

"A Novel Approach to Obsolescence Management"

A discussion on obtaining reliable and accurate reproductions of older integrated circuits

John O'Boyle
QP Semiconductor

Introduction

Many devices still in use today for military specification (mil-spec) and high-reliability applications were designed and fabricated in the 1970s and early 1980s. During the 1990s, Commercial-Off-The-Shelf (COTS) Integrated Circuits (ICs) came into fairly widespread use. Both mil spec and COTS devices have unique issues concerning obsolescence management. The mil spec parts, although engineered and/or designed over thirty years ago, continue to be specified in new applications due to their reliability in harsh environments and their high performance. For similar reasons, and the key consideration being cost, COTS parts from the near past have also been designed into applications.

Unfortunately, the same environmental ruggedness has been difficult to achieve in the more advanced ICs developed and manufactured today. As a result, industrial, aerospace and military manufacturers have relied on mature technologies for reliable devices for many taxing applications. 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 design.

A satisfactory solution for these problems is to specify the newer high performance/high density COTS devices for advanced processing/decision making while using the older devices for interface and line driver/controller types of applications. This is the ideal solution in many cases, though it can often be thwarted by lack of supply for both the COTS and the older devices.

This situation illustrates an old and recurring problem. Manufacturers have faced the obsolescence of hardware and components 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 likely due to the complex nature of IC design and manufacture and the associated costs of running a manufacturing plant to fabricate older devices.

Ironically, one of the issues which COTS sought to address, namely availability of commercial parts, is the main reason for the shortfall. When a commercial device becomes obsolete, typically the device and the application become obsolete together. For example the CDMA cell phones of 3 years ago have been superseded by GSM phones today. The one technology making the other virtually obsolete in a short three-year cycle (this doesn't include the cell phone technology permutations which enabled the introduction of color screens, Internet access, Short Message Service (SMS), Instant Messaging (IM), and cameras into the phone). COTS parts can reach end of life very quickly, especially when compared to mil-spec parts.

But the mil-spec parts are not immune to rapidly declining supplies. In some cases the mil-spec parts just vanish without any warning. Recently a major military IC manufacturer found demand

had far outstripped their forecasts. The company unilaterally discontinued supply of these devices because they ran out of wafer stock before they could announce a last time buy.

OEMs contemplate making a replacement device when faced with such sudden obsolescence. But 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 Contributing to Parts Obsolescence

Faced with diminishing supplies, 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 rugged enough to meet the original specification and they will run out too, sometimes even more quickly than older high-reliability devices. In both cases the OEM faces the same dilemma. Beginning at the foundry level, several factors contribute to the reduced desire of modern semiconductor fabrication facilities to address the demand for older technologies.

1. Low Volumes – A Fundamental “Disconnect”

The necessarily low wafer volumes needed to supply this market are totally unattractive to semiconductor foundries accustomed to thinking in terms of 200 to 300 wafers per month per part number for what they might call a “small job.” Compare this large volume to what a typical military buyer requires, and the disconnect becomes evident.


A Case in Point: Satellites

As an example, several years ago I worked for a major semiconductor manufacturer. A satellite maker informed me that they needed volumes of parts for an up coming LEO Satellite Program. They needed hundreds of devices in each satellite, and were projecting up to a total of 1000 satellites. Now to understand the challenge, let's look at the math: assuming there were multiple use die (i.e., many with the same part number) in each satellite. Here's how it breaks down:

Assumptions:

1. 1000 satellites (this came down to under 500 before the program was closed) – See Figure 1.
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, an estimated die size of 5mm x 6mm
6. Technology was 0.25 micron on 8 inch wafers
7. Total fab, assembly, test yield = 85%

Figure 1
LEO Satellite Concept



The reality:

1. 1000 satellites times 100 devices = demand of 100,000 die
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.

The total production capacity at my company alone at that time was just under 1 million wafers per year. When adding in the pure play foundries plus the top 10 IDMs the total industry capacity exceeded 20 million wafers per year. So quickly one can arrive at 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. Without this “systems” approach to IC design we would never have been able to design the microprocessors, DSPs, and other complex circuits we have today. Designers for the most part do not understand the process used to make the transistors at the foundry of their choice. Additionally they are limited in what they can do to alter performance. The amount of change they can make is constrained by design rules to the basic elemental structure and changes are limited within a characterized range.

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

Case in Point: An Ineffective New ECL Design

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 ECL process. 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 would compensate almost exactly for the negative temperature coefficient of the process. See Figure 2.

Recently a 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 tracked improperly with temperature. He followed modern device rules and placed polygons rather than having a process knowledge from which to draw.

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

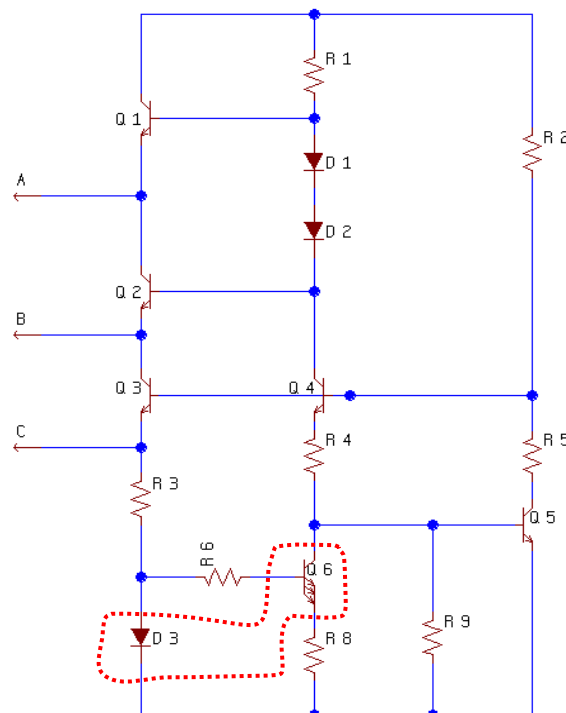


Figure 2
On chip bias
driver with
ratioed Q6:D3

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

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. Today wafer throughput and yield, i.e. large-scale semiconductor manufacturing economics, drive the fab. This results in 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. The driving force here is the change from the product centric business model of old to the new process centric model.

A semiconductor plant manager today would be measured on how well he/she covered fixed and variable costs. So they 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 known 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.

4. Scarce Expertise – “Well, back in my day...”

Finally, the “graying” of IC designers impacts the ability of companies to make the older 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, hand calculated Karnaugh maps to simplify the logic and then the schematic to achieve the logic elements. Mylar grids with Mylar paste-up transistors were used with hand drawn resistors to match the schematic. 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 planned to be 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 requires significant 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 enumerate solutions available in the market today:

1. Size Does Matter

Large pure play foundries aside, there are interesting, small, dedicated foundries focusing on research like projects. They are more willing to take on small jobs of a few dozen to a few hundred wafers. 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. The good news is that they are willing to supply relatively small quantities.

2. The Emulation Approach

Emulation has shown itself to be useful in select applications where the device is no 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 involves an emerging engineering-intensive approach. Similar to the well-understood Multi Project Wafer (MPW), this approach expands on the MPW technique by deriving several different “options” from a given die. 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.

Case in Point: 34 Parts from One Design

This new Multi-Project die approach which we engineered at QP Semiconductor began simply as four devices made from the same core and each different part number being merely a metal mask option. See Figure 3. Only the metal mask changes for each specific product type. Ultimately we evolved this approach to where the most recent design holds 34 different parts numbers.

Figure 3
QP1631, QP1632, QP1633 and QP1634



To further illustrate the concept, the following set of figures (Figures 4 through 7) show two different output structures for two different TTL devices – note in the layout drawings the disconnected devices on the two different images.




Figure 4
Schematic of
Output Structure 1




Figure 5
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 core device sizes. 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 34 different devices) and the volumes were sufficient to cover the production costs.

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 located 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), waiting for us to define the metal mask to build out.




Figure 6
Schematic of
Output Structure 2



Figure 7
Physical Layout of
Output Structure 2

When the average wafers on hold drops to 12, they start another 24 wafers so volume fluctuates between 12 and 36 wafers in the line. This may have been a small volume to the big foundry, but it is quite acceptable to the selected foundry. The success of this approach helped raise wafer volumes, making the partnership attractive to our foundry partner.

Design Process Expertise and Foundry Selection is Key

Large foundries are not interested in building obsolete parts. 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 lack of economic motivation at the fabs requires that the obsolete parts maker think in new ways to obtain older parts.

Today, the successful device designer of obsolete ICs must have sufficient process knowledge in order to tailor the IC design to exploit the process anomalies and/or peculiarities.

This new approach focuses on the factors important to the obsolete device customer:

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

At QP Semiconductor, in our experience, we now target smaller foundries and adapt designs and layout to allow for more flexibility in die size while keeping the costs within reason, allowing us to center new designs on achieving virtually identical performance to the replaced device.

###

John O'Boyle manages the development of new business relationships as director of business development for QP Semiconductor, Inc. He also directs the company's strategic planning in support of firm's core business, mission-critical military/aerospace and high reliability industrial programs. O'Boyle's semiconductor technology and business experience ranges from worldwide strategic marketing responsibilities at Samsung and Tower Semiconductor to commercial space business development at National Semiconductor. His background also includes a stint at Dataquest, where as a vice president and director he developed an entirely new business unit (Consulting Services) that addressed the needs of government, industrial and consumer customers. Earlier in his career, O'Boyle held engineering positions with Color Planar Displays Inc. and Fairchild Semiconductor. He holds BSEE, MSEE and MBA degrees from Santa Clara University.

About QP Semiconductor:

Founded in 1985 as a test services company, QP Semiconductor ranks today as one of the largest fabless manufacturers of high-reliability ICs for the military, aerospace and industrial systems. From its headquarters in Santa Clara, California, QP Semiconductor applies design, re-engineering and a full range of manufacturing capabilities to extend the life of classic ICs for a host of mission-critical programs. For more information, visit www.qpsemi.com.