when the higher density parts will be available at a lower cost per bit than established product, careful consideration should be given to a design approach that will allow a single board to use both high and low density devices. However, this approach should only be followed if an analysis of the devices and vendors shows that the compatible devices offered meet all of the performance and cost goals. The cost of designing a card for each type is not great enough to warrant compromising performance or reliability with a single card design. Thus, the designer should feel free to make an unbiased choice of device for high and low density products as he starts the device selection phase of his design.

DEVICE SELECTION

A survey of product available shows several choices of 4k x 1 dynamic configurations, an 8k x 1 dynamic device, and a couple of 16k x 1 dynamic configurations. In addition, there are several choices of 4k x 1 static configurations. (In order to reduce the choices to a more manageable number, the designer must compile a list of factors that will impact the potential cost of a device).

The first and major factor is pressure of competition. The availability of several qualified vendors committing to volume production ensures a long-lived product with an industry-wide commitment to maximize yields. Because volume production commitments from several vendors is dependent on large users choosing a particular configuration, a knowledge of who is committing to use the device types selected, and when, is significant in the selection of a device.

The second factor is die size. This is more difficult to evaluate because the smallest die sizes might be obtained with special processing technology that has not been optimized and therefore will take longer to reach an optimum yield. This would have a serious impact on a product that must be introduced into an established market place as soon as possible. It is necessary to weigh projected die cost by a process learning factor. (Learning experience on a product typically leads to a 3:1 improvement in yield over a one-year period of volume production). Thus, the *means* employed to reach a small die size is more important than the actual size. A small die achieved by unique chip design without resorting to especially tight masking tolerances or critical processing techniques certainly would give a vendor a definite cost advantage which would be available to a system designer. The variation in die size offered by the vendors who scored high on the other factors outlined in this paper was not significant enough to influence the potential die cost.

The third factor is the type of package that will be offered. Several considerations can create package limitations that may not be obvious from the initial evaluation samples which are usually supplied in a hermetic side-brazed ceramic package which, although expensive, places the fewest restrictions on the device's performance. The lowest cost product will be one that can be packaged in CERDIP and plastic. CERDIP can put a restraint on the die size, and plastic can restrict its operating speed. Thus, the choice of a high-performance part with a wide die may be a poor choice from a low-cost standpoint, because its die may not fit in CERDIP and a high performance device may suffer degraded performance in plastic.

The fourth factor really concerns the choice of vendor rather than a specific product or process, but it is a key consideration from a cost viewpoint. An evaluation of the component supplier's process control program and his level of quality and reliability activity is important. Absence of a meaningful program will affect the level of yield improvement that can be expected within a certain time frame, and will raise the cost of administering product through the system designer's facility. Troublesome vendors with marginal product create extra costs.

Application of the above factors to the devices, in our case, allowed the device selection to be narrowed down to a 16k x 1 RAM from 4 viable vendors, and a 4k x 1 static from 4 viable vendors. These 2 parts offered the lowest potential cost per bit through 1979. This choice did not take into account the high performance static product announced at the Solid State Conference in Philadelphia in February 1977 because that level of performance was not needed. In addition, the technologies outlined, although very significant in terms of superior performance at a lower bit cost, will not be cost effective in lower performance applications in the near future.

Having chosen the devices and vendors, the designer must develop a detailed specification for the parts he needs to procure. This leads him to the next phase of his device selection.

CHARACTERIZATION

The error-free performance of any memory device is a function of the stability of the memory element and its ability to function correctly in the presence of electrical noise and temperature variation. In order to ensure that a system design does not exceed the tolerance of the memory device, a designer must have these limitations defined so he can structure his design correctly. The purpose of the characterization phase of the system design is to define these optimum operating conditions.

Each potential vendor chosen during the device selection stage submitted a small quantity of his devices for initial evaluation. This evaluation merely allows one to gain a familiarity with the device and is used as a vehicle to set up a characterization program. Any potential problems that would eliminate further consideration are noted, but this sample cannot be considered as suitable for final characterization. This device characterization must be performed on a group of devices comprised of samples from each speed range offered by the vendor. Each vendor's composite group is split into 2 groups. One is scheduled for a 1000-hour life test and the other for board characterization. This characterization and life test evaluation is the first and most critical phase of any memory system design. It must establish, as much as possible, each vendor's distribution. This will allow the designer to optimize his system around the chip design. Observing parts from the entire distribution allows him to identify operating region limitations with temperature, patterns, and voltage levels without

regard to specific data sheets. Potential pattern problems can often be identified by investigating anomalies in the device operating regions that are far outside the range specified in the data sheet using fairly simple test patterns. Once identified, these same anomolies can often be seen to move within the specified operating range when subjected to more complex patterns.

During the earlier selection phase every vendor supplies some level of characterization and life test data. Unfortunately in the early stages of new devices, the vendor has very little to offer. The most meaningful data is generated by device users and is not available for review. This is especially true for life test evaluations where it is very difficult to measure the need for extended burn-in from vendor's data. Most vendors offer some level of burn-in as part of their standard process flow, but do not define the level of infant mortality that might remain in the product shipped. This criteria is very critical to systems planned for high reliability. Therefore, the system designer must structure a life test evaluation that monitors infant mortality as well as long-term shifts in operating parameters. Infant mortality problems in semiconductor devices are shipping-lot dependent, thus lots should be continuously sampled for evidence of infant mortality problems remaining after the vendor's standard burn-in. This problem can be avoided by providing 100% extended burn-in, but this is prohibitively expensive for a low-cost product. An incoming program that required a sample burn-in of all lots can effectively identify lots requiring extended burn-in. In addition, it gives valuable information on the control levels implemented by each vendor. So armed with the data collected from the characterization phase the system designer can proceed to the design of his peripheral logic, and the structuring of his system. He can weigh the impact of interface logic design against chip performance and arrive at the most cost effective and reliable memory system. His system test plan can be designed to effectively catch marginal devices that could create field problems, and achieve the highest levels of reliable operation.

CONCLUSIONS

The rapid development of the semiconductor memory market has made available a wide range of products, offering potential performance and cost advantages. However, the market place is limited and not all of the devices offered will reach their cost goals. Thus, the system designer's task is to first select a device that has wide acceptance, and has several viable vendors. Device performance must be obtained without resorting to state-of-the-art technologies. He must then accept the fact that to obtain the optimum cost device it is necessary to precisely specify the performance needed and he must assure himself that each vendor's device comfortably meets that performance. He must accept the fact that he will be required to perform a significant amount of device characterization so that his design can be optimized to meet the twin goals of low cost and high reliability. Finally, his system test plan must be structured to thoroughly exercise all discovered pattern sensitivities in the basic devices. The choice of static or dynamic is one of system interfacing and not cost, and the use of static devices does not eliminate the need for careful device evaluation.

With the advent of the 16k x 1 dynamic device, the semiconductor industry is moving toward a standard configuration which will remove a significant amount of the present confusion of choice. The industry is also starting to extend the life of MOS products with cost reductions obtained by optimizing and shrinking existing designs. This will further simplify the designer's choice, as cell designs and chip structures will remain the same. However, it appears that the user of memory devices will still be faced with a need to carefully evaluate the devices he intends to use and perform some device screening if semiconductor memory systems are to reach the highest levels of reliability.