Introduction

Many established large system companies as well as new startups are looking to build custom chips to meet specific market needs that cannot be met with off-the-shelf applications processors. This is especially true for applications that require unique IP, high-bandwidth memory, and cutting-edge workloads like AI. Many capable application specific standard products (ASSP) vendors build very capable chips, but these ASSPs are a compromise to address multiple applications or markets. Yet these are some applications that are best served with a chip designed specifically to the application. These custom chips are commonly called application specific integrated circuits (ASICs), and a number of companies have been building ASICs for decades.

Over the years, ASSPs and field programmable gate arrays (FPGAs) have played a more visible role in the industry. Yet, some applications require specialized or proprietary intellectual property (IP) and chips that need a specific architectural design that no ASSP vendor can supply, nor can be efficiently implemented in an FPGA. For those challenging designs, an ASIC is still the preferred solution.

These ASIC vendors vary in capability. Even traditional chip companies like AMD have entered the ASIC business with its semicustom division, mostly to build chips for video gaming consoles. One vendor TIRIAS Research has talked to in detail is eSilicon, which has built some world-class chips for leading systems manufacturers and high-performance startups. Top companies rely on eSilicon to build their toughest chips.

The challenge for eSilicon is that often its customers do not want it known that eSilicon built their chip. The reasons for the obfuscation vary; it could be for corporate pride, or to tip off the competition. TIRIAS Research was able to talk to a number of eSilicon customers and other system designers to get a perspective on the options available to startups and established system companies for building complex ASICs. In deference to the requests of the customers, their identities are anonymized.

Why Build an ASIC?

System companies and OEMs traditionally design their own chip in house with a team or turn to ASIC companies when they need something unique or differentiated. Most OEMs do not have a silicon design team capable of building and managing an ASIC program internally. Some may use an FPGA or a non-optimal ASSP if they feel a need to rush the design to market. If that product proves successful or if costs are too high, the OEM can look into building a cheaper, more efficient ASIC. But if there is no ASSP alternative nor an FPGA capable enough, building an ASIC is required.

Building an ASIC in house may be difficult without the right team. If it is an ASIC with cutting edge technology, then experienced ASIC companies can take the design through production at competitive prices. ASICs can also provide better IP protection; it is harder to copy a chip than a printed circuit board (PCB). ASICs also add additional functionality and create differentiation.

As a result, most major semiconductor vendors and foundries offer custom ASIC design services to select customers.

For some startups, building a circuit and back end support team is thought to be a requirement to be in the chip business. But increasingly, it's more important to build a board-level or system-level product to deliver the most value. If this is the case, then it makes more sense to off load the back end to an experienced ASIC company and focus on system level design concerns and software tool support. It's often software that is more important to deliver the product value and often there are more software engineers involved in a high-performance system product (like machine intelligence) than hardware engineers.

Silicon Business Models

Companies may choose among multiple paths for silicon design. A university startup may choose a path using free-to-license cores, because they 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 they have little negotiating room for amortizing design costs over production silicon. One such example of this path is the SiFive team that is using the open source RISC-V processor instruction set. These are not usually customers for eSilicon's high value capabilities, but some more advanced AI applications should consider their services.

A more established OEM will have a 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 an ASIC company. For an OEM, the chip is designed to support a systems product, their costs and expertise are around the systems design, not the chip.

But not all OEMs have the expertise or resources to build an ASIC to fit their needs. In these cases, the OEMs make do with standard ASSPs with additional components needed for the specific applications. These companies can now explore building their own ASIC as the barriers to entry are coming down.

Additionally, proprietary circuit designs and other intellectual property can be better protected in custom silicon instead of being available for all to see on a 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 time- and cost-sensitive. One advantage of using a well-proven ASIC company is that part of the risk is mitigated.

For some OEMs, the key to success is monetizing the entirety of system solution, where the ASIC design costs are not the biggest concern. But for any company, there may be only one shot at market success. A proven ASIC company will have the methodology and design flow critical to a successful design. The eSilicon team is experienced and understands risks. OEM customers can benefit economically (both top line revenue growth and bottom line profit), technically (added product features), and legally (IP protection). The bottom line can be improved by

reducing manufacturing costs with a smaller PC board and reducing the number of discrete components.

Building an ASIC invariably requires collecting specialized IP for the chip. This IP includes standard I/O, high-performance SerDes, clocks and PLLs, memory controllers, *etc*. For a systems company or a startup, collecting the best IP can be a challenge. The experienced eSilicon design team has access to this IP and makes the collection process easier. In particular, eSilicon is known for fast serial / deserial (SerDes) chip-to-chip communications link designs.

Experience Counts

With experience building chips using TSMC's 2.5D CoWoS technology, eSilicon has extensive 2.5D experience, including IP design, systems in package (SiP), manufacturing, assembly and test. The company has projects that include HBM2 memory systems and other 2.5D implementations. These design techniques are used in many new AI, datacenter, and communications/networking applications.

Figure 1: 2.5D Packaging Technology

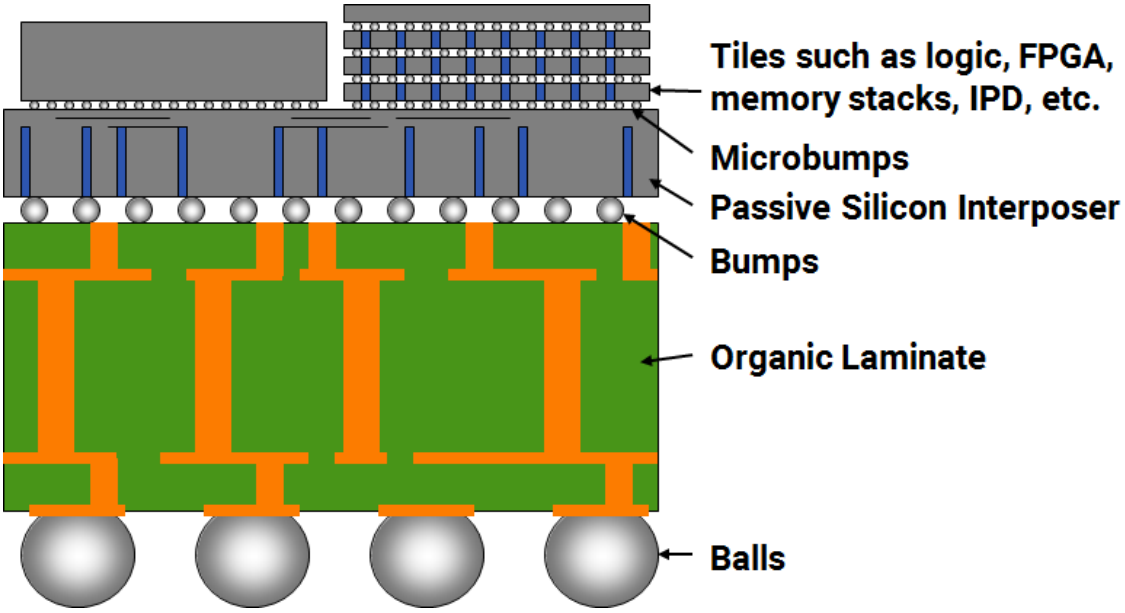

Source: eSilicon, 2017

The company also offers leading-edge internal memory designs such as a ternary content-addressable memory (Ternary CAM or TCAM), which is an excellent choice for packet forwarding and classification in internet routers, as well as in a variety of other applications that require high-speed table lookup. TCAMs are unique: they search an entire lookup table in just one cycle. The eSilicon TCAM compiler creates customized TCAMs that provide up to 2.5 billion search results in one second (GSPS), enabling high-efficiency, cost-effective solutions for applications such as network search engines, cache for network processors, QoS services, classifications, Ethernet, ATM switches, and other diverse networking applications.

The eSilicon design team has extensive experience in analog, mixed-signal, and SerDes designs having successfully proven two 56G SerDes architectures in silicon at 28nm. This team is now working on a 56G long-reach SerDes in 7nm technology. Architectural work on a 112G design has begun as well.

One of eSilicon's advanced capabilities is its fast cache which offers speeds that exceed 3.3GHz under worst case operating conditions. The company also has several multi-port SRAM architectures to support parallel operations to increase bandwidth. These include dual-port SRAM, as well as pseudo 2-port and 4-port SRAMs.

The company's new business is now 70% of the cutting edge technology FinFET process nodes (16nm and below) and quite a few 2.5D packaging technology projects. These will become even more important as the industry is challenged by slowing Moore's Law transistor shrinks, which will require multi-die solutions to scale performance and better yields. The multi-die approach splits large die into smaller die and then connects them through high speed SerDes on a silicon substrate. This technique can improve yields and can scale to areas larger than silicon reticles.

While a lot of startups can use eSilicon's expertise, the challenge is that many startups do not make it or do not reach volume manufacturing. One recent project was for a startup building a 2.5D chip with high bandwidth memory on a silicon substrate. eSilicon successfully delivered on the project but the startup was bought by a larger semiconductor company. While eSilicon will take this design into production, future generations may be brought in-house at the larger semi company. Still, the company can deliver 2.5D chips that meet the standards to a top-tier chip company.

To support top-tier companies, eSilicon can deliver these advanced technologies in volume and also meet top reliability standards.

Recent Projects

The eSilicon design at an AI startup required fast access to DRAM. Much like GPU vendors AMD and NVIDIA, memory bandwidth can be a constraint on processing the workload. Using TSMC's CoWoS 2.5D technology required a close working relation with the foundry. Having close access to TSMC's latest process and packaging technology is a big plus.

With one big system house, eSilicon's fast SerDes technology was a critical component. Their ASIC is used in very high end networking equipment and having fast, on-die SerDes reduces power, area, and costs. The networking chip has to handle a lot of fast data, so fast memories and leading edge FinFET process nodes are a must. In addition, system companies often have unique IP that needs to be included in the chip. Only an ASIC company can deliver that custom experience without building an entire design team in house. For the system house, building a full design team in house for a limited number of projects often does not make sense. Also, finding a team with the experience with leading edge packaging and process nodes is hard to do.

And the design work at eSilicon is more than just producing the chip. There's full system design that takes into account packaging, thermal profiling, testability, assembly and yield. This is

especially challenging with the new high bandwidth memory and 2.5D silicon interposers (Figure 2).

Figure 2: 2.5D Packaging Design

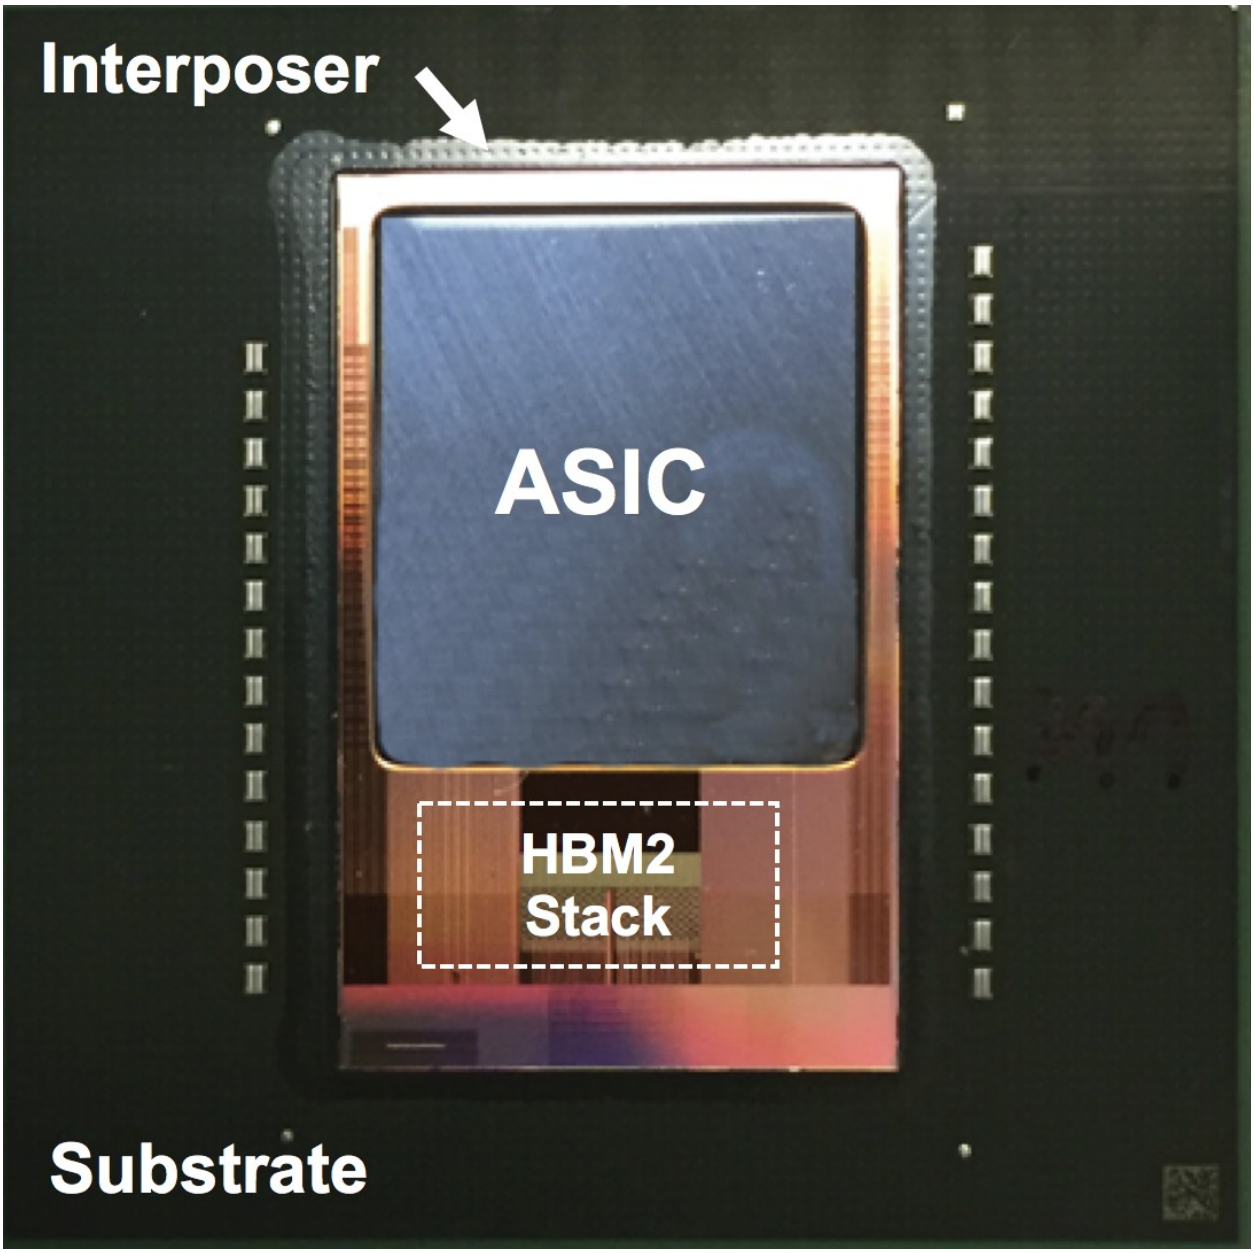

Source: eSilicon, 2017

Having worked with very demanding customers, the team at eSilicon can handle the stringent requirements these companies demand. Yet, eSilicon is also shown it can be flexible by bidding on jobs where the technologies are so new, the learning process has to take place during the development. The company must make a deep engagement with the customer's design team and be flexible to take new netlists on a regular basis. Because eSilicon is still a relatively small company, it still has that small company flexibility, yet it offers the technology capability of a

much larger company. In all these cases, its eSilicon that takes the risk on hitting the financial milestones and unit costs. There does have to be a high level of trust between eSilicon and the customer because this is a sole source chip and eSilicon needs to deliver or the product often cannot be made. To be sure the design will be successful, there is a high level of collaboration between eSilicon and the customer. The shared goal is to bring the product to production as soon as possible. The company has a sophisticated supply chain network which can help ramp products into market.

These deals also require eSilicon to be expert as a system-level designer, not just at the silicon level. The real revenue stream comes with unit shipments, and that requires a successful product. And that requires timely delivery and a working product.

Go Big or Go Home

Jack Harding, eSilicon CEO, says the company decided four years ago it needed to go from good to great. The company focused on leading edge products that required advanced design expertise. It also required the company to compress the research and development cycle. Concurrent design, R&D and manufacturing – it would learn new techniques as it was developing the products. Compressing the research and design cycle had a risk, but it would also separate eSilicon from the pack.

Presently eSilicon has multiple 2.5D products in development and production ramp, something only much larger companies like AMD, NVIDIA, and Xilinx would attempt. The company is driving 2.5D design to monolithic yield rates. The company pivoted from broad based to narrow focus to Tier 1 to allow it to focus on more established opportunities, but it still will take a risk on some cutting-edge startups.

Conclusion

Once you decide that an ASIC is appropriate for your design and you need a partner with advanced system level expertise, you should include eSilicon. The company is well positioned as the premier builder of high-performance ASICs.