"Digital Design Obsolescence"

John O'Boyle
QP Semiconductor

Introduction

Many devices still in use today for military specification and high reliability applications were designed and fabricated in the 1970s and early 1980s. Integrated Circuits (ICs) engineered over thirty years ago continue to be specified in new applications due to their reliability in harsh environments and their high performance. Many industrial and military applications are adequately served with devices based on 3-micron bipolar process technologies, for example, down hole instrumentation for oil well drilling. Why? Because these devices have been shown to be rugged and reliable in mechanically and thermally difficult environments such as oil well drilling.

Unfortunately, the same environmental ruggedness has been difficult to achieve in more advanced (state-of-the-art) ICs developed and marketed recently. Therefore, industrial, aerospace and military manufacturers have relied on mature technologies for reliable devices for many applications. Of course, the performance capabilities of today's modern ICs are not overlooked by design engineers. And today we find numerous solutions incorporating relatively new DSPs, microcontrollers, microprocessors and high-density memories, but we also still find the "old reliable" devices incorporated for interface and driver needs. Additionally, when repairing an older system, or deploying additional older designed systems, it is easier to locate and employ an obsolete device rather than having to re-qualify a totally new board design.

Naturally, this would be a good combination of IC solutions – newer devices for advanced processing/decision making and older devices for existing interface and line driver/controller types of applications. In many cases it is an ideal solution except the quantities of the older devices are diminishing. And this does not even begin to delve into the problems with the obsolescence of analog parts where even entire process technologies are being eliminated with no replacement alternatives today.

The problem illuminated by this evolving situation is not new. Manufacturers have faced the obsolescence of needed hardware and componentry since the beginning of manufacturing time. Usually it was possible to find an artisan or skilled manufacturer that could make a duplicate of the original. But today that is less of an option due to the complexity of the ICs and, more importantly, the cost of running a manufacturing plant to fabricate these devices. Older integrated circuits are no longer profitable to manufacture and thus are rapidly becoming obsolete.

In some cases the obsolete parts just vanish without any warning, see Figure 1. Take for example National Semiconductor's dual RS-422 Line Driver; on April 21st, 2005 National unilaterally discontinued supply of these devices because demand was higher than anticipated and they ran out of wafer stock before they could announce a last time buy.

Figure 1
National Semi Announces Immediate Device Discontinuation

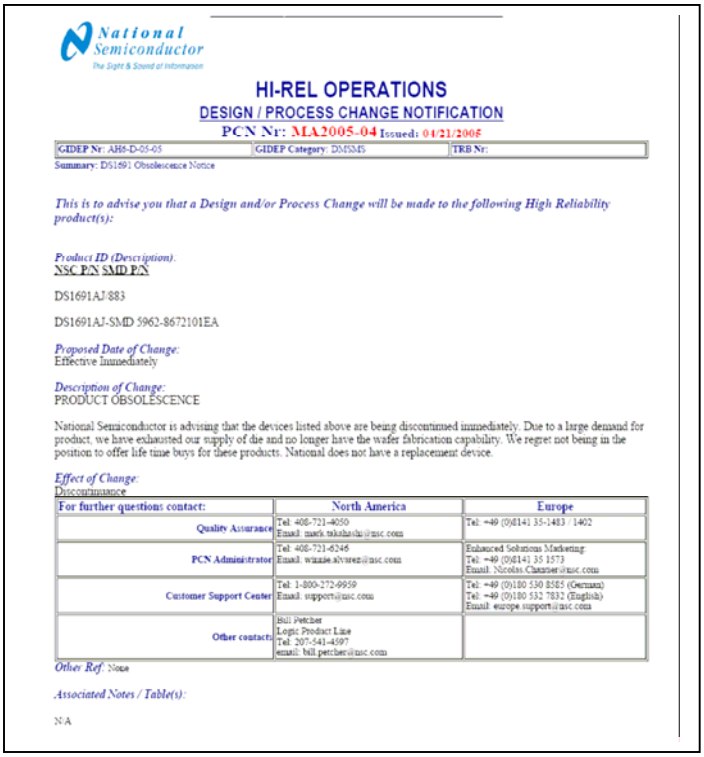

National Semiconductor

HI-REL OPERATIONS

DESIGN / PROCESS CHANGE NOTIFICATION

PCN Nr: MA2005-04 Issued: 04/21/2005

GIDEP Nr: AH6-D-05-05	GIDEP Category: DMSMS	TRB Nr:

Summary: DS1691 Obsolescence Notice

This is to advise you that a Design and/or Process Change will be made to the following High Reliability product(s):

Product ID (Description):
NSC P/N SMD P/N

DS1691AJ/883

DS1691AJ-SMD 5962-8672101EA

Proposed Date of Change:
Effective Immediately

Description of Change:
PRODUCT OBSOLESCENCE

National Semiconductor is advising that the devices listed above are being discontinued immediately. Due to a large demand for product, we have exhausted our supply of die and no longer have the wafer fabrication capability. We regret not being in the position to offer life time buys for these products. National does not have a replacement device.

Effect of Change:
Discontinuance

For further questions contact:	North America	Europe
Quality Assurance	Tel: 408-721-4050 Email: mark.nikkhahi@nsc.com	Tel: +49 (0)141 35-1483 / 1402
PCN Administrator	Tel: 408-721-6245 Email: winnie.alvarez@nsc.com	Enhanced Solutions Marketing Tel: +49 (0)141 35 1573 Email: Nicolas.Channer@nsc.com
Customer Support Center	Tel: 1-800-272-9959 Email: support@nsc.com	Tel: +49 (0)180 530 8585 (German) Tel: +49 (0)180 532 7832 (English) Email: europe.support@nsc.com
Other contacts:	Bill Petcher Logic Product Line Tel: 207-541-4597 email: bill.petcher@nsc.com	

Other Ref: None

Associated Notes / Table(s):

N/A

Any newly fabricated replacement part must meet the original device specifications while also being fabricated in modern high volume semiconductor facilities. That is the crux of the problem – how can a manufacturer obtain reliable, accurate reproductions of older parts or where necessary, entice the IC foundry to build such replacement parts?

Factors for Consideration

Faced with diminishing or obsolete parts, what can an industrial, aerospace, or military OEM do? Certainly one could buy the entire remaining inventory, but given the number of individual devices this would be an expensive solution. Commercial parts are not an option because they are not rugged enough to meet the original specification and they will run out too, sometimes even sooner. (This is not to imply inadequacy in the case where a commercial part was originally specified.) So, in both cases the OEM faces the same issues.

Beginning at the foundry level, several factors contribute to the reduced desire of modern semiconductor fabrication facilities to address the demand for older technologies directly:

1. Low Volumes – A Fundamental "Disconnect"

The necessarily low wafer volumes needed to supply the obsolete parts market are totally unattractive to semiconductor fabs/foundries. These facilities are accustomed to thinking in

terms of 200-300 wafers per month per part number for what they call a “small job.” Compare this to what a typical military buyer needs and you can begin to see the disconnect.

As an example, several years ago I worked for a major semiconductor IDM (integrated device manufacturer). I was approached by a friend, who was also a U.S. Government employee. He informed me that he had large volumes of parts needed for an up coming LEO Satellite Program. He related the need for hundreds of devices in each satellite, up to a total of 1000 satellites. To give him an idea of what his challenge was, I took him through the math, assuming there were multiple use die in each satellite (that is the same part number). Here’s how it breaks down:

Assumptions:

1. 1000 satellites (this came down to under 500 before the program was closed) – See Figure 2.
2. 100 “common” ICs per satellite (same part number, or made from the same wafer/die)
3. 1 year build out cycle (not realistic, but it makes the math easier)
4. Spare parts at 100% the number of built out parts
5. Complex die, so we guessed a die size of 5mm x 6mm
6. Technology was 0.25 micron on 8 inch wafers
7. Total fab, assembly, test yield = 85%

Figure 2
LEO Satellite Concept

The reality:

1. 1000 satellites times 100 devices = demand of 100,000
2. Plus spares at 100%, so total demand = 200,000 net die
3. Manufactured need = 200000/0.85 = 236,000 gross die
4. Gross die per wafer about 1000 (allowing for scribe lines and PCM die)
5. Total number of wafers needed to meet the demand = 236,000/1000 = 236, round up to 250 wafers.

I further explained that the total capacity of the fab at my company alone at that time was just under 1 million wafers per year. When adding in the pure play foundries at that time, plus the top 10 IDMs the total industry capacity exceeded 20 million wafers per year. We quickly came to the conclusion that the 250 wafers represented by the satellite program was truly an insignificant number for the foundries (250/20 million = 1.25×10^{-3} percent of capacity).

2. Lack of Process Knowledge – a real gap.

Modern IC designers place blocks (polygons) similar to the system designers of thirty years ago. That's both good and bad. Without the "systems" approach to IC design we would never have been able to design the microprocessors, DSPs, and other complex circuits we have today. But the designers today enjoy no real process knowledge. Designers for the most part do not understand the process used to make the transistors at the fab/foundry of their choice. Additionally they are limited in what they can do to change performance. They are constrained by design rules to the amount of change they can make to the basic element structure and are held to changes within some limited scaling range which has been characterized.

Historically a knowledge of the process was important when designing a device. There were many nuances of a process that designers of the period would exploit in the design of their devices. And, unlike the designer of today, designers of old could adjust the layout to change performance. That same awareness is important in today's re-designs.

For example, early on in my design career we used to have a 3.0 micron bipolar flow – very state of the art at the time – which was the core process for our TTL lines (both 54 and 74, plus LS). It also formed the base for the 10K ECL process.

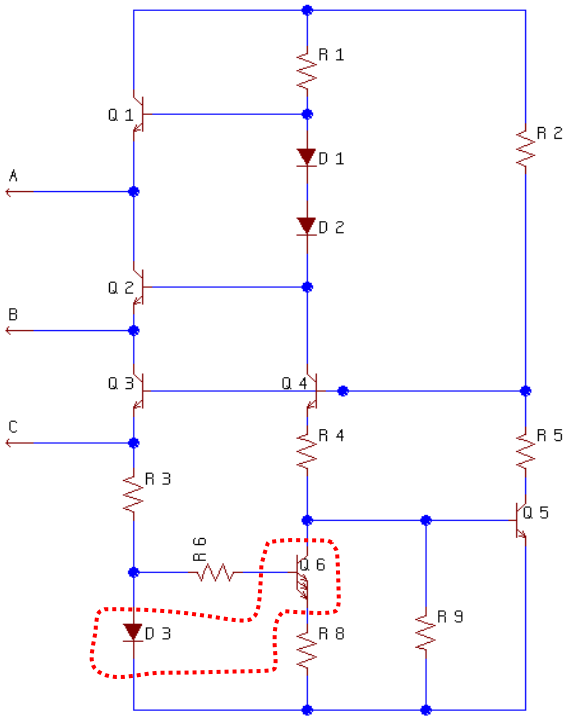

Figure 3
On chip bias driver with ratioed Q6:D3

In addition it had a linear negative temperature coefficient for V_{BE} as a function of rising temperature. By forming a ratio between multiple emitters on a current source transistor (Q6) and a single emitter on an adjacent base-to-ground diode (D3) in the on-chip bias driver, we could obtain a positive temperature coefficient differential. This could almost exactly compensate for the negative temperature coefficient of the process.

Recently an unnamed designer tried to re-design the ECL part with the bias driver above but he did not understand the process trick that was used. Instead, he placed a MOSFET in the bias driver in an attempt to make a temperature compensated device – it failed to work and tracked improperly with temperature. He followed modern device rules and placed polygons rather than thinking about the process itself and how it might work. As mentioned, if every designer worked to the transistor level, we would never have achieved the IC developmental progress of the past 4 decades.

The successful device designer of obsolete ICs has sufficient process knowledge such that he can tailor the IC design to take advantage of process anomalies and/or peculiarities.

3. Economic Imperatives – there's just no room for "Specials"

Today wafer throughput and yield, i.e. large-scale semiconductor manufacturing economics, drive the fab. The end result is a process flow optimized for the "few" as compared to the past when numerous flows were available to provide optimal product performance for many different part numbers. We had lots of different die, but the yields were low and the total revenue produced by the fab did not reach theoretical high (wafers-in/wafers-out x ASP per wafer). The driving force here was the change from a product centric business model to a process centric model – driven by the large economic powerhouses like TSMC. In the new model the process must be capable of being sourced by both the in-house fab as well as the foundry. Historically the process was unique to a given manufacturer but today the advent of System-on-a-Chip products from fabless semiconductor companies requires a controlled and well defined process.

As mentioned on the first point of this segment, semiconductor IC manufacturing is a volume business. The only way one earns a profit in a volume business is by making lots of parts to cover the fixed overhead first and then the variable costs.

In semiconductor manufacturing today, the cost of constructing a state of the art 12-inch fab for 130 nm and below, is on the order of 3 to 4 billion dollars. That's just for the building, air handling and filtering equipment, environmental monitors and tools, and the manufacturing equipment itself.

Then you have the process related expendables (like HEPA filters), the process related tools (like the SMIF boxes), the specialty chemicals and photoresists, the masks (reticles), the wafers, and the people – the "variable" costs. So it is easy to see why the processing of these wafers is so expensive.

As a plant manager you would be measured on how well you covered your fixed and variable costs. So you would want to:

1. Keep the fab as full as possible ("Fill the fab") – maximum wafers out.
2. Run as few different processes as possible (time is lost when switching from one process type to another) and time is throughput – or the fab metric; wafers out divided by wafers started. Also know as, "Cut the Scrap."

Clearly, stopping a process that is running at full capacity with few variations is counterproductive to achieving the lowest possible per wafer costs. Stopping one to make a small run of a different part is clearly not cost effective. This is the case for all fabs, not just the 12-inch in this example. Although the true costs are different for an 8-inch or 6-inch line, the basics apply and at the end of the production day the plant manager has a high fixed plant and significant variable costs to cover.

There are now over 68 wafer foundries listed by the FSA Fabless Semiconductor Association (FSA), not counting the foundry operations of some of the major IDMs like Toshiba offering custom system-on-a-chip solutions.

When one looks at the magnitude of wafer volume required just to achieve financial breakeven it is surprising.

As shown in the Figure 4 graph (illustrating a foreign foundry) the modern 8 inch, 0.18 micron fab needs almost $15 million a month just to reach breakeven at 15,000 wafer starts per month (WSM) with wafers priced at $1000 each. The breakeven point grows as

Figure 4
8 Inch Revenue to Breakeven

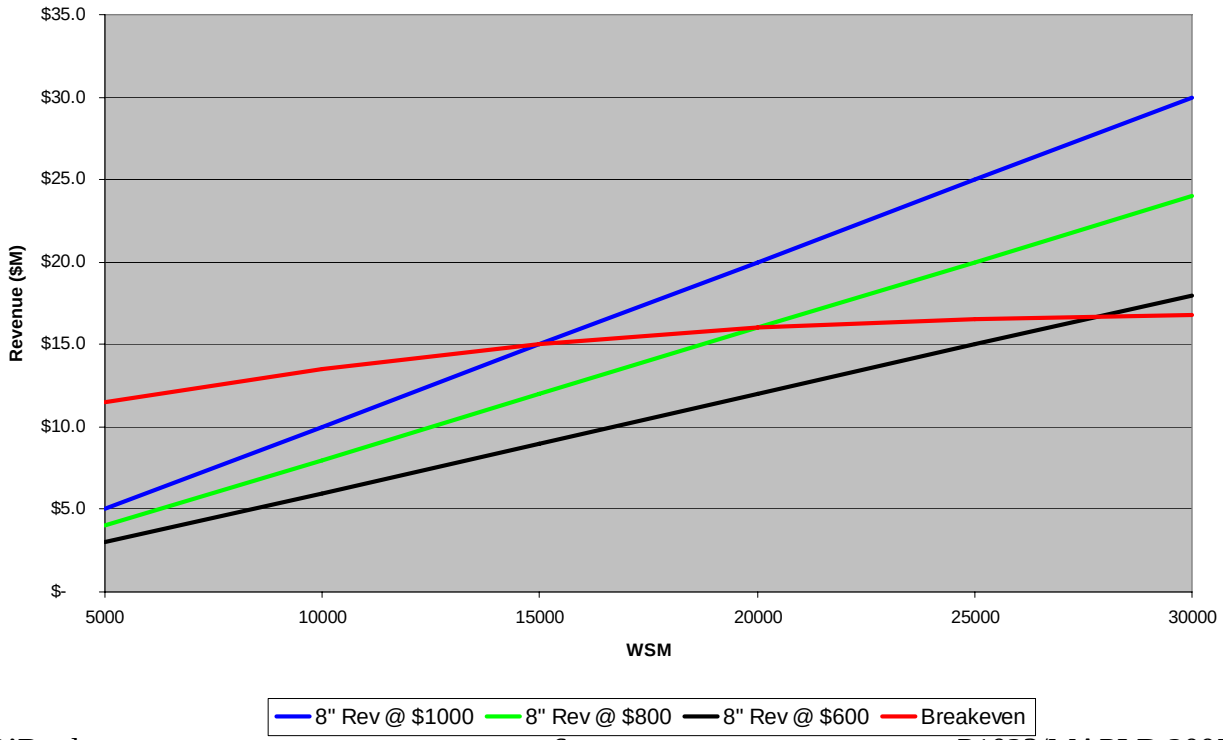

a function of variable costs such that at $800 per wafer the breakeven is about $16 million and 20,000 WSM.

Similarly a new 12 inch, 90 nm fab would need about $40 million a month at an ASP of $2500 per wafer at 17,000 WSM to achieve breakeven. (See Figure 5) This is what drives the managers of these fabs to go for the volume accounts.

Figure 5
12 Inch Revenue to Breakeven

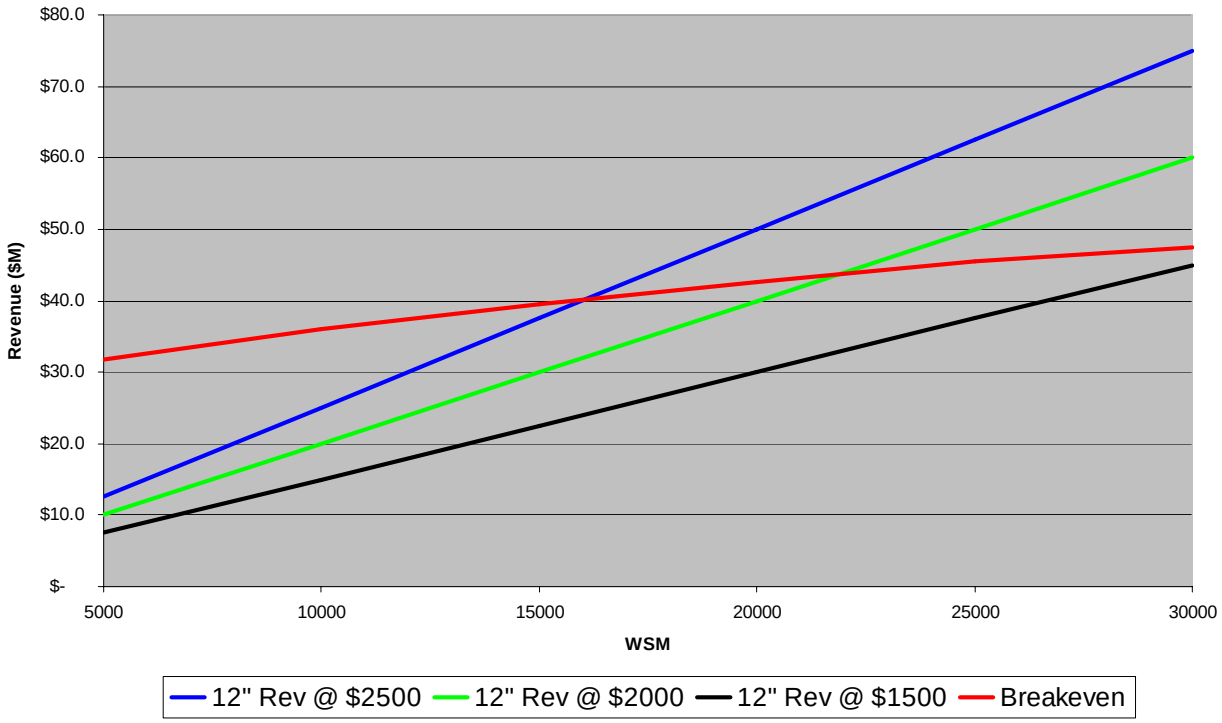

At this point you may be wondering how will anyone ever get special parts fabricated? Bottomline, the lack of economic motivation at the fabs requires that the obsolete parts maker think in new ways to accomplish the goal of making older parts.

4. Scarce Expertise – "Well, back in my day…"

Finally, the "graying" of IC designers impacts the ability of companies to make the older specification parts. There are fewer process engineers around who understand the architecture of designs invented in the 70s. As noted earlier it is important to have a device designer that understands the process, at least well enough to figure out ways to exploit its anomalies and hidden features.

In the past, designing ICs began with paper and pencil, with the use of hand calculated Karnaugh maps to simplify the logic and then form the schematic to achieve the logic

elements of the truth table. Mylar grids with Mylar paste-up transistors were then used to draw in the resistors (at 60 ohms per square, or whatever) to match the schematic. In those days, rubyliths and huge 500x cameras that floated on mercury were used. To make emulsion glass plates the size of the wafer, glass plate chrome masters were utilized. We were inventing it as we went along and unfortunately, lack of documentation does not help the modern obsolete designer in re-engineering these "classic" devices.

By comparison, today's designer of obsolete ICs is becoming an artisan in the true literary sense. He must refer to the old design information, the truth tables and the schematics to see what was actually done. In reality the schematic is typically just an approximation of the true device since it was probably revised between first design and final release to production. So the modern designer must think like the designer of the past and copy as much as possible the performance and characteristics of the obsolete design and process in a new process. This is no small task and surely requires technical talent and experience.

Possible Solutions of Today

Tackling the sourcing challenges of obsolete or diminished supply devices is challenging and calls for creative solutions to balance between the economic/volumetric burdens of the foundries and the very real aspects and objectives aligned with the utilization and design of classic devices to achieve the objectives of given programs. The following sections enumerates solutions available in the market today.

1. **Size Does Matter –**

Besides the very large pure play foundries such as TSMC and UMC, there are also interesting, small, dedicated foundries focusing on research like projects. They're willing to take on small jobs of a few dozen to a few hundred wafers. Additionally, many of them are willing to customize a flow for a given customer. In some cases these small fabs have affiliated designers who, when working with a customer, are able to design to an older flow. Their prices are typically higher, however, than when the fabs were originally running and higher on a per wafer basis than today's "pure-play" foundries. But the fact is they are willing to supply small quantities.

The only precaution is processes may not be as well controlled as the "pure-play" foundries and the major IDM's. As a result, there may be some process induced shifts from lot to lot. That's not a big problem and it will be caught in wafer sort and rectified quickly. Plus, because they typically don't have the original designs or the original process, a significant amount of non-recurring engineering charges may be needed to yield a device working to the original specs.

2. **The Emulation Approach**

Emulation has shown itself to be useful in select applications where the device is not longer manufactured and the original company either has closed or has released interest in the device.

The solution has proven that programmable devices can be made which can provide a similar solution to an obsolete device in a potentially cost effective fashion. The present day emulation approaches focus on digital but some analog is beginning to arrive, albeit for the simpler analog devices. This approach could still be realistic for a great many future low-volume devices, and many may find resurrection in the emulation methodology.

3. The Multi-Project Die - A New Approach

A relatively new solution to the problem involves an engineering-intense approach that is just evolving. It's a technique like the well-understood Multi Project Wafer (MPW). This new approach expands on the MPW technique by deriving several different "options" from a given die. That is effectively the approach.

As we embarked on this approach, we began with small PROMs – 16k and smaller. We noted that there appeared to be a growing shortage of the older low-density PROMs ranging from small 256b up to the 16k sizes. Upon further investigation we discovered that the shortage was real. Unfortunately, the demand for any one part was not sufficient to cover even the basic mask tooling costs for a mask set for the one part let alone the foundry processing costs.

So we took a page out of the ancient IC designer's handbook and investigated how we could make a family. Even with the original devices, it was common practice to make several devices from a single layout and either bond out or just change the metal mask to make a different device. Typically these optioned parts were limited to 2 to 4 additional part numbers from the single core die.

This new approach began simply as illustrated in Figure 6 where four devices are made from the same core and each different part number is merely a metal mask option. In the illustration metal mask options are used and the base silicon contains all the elements needed to form the logical variations AND, NAND, OR and NOR. The metal mask (in this single metal flow process), configures the logical operation desired such that all the masks below and above the metal are identical for all 4 devices. Only the metal mask changes for each specific product type.

The investigative process led further to an interesting discovery as we explored an even more audacious design approach for the PROM. Since we were not concerned with absolute maximum gross die per wafer (GDPW) we could build 3 "families" using the approach that one larger PROM could begat several smaller devices. We came up with definitions for small, medium and large families and the constraints which developed were more directly related to package and number of I/Os than absolute die size. We were able to create these families with numerous options for outputs, for example TTL totem pole outputs, various input and other output structures too. When we wanted a particular PROM we just needed to specify which metal mask to use and the part was done. The core remained unchanged. Imagine a solution where the options are all on a given reticle and by selective pinouts and/or metal mask options different devices are selected.

Figure 6
QP1631, QP1632, QP1633 and QP1634

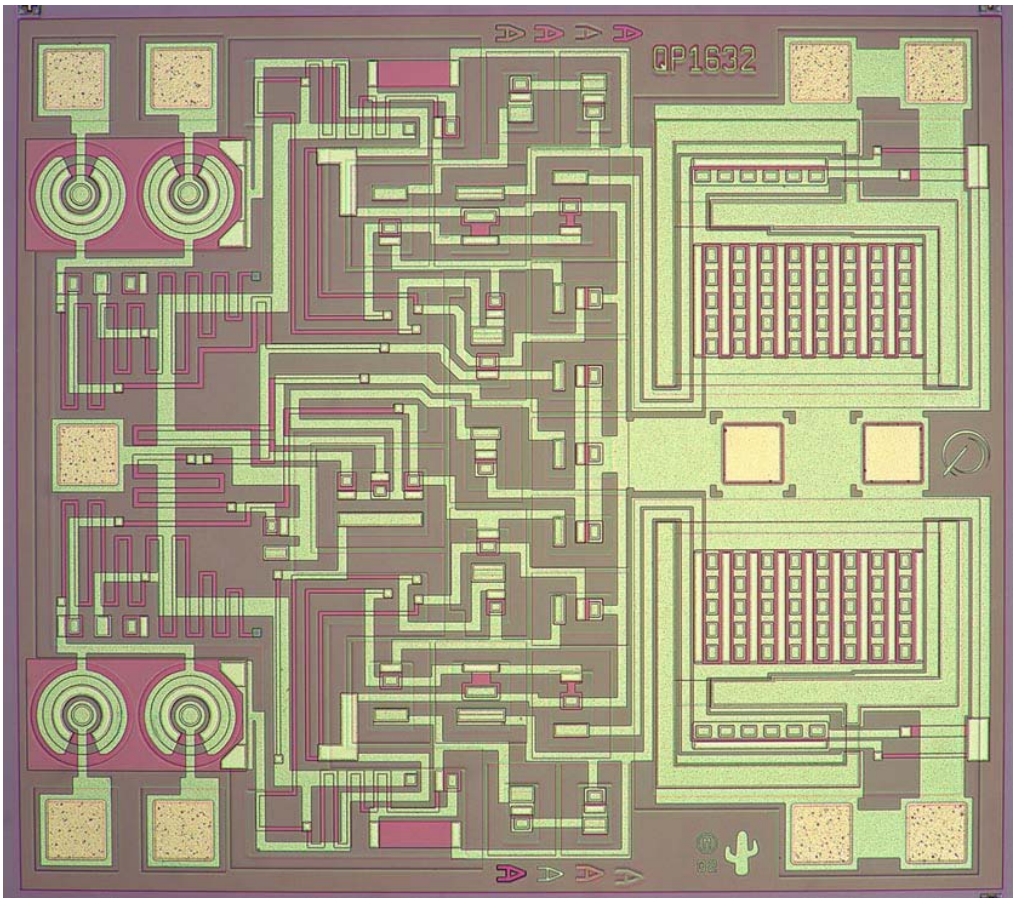

To further illustrate the concept, the following set of figures (Figures 7 through 10) show two different output structures for two different TTL devices – note in the layout shots the disconnected devices on the two different shots.

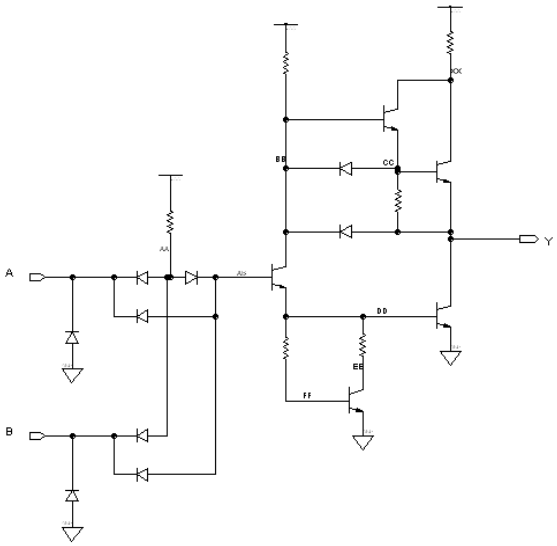

Figure 7
Schematic of
Output Structure 1

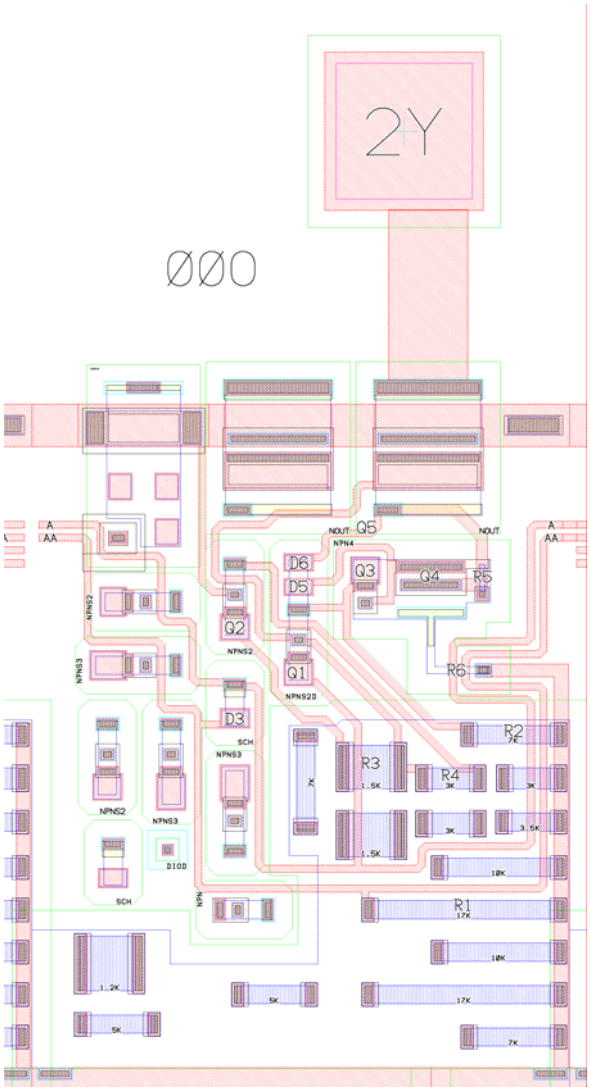

Figure 8
Physical Layout of
Output Structure 1

On the packaging side we ran into a few small difficulties. The dual inline packages (DIPs) were the most usable and flexible since there were numerous cavity sizes available to meet the 3 core device sizes. Other packaging approaches were more constraining and some posed a problem, especially the surface outline (or small outline) SO type packages. The advantage we had was that the families all had common power and ground pin placement which helped package design immensely.

At this point, we now had a solution to the manufacturing mask tool cost (since we could amortize the mask costs across several different devices) and the volumes were sufficient to cover the production costs. At the same time, we also tackled the packaging issue. However, the total demand was still small by commercial pure play foundry standards, so we needed a smaller foundry willing to work with us. We found a foundry partner and arranged a program whereby we have an average of 24 wafers in the line, on hold at metal (this is a single metal process),

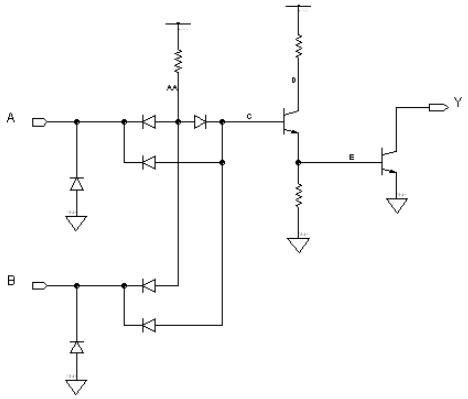

Figure 9
Schematic of
Output Structure 2

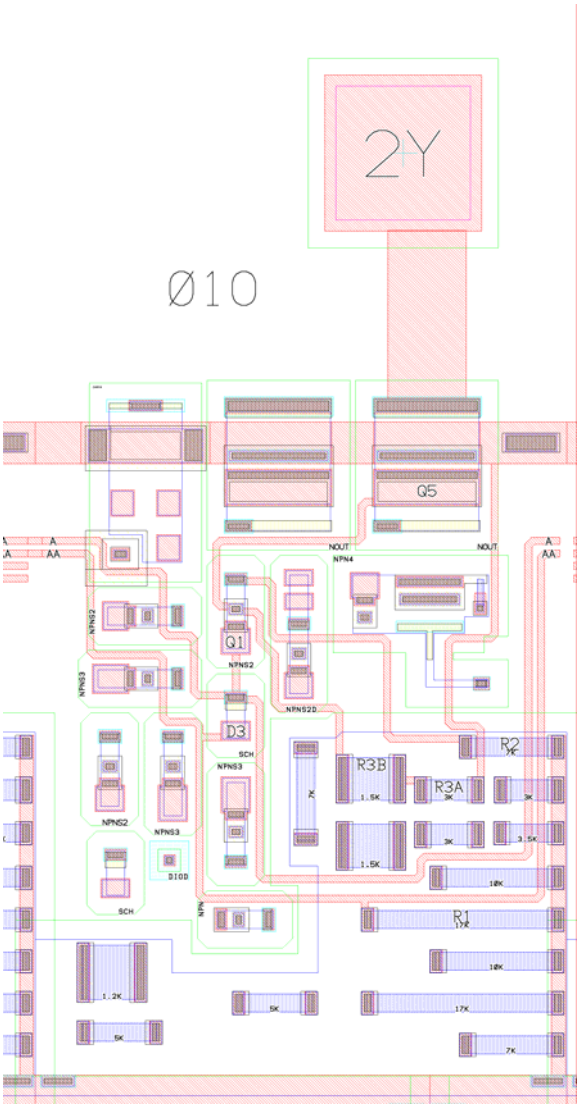

Figure 10
Physical Layout of
Output Structure 2

waiting for us to define the metal mask to build out. When the average wafers on hold drops to 12, they start another 24 wafers so we fluctuate between 12 and 36 wafers in the line. This may have been a small volume by the big foundry standards but it is quite acceptable to the selected foundry.

Today, the PROMs are designed and fabricated and now we're optioning out sequential circuits to be followed by fused programmable memories. The success of this approach helped us raise our wafer volumes, also making the partnership attractive to our foundry partner.

Lesson Learned: Plan on 2 spins through the foundry – the first run will be close, but the design may need a few small adjustments to hit the target.

Summary –

Large foundries are not interested in building obsolete parts for many reasons. The economies of scale that drive the modern pure play foundry are exactly the opposite of what the obsolete device maker desires. There is very little common ground.

The successful device designer of obsolete ICs must have sufficient process knowledge such that he can tailor the IC design to take advantage of process anomalies and/or peculiarities.

Bottomline, the lack of economic motivation at the fabs requires that the obsolete parts maker think in new ways to accomplish the goal of making older parts.

Therefore it is necessary to take a different approach focusing on the factors important to the obsolete device customer:

1. Availability of devices
2. Electrically and functionally equivalent
3. Pin-for-pin equivalence

Focus on what is important to achieve the performance, target smaller foundries and adapt designs and layout to allow for more flexibility in die size while keeping the costs within reason.