Technology Brief…

March 15, 2022

J.Gold Associates LLC, 6 Valentine Road, Northborough, MA 01532, USA
www.jgoldassociates.com +1-508-393-5294

Research, Analysis, Strategy, Insight

Why Chiplets are Critical to Future Processors

Recently a group of chip companies and infrastructure vendors, including Intel, Advanced Semiconductor Engineering Inc. (ASE), AMD, Arm, Google Cloud, Meta, Microsoft Corp., Qualcomm Inc., Samsung and Taiwan Semiconductor Manufacturing Co., announced the establishment of an industry consortium to promote an open die-to-die interconnect standard called Universal Chiplet Interconnect Express (UCIe). Conspicuously absent from this announcement is Nvidia, but we suspect it's more of a timing issue than reluctance on its part, and it will likely join in the near term.

UCIe was originally developed by Intel, who has donated the specification to the group to define an industry-standard way to interconnect between "chiplets" on SoCs. The goal is to create an open system whereby various chiplets from various vendors can be "mix and match" compatible for packaged devices. This is an important step forward, and not just for the benefit of Intel or the other chip companies, as the support of non-chip vendors demonstrates. The entire industry will benefit.

As we move to next gen SoCs, they will require components (Chiplets) from a variety of vendors. Much as we build computers today with standard buses that let companies pick and chose board level components or backplane components to plug in, the same dynamic is moving to the SoC world as chiplets will be applied, potentially from competing vendors, to create either a standard device or a customized SoC. In fact, the customization piece is a huge reason why this is important as it lets specialty components be included that wouldn't be economically feasible otherwise, much as can be done today on a PC board.

It will also significantly reduce the development time of new SoCs, as components can be reused (e.g., network interfaces, video components, security processors, memory interfaces, etc.). They won't need to be redesigned each time companies move to a new chip process node as often there is no real benefit to doing so for many computing functions. That makes it much easier to redesign and enhance devices, and shortens time to next generation products. It also allows use of older process technology (e.g., 16 nm or larger) for some of the chiplets without paying a penalty in performance and without impacting the yield on the newer, (e.g., 5nm), processes. The more complex the chip being produced, the harder it is to get good yields, so anything that can be done to simplify is a major manufacturing benefit. And smaller chips mean more chips per wafer, thus potentially reducing supply issues and lowering production cost.

All of the chip foundries will benefit from this work. If a chiplet is designed to this spec, it can then move to any assembly point. And not all SoCs will be built in the foundry. There are special assembly "factories" where chips are built, often separate from the foundry. If this new standard is followed, a vendor can grab a chiplet produced at any foundry and then assemble it on a substrate at any assembly location, at least in theory. It's very much like Nvidia or AMD building a graphics card and being able to plug into a Dell or HP computer because there is a standard interface (e.g., PCIe). So imagine a Samsung as an example, being able to buy a chiplet from Intel and then build out an SoC with the Intel processor but Samsung components like memory, 5G modems, etc. Or it could also make

Inside This Issue

"….UCIe, if fully implemented and standardized, will have a profound impact on how processors are produced and what functions can be included. While there is still a lot of work to be done to make components fully compatible, and its likely to take 2-3 years to get to shipping products, but chiplets are a revolutionary step in an effort to keep the promise of Moore's Law alive. It's likely that this technology will usher in a new world of more complex, more quickly updated and more customized chips…."

the processor chip interchangeable with an ARM device, while still keeping all the other components on the SoC the same (there are of course some other issues around full compatibility). In theory this is a benefit to all chip makers.

This can also help in the supply chain constraints currently being experienced and that are likely to continue for the foreseeable future. Chiplets can be built on older process technology, and there are many foundries that still have 2-4 generations old processes running, including in the US as well as other parts of the world like Europe, Japan, and Taiwan. Building chiplets on older process nodes opens up the leading edge fabs to produce more chips as some of the load has been removed. And it also means not having to rebuild fabs for new process nodes which is time consuming and expensive.

But this doesn't necessarily come for free. Moving to a chiplet world requires having assembly plants that can actually do chiplet assembly. Not all of them can at this point – many just take current monolithic chips and place them into a package. So, to move to a more chiplet-centric world, it means upgrading assembly plants, including the need to build substrates that include busses/interconnects (much like a PC Board) that the chiplets can be mounted on. That is a specialty in itself. But assembly, relatively speaking, is easier to do and less expensive to build out than a full blown semiconductor fab, and takes less time to build. That is a benefit that can be leveraged to build out more plants in the US if desired.

The UCIe approach can help spread out chip manufacturing to regions outside Asia to avoid things like losing TSMC and all of Taiwan chipmaking if China gets aggressive. Making it easier to produce chips, in the long run, makes it easier to build fabs. And it enables spreading out assembly to other geographies. Many are still in the Far East, although some are also in places like Vietnam, Singapore, Malaysia, etc. This could become a major benefit for countries that want to be large chip producers, like India or Eastern Europe. The cost of entry is now lower if a vendor can buy the most difficult chiplets to create and then produce the easier ones and package them locally.

Bottom Line: UCIe, if fully implemented and standardized, will have a profound impact on how processors are produced and what functions can be included. While there is still a lot of work to be done to make components fully compatible, and its likely to take 2-3 years to get to shipping products, but chiplets are a revolutionary step in an effort to keep the promise of Moore's Law alive. It's likely that this technology will usher in a new world of more complex, more quickly updated and more customized chips that will benefit all industries and consumers over the next couple of years.

T-Mobile T-IoT Enables Geographically Dispersed Connected Things

"…Connected IoT devices over a 5G network show much promise in making systems safer, more intelligent and offering two way communications in real or near real time. But the complexity of deploying IoT assets across multiple geographies and particularly for those that travel wide and far has limited the attractiveness. The T-IoT service holds the promise of simplified deployment and management…"

T-Mobile US and Deutsche Telekom (owner of a significant share of the US company) have announced a partnership to make connected IoT available across multiple countries. It includes a single point of contact, a unified management platform, and simplified cross-border licensing across all markets that T-Mobile serves, and includes unlimited data for each device. This partnership provides a solution to a vexing problem for many vendors of connected things – how to build a product that can easily be provisioned and operate across country borders without requiring a complex set of individual country-specific products and/or services.

5G is ushering in a whole set of new services and capabilities, but the complexity of cross-border implementations has made it difficult for vendors like car makers, industrial machine makers, transportation and shipping companies, etc., particularly in places like Europe where individual services must often be procured for each country. The T-IoT capabilities offer management across 188 destinations and 300 networks, taking this burden away from companies. It provides both cost benefits to those companies as well as ease of operation to the end users served as they no longer need to be concerned with where they travel to. Further, the need to meet varying compliance and privacy regulations across multiple geographies makes deployment extremely complex. Its imperative a service is available to manage and comply with all such regulations through a single interface rather than multiple control points, and that also provides expertise to advise its customers on the pertinent compliance processes to implement.

Current IoT connectivity solutions are hard to employ and manage, causing many companies to limit their ability to serve multiple geographies around the world. Indeed, the current need to deploy and manage different networks, contracts, platforms, multiple portals, multiple look and feels, establish multiple customer care capabilities, manage multiple customer touch points, etc requires a lot of resources that takes away from a company's ability to concentrate on its core business. The ability to centralize sales, contracts, global connectivity, platform choice, visibility of assets, and lifecycle management to one universal management point represents a major cost and business benefit. Further, companies also need flexibility in pricing structure, using various technology platforms and offering the right connectivity options at the right price, while simplifying deployments with one interface and a single bill across all geographies. And finally, while not yet a part of the T-IoT service, many companies will require a full device lifecycle management capability that will eventually enhance such single point of contact and procurement capabilities and make it far simpler to deploy IoT solutions.

While this service is applicable across many vertical industries and solutions, early adopters will most likely be centered on industries such as;

- *Automotive* which currently represents the single biggest piece of the IoT market and where car makers can create a single connected vehicle in one location and deploy it almost anywhere to enable telematics and services like streaming and music which are growing more important;
- *Healthcare* where connected devices must travel with the users of those devices as they move around and must meet stringent regulatory criteria; and
- *Transportation*, particularly for high cost goods like pharmaceuticals and technology products that get shipped to many different countries.

But we expect the service to make it attractive for many other industries over the longer term, with cost flexibility meaning sensors networks can be deployed using only the amount of connectivity they need (e.g., eventually deploying network slicing solutions to limit costs for lower data requirement IoT deployments).

Bottom Line: Connected IoT devices over a 5G network show much promise in making systems safer, more intelligent and offering two way communications in real or near real time. But the complexity of deploying IoT assets across multiple geographies and particularly for those that travel wide and far has limited the attractiveness. The T-IoT service holds the promise of simplified deployment and management, especially of non-fixed IoT assets. This service can alleviate much of the challenge in making connected intelligent devices a reality. We expect to see a major uptake of such services over the next 2-3 years as companies realize the benefits while also taking advantage of the reduction in complexity and cost advantages such a service offers.

RISC-V Gets a Push from Intel and Partners

"….RISC-V is an ascendant processor technology that will take 2-3 more years to become mature, but that holds promise for many current and emerging companies to offer products compatible with its open source instruction set…. RISC-V is beginning to reach a critical mass that makes it an attractive option. We expect a large number of both complete processors and "chiplet" designs to become available in the next 1-2 years…."

With an increasing need to build specialized processors for a variety of mobile, IoT, connected car/autonomous vehicles and all manner of connected things, there continues to be incentives to creating competing platform options to the dominance of the Arm and the x86 platform, while also providing more compute power than many current microcontrollers. Just as with the software open source movement taking hold a few years ago and producing fertile ground for arrival of cloud computing, organizations are looking for an open sourced version of computing platforms on which they can build their products.

RISC-V International was founded in 2015 and is based in Switzerland. It has grown from its original 29 members to 2k+ members in more than 70 countries. It was established as a non-profit organization supporting the free and open RISC instruction set architecture and extensions. They enable open community collaboration, technology advancements in the RISC-V ecosystem, and visibility of RISC-V successes. By creating an open standard instruction set, companies are free to produce devices that meet the spec and are totally compatible, but that can contain their proprietary chip designs. These chips can then be produced in a foundry, and RISC-V participating companies thereby join a growing list of other fables semiconductor companies, many of which currently support the Arm IP licensing market, but many are also exploring a RISC-V option.

The Arm acquisition attempt by Nvidia and the potential challenges it posed for many companies licensing Arm tech has shown that moving to a more open architecture not controlled by a single entity may be a safer move. Although RISC-V is being considered by some as an alternative to Arm, it is nowhere near as complete from an IP standpoint (e.g., no GPU or AI acceleration). But it provides an avenue for companies to be self sufficient while still gaining some IP to get them started. And because of the low cost of entry, it provides an opportunity for many startups to design specific workload processing functions and gain entry to the processor marketplace. Over time, we expect the ecosystem will build out and there will be more components and features available in RISC-V, making it a more compelling choice and more competitive with Arm, although clearly Arm is not sitting still. But there is enough promise that companies like MobileEye has chosen RISC-V as its core CPU engine for its next generation of autonomous vehicle chips.

Getting a Boost from Intel

Recently, as part of its IDM 2.0 initiative, Intel has been expanding its market presence, especially with a renewed interest in becoming a major foundry. Intel announced that its will be joining the RISC-V community and will be dedicating resources to manufacture chips based on RISC-V. It will be producing its “Horse Creek” RISC-V developer platform on its Intel 4 process technology for release later in 2022. Intel is not new to RISC-V, having already produced Nios® V processors based on RISC-V. Looking to accelerate adoption, Intel Foundry Services (IFS) will sponsor an open-source software development platform and will provide IP optimized for Intel process technologies, making it the only foundry to offer IP optimized for all three of the industry’s leading ISAs: x86, Arm and RISC-V. Intel stated that it has already seen strong demand for RISC-V offerings.

Bottom Line: RISC-V is an ascendant processor technology that will take 2-3 more years to become mature, but that holds promise for many current and emerging companies to offer products compatible with its open source instruction set. While Intel brings some needed heft to the platform, it’s not alone, as TSMC and Samsung are also supporting RISC-V. Working with partners like SiFive, Andes Technology, Esperanto, Ventana and others, RISC-V is beginning to reach a critical mass that makes it an attractive option. We expect a large number of both complete processors and “chiplet” designs to become available in the next 1-2 years.

J.Gold Associates, LLC
6 Valentine Road
Northborough, MA 01532 USA

Phone:
+1-508-393-5294

Web:
www.jgoldassociates.com

Research, Analysis, Strategy, Insight

About J.Gold Associates, LLC.

J.Gold Associates provides advisory services, syndicated research, strategic consulting and in-context analysis to help its clients make important technology choices and to enable improved product deployment decisions and go to market strategies. We work with our clients to produce successful new product strategies and deployments through workshops and reviews, business and strategic plan coaching and reviews, assistance in product selection and vendor evaluations, needs analysis, competitive analysis, and ongoing expertise transfer.

J.Gold Associates provides its clients with insightful, meaningful and actionable analysis of trends in the computer and technology industries. We have acquired a broad based knowledge of the technology landscape and business deployment requirements, and bring that expertise to bear in our work. We cover the needs of business users in enterprise and SMB markets, plus focus on emerging consumer technologies that will quickly be re-purposed to business use.

We can provide your company with a trusted and expert resource to maximize your investments and minimize your risk. Please contact us to see how we can help you.