Lowering Barriers to Entry for ASICs

Whitepaper by TIRIAS Research
June 20, 2017

There has never been a better time to build your own custom application specific integrated circuit (ASIC). Despite the talk of Moore's Law slowing and the cost of new chips rising, there are many opportunities to turn your sensor-driver design or your multi-chip controller into a small ASIC to lower costs and protect your intellectual property (IP).

This paper will explore the different ways in which companies are building chips that reduce cost, space, power, while adding features, and protecting the designer's IP.

Why ASICs?

OEMs traditionally designed their own chip or turned to semiconductor design houses for custom SoCs (ASICs), either because they needed processors that were not available or they desired differentiation from competitors. However, only a small portion of the OEMs have the in-house capabilities to build their own ASIC. Most other companies come to market with a less than optimal design in the interest of time-to-market or due to the lack of capabilities. But if the product is successful or if costs are too high, it is always possible to consider an ASIC solution designed in-house or through a design partner.

ASICs have typically included mixed signal devices where a range of analog circuity is integrated alongside a CPU to reduced components, board area, power, etc. ASICs can also provide better IP protection (it is harder to copy a chip than a Printed Circuit Board (PCB)), add additional functionality, and create differentiation. Going forward, many Internet of Things (IoT) companies are considering custom ASICs both in the form of single-die solutions and multi-chip modules to include an increasing amount of system IP, such as processing accelerators, memory, wireless interfaces, and sensors. As a result, most major semiconductor vendors and foundries also offer custom ASIC design services to select customers.

Silicon Business Models

When it comes to silicon designs, there are multiple paths a company may choose. In the case of a university start up, the team may choose a path using free-to-license cores because they may have little seed money and are willing to take a riskier path. To them, budgets are tight and taking higher risks is required. Also, design volumes are not a given, so the there's little negotiating room for amortizing design costs over production silicon.

A more established OEM will have a much better sense of the potential volume of the production design and can better judge the design investment risks. An established OEM also has the financial resources to invest in either a design team or to hire a design services company.

But not all OEMs have the expertise to build an ASIC to fit their needs. In these cases, the OEMs make do with standard Application-Specific Standard Processors (ASSPs) or microcontrollers (MCUs) with additional components needed for the specific applications. These companies can now explore building their own application specific integrated circuit (ASIC) as the barriers to entry are coming down (see Figure 1). In the case of processor IP, new programs like ARM's DesignStart offer low-risk access to proven processor cores and come with support logic.

Many mixed-signal designs can use the older process nodes and still extract good performance from small CPU cores, like the Cortex-M3 or Cortex-M0. Integrating the mixed-signal designs onto a single die can reduce power, signal noise, and component count, including supporting passive components. Additionally, proprietary circuit designs can be better protected on silicon instead of being available for all to see on a printed circuit board (PCB).

Building an ASIC always involves some risks – product delays, sub-optimal execution, cost overruns, etc. The challenge is to manage those risks as most designs are sensitive to time-to-market and costs. One advantage of using a well-proven set of IP cores, is that part of the risk is mitigated. The risk can be lowered even more by using an experienced ASIC design house.

Figure 1. Lowering the entry costs to ASICs.

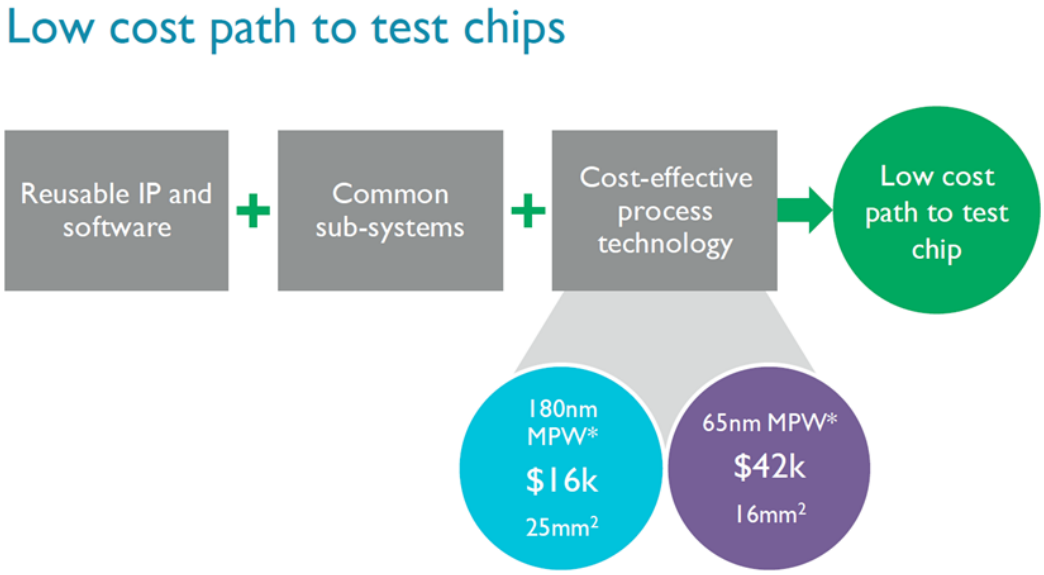

(*45 Samples. Source: Europractice)

Design House Choices

We talked to a couple of ASIC design houses to gather more information on the process companies go through to decide if an ASIC is right for them. Here's what we found.

Sondrel

We asked Sondrel's Director of Engineering, Paul Martin, about the primary reasons a company considers a custom ASIC, versus using an ASSP or MCU. According to Mr. Martin there are many factors, including the targeted industry, the amount of proprietary IP, the company's experience in chip design, etc. In the IoT market, which uses many of Cortex-M CPUs, companies often have a specific set of proprietary IP that they want to wrap control and support logic around to turn it into a product or cost reduce an existing product. With an ASIC design, the company can reduce the number of components, which reduces the Bill of Materials (BOM), makes the form factor smaller, and reduces power consumption. The integration of components also reduces the system complexity and manufacturing costs, as a result of smaller boards with fewer components.

Sondrel is experienced in pushing the performance of older process nodes. These older nodes offer lower mask costs and cheaper wafers due to depreciated equipment and facilities.

Another reason for a custom ASIC design is transaction security. Often important transactions need to be protected and it is easier to protect them when kept on-chip. In the same way, critical and valuable IP can be protected from potential copy cats inside a custom ASIC. Many sensor ASIC designs include signal processing and proprietary IP.

Often the challenge for OEM customers is a general lack of knowledge of ASIC design. These OEMs may not have a silicon background or may have never developed a specific product before. For these customers, a field programmable gate array (FPGA) is also too complex to attempt. In addition, the investment and time required to build a design team is too high and takes too long. The design could also could be a one-off design eliminating the possibility of monetizing over multiple designs and products.

But a design house like Sondrel can guide these OEMs through the supply chain and provide a full documented design. A design house can also offer additional IP sources for the design and monetize that IP over multiple designs for multiple customers.

The costs for manufacturing an ASIC are not actually very high. It can be as little as $16K for prototype quantities (Figure 1). Costs are heavily dependent on the process node. Mask costs can double from one process generation to the next, such as 40nm to 28nm and 28nm to 16nm. Unusual nodes (e.g. high voltage) can also increase the cost. But many IoT designs are well suited for older

process nodes larger than 40nm; indeed, for mixed-signal devices, the analog components are often better served by 90 or 180nm. The low transistor count of cores like the ARM Cortex-M0 and Cortex-M3 cores are also well suited for older process nodes, and the volume required for an ASIC can be as low as about three-thousand units.

For some OEMs, the key to success is monetizing the entirety of solution, where the ASIC design costs are not the biggest concern. But for any company, there may be only one shot at market success, so why risk the design using unproven IP? Working with a design house, they have the methodology and design flow that's critical to a successful first-time design. Sondrel touted their well proven project management flow for this. Their people are an experienced team who understand risks. They also have local partners in Europe and now own their own silicon validation labs.

While Mr. Martin tried to provide more detail on specific designs, the more complex designs are under NDA.

S3 Group

Another design house we talked to was the Ireland-based S3 Group. We talked to Darren Hobbs, Director of Marketing and Strategy, The S3 group is an expert in mixed-signal and RF design. While the company may be a little under the radar, it has a 30-year history that includes a spin out from Philips. The S3 Group has 250 people, with 120 in semiconductor design. Its target market is mostly non-consumer - communications, satellite, automotive, and industrial. These are markets that are still slowly adopting wireless and may still require wired connectivity.

The S3 Group's mixed-signal expertise is a specialty that requires a longer design cycle. As they pointed out, there is no one size fits all design. Its customer base is more mature, but are concerned about size and power over cost.

The benefits to OEM customers are largely economic - both top line revenue growth and bottom line profit. An ASIC may be able to add more features to the product for an up sell. The ASIC also protects IP. The bottom line can be improved by reducing manufacturing costs with a smaller PC board and reducing the number of discrete components.

Darren indicated that the design costs break down to NRE and volume. Design costs can be the highest when mixed signal is required. Choose badly and it will be risky. But manufacturing costs can be quite low as wafer costs on 180nm and 65nm have never been lower. Those nodes have a broad ecosystem as well.

Most of the OEM customers S3 Group supports uses existing ARM cores – it has really become the de facto standard.

One design example S3 Group could cite was in the area of industrial control (oil and gas industry) - a controller design converted to an ASIC in order to add functionality and more connectivity. It used an ARM Cortex-M3 for temperature control.

When considering alternative cores, it's difficult to rationalize using anything but ARM. S3 Group did have a design that used the open sourced LEON core seven years ago, but there were concerns with warranty for the core. How can you guarantee correct operation of an open sourced core? This then places the burden on the OEM to fully validate the core design.

In S3 Group's experience, the typical ASIC runs 50K units or higher over a 5- to 7-year life cycle. The company works with customers on yield models for chip costs, and can share in any yield upside.

Figure 2. ARM Cores Are Shipping in the Billions

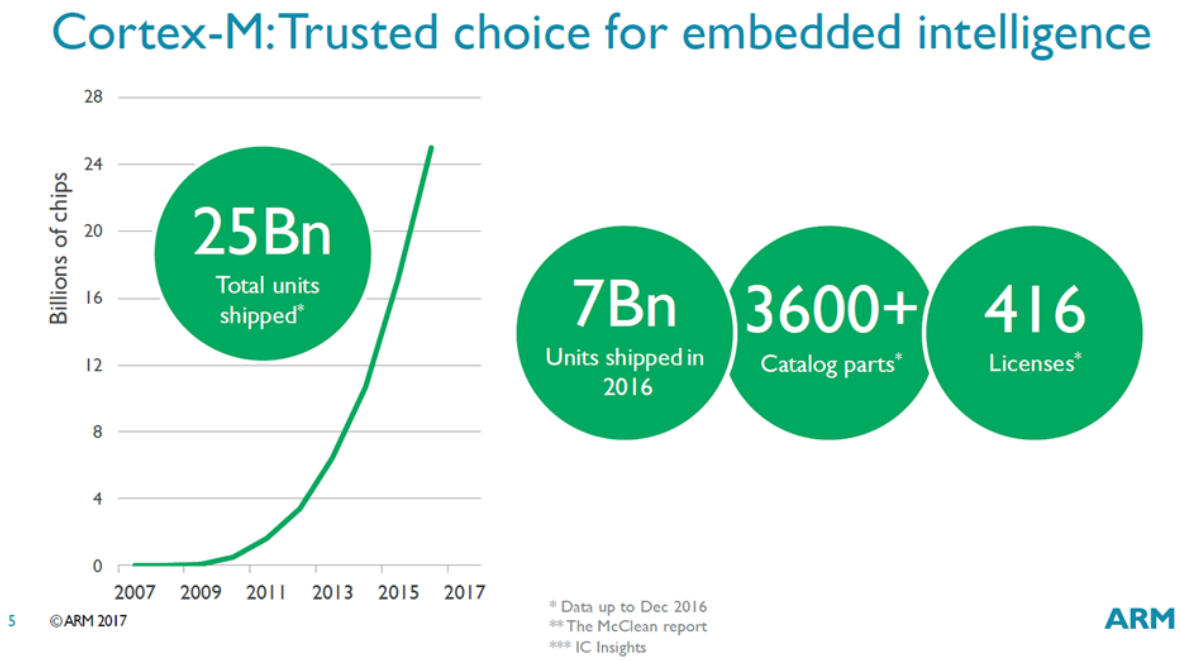

(Source: ARM)

ARM DesignStart Program Evolved

Along with all of the many Physical IP libraries that ARM has, the DesignStart program initially included just the Cortex-M0, but recently it added the Cortex-M3 and a verified subsystem, making building an ASIC easier than ever. Adding the M3 allows the program to address new markets.

Figure 3. DesignStart Program Evolution

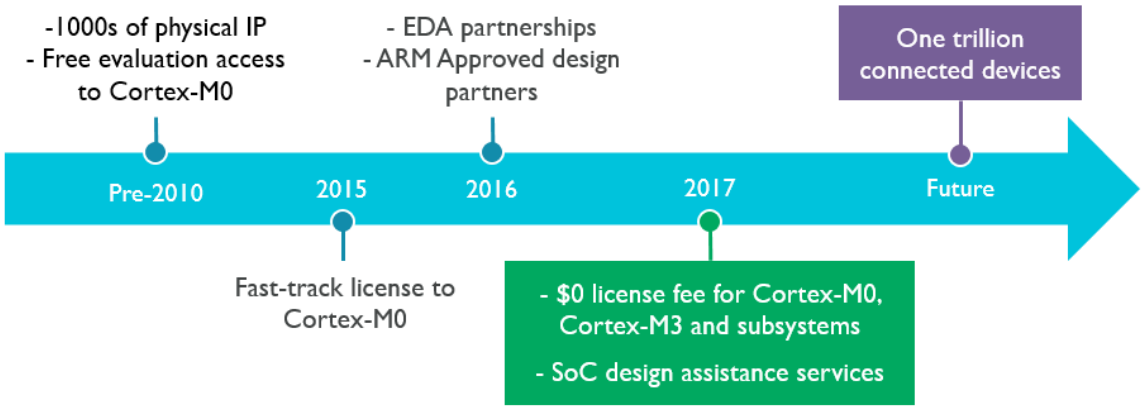

(Source: ARM)

For companies that have chip design expertise in house, ARM DesignStart provides a way to get started with a minimal investment. For evaluation purposes, ARM offers free instant download of a fixed configuration CPU and subsystem IP to all. Designers can evaluate, develop and prototype. ARM provides out of the box support on an ARM supplied FPGA platform ($495), or designers can use their own preferred FPGA if they prefer.

Once the project begins full commercial SoC development, the DesignStart program has a $0 license fee and is built with a success-based royalty model, which lowers the monetary risk on a new design. Customers can download the full product IP within a small number of days of registering. As ARM's program is not tied to specific foundries, low-volume designs can leverage an aggregator like IMEC or Mosaic.

The program includes the processor core and a verified subsystem. Additional subsystem IP includes interconnects, and bridges, and small amount of peripheral IP, such as UART and timer, and a true random number generator (TRNG) for the Cortex-M3. For IoT designs, the program also includes access to mbed OS.

As part of the DesignStart program, ARM offers design support services for a fee or an enhanced SoC development services option in conjunction with ARM Approved design partners. The ARM Approved program has seven design houses today that are audited by ARM to assure a quality

product. These design houses are scattered around Europe and Asia. Some of these design houses have expertise in mixed signal, RF, and other specialties. These design houses have the resources and expertise to deliver a quick time-to-market.

ARM IP Mitigates Risk

The reasons ARM cores are popular is because they have highly validated processors and support IP, which can reduce the risk of a new design. The cores are "battle proven" and come from a trusted source (ARM), allying any concerns on the overall support and reliability of the core. In addition, the completeness of the software tools (including valuation suite) is hard to beat (See Figure 3).

Figure 3. Extensive ARM Cortex M ecosystem

(Source: ARM)

Conclusion

An OEM should always consider if an ASIC is the right solution for its design. Even if an OEM does not feel it can build a ASIC on its own, there are design support services to help – letting companies of any size to build custom chips. Working with a proven IP provider, a OEM can be assured that their design will be correct and have long-term support.

For an OEM, an ASIC can significantly reduce manufacturing costs by reducing printed circuit board (PCB) size, which can reduce packaging. It also reduces the number of components that need to be sourced, inventoried, pick and placed on the PCB. The ASIC design can also lower power consumption by integrating passive components, which can reduce power supply costs and reduce thermal issues.

Using a proven IP provider like ARM and its extensive and capable ecosystem will give your products a better chance to succeed in the market. Lower cost ARM licenses for ARM Cortex M0 and M3 IP should further lower the bar for designing and manufacturing cost effective ASICs. Even start-up that are bootstrapping with little capital can develop their designs using the ARM DesignStart program with little risk.