

# Chiplets - Taking SoC Design Where no Monolithic IC has Gone Before (WP016)



---

April 04, 2019

White Paper

---

## Introduction to Chiplets

The well-documented slowing of Moore's Law is the key driver behind the movement towards the use of chiplets in the design and manufacture of new, high-performance semiconductor devices. While the monolithic IC has been the ultimate design target for many decades, there have always been reasons to build certain devices with multiple die using multi-chip modules (MCMs), whether it was for additional memory capacity or to fabricate chips based on IP blocks that require incompatible IC processes. An example is the pairing of bulk CMOS die with high-performance analog die that implement high-speed SerDes ports or high-performance ADCs and DACs.

The chiplet movement is a reaction to the rapidly changing IC landscape and the current IC fabrication realities. Engineers are increasingly realizing that it makes little sense to integrate every IP block in a system on one piece of silicon if the fit is poor. There are many advantages with monolithic silicon integration, but those advantages are rapidly being outweighed by the economics of building advanced technology ICs. It is extremely expensive and time consuming to integrate, validate, and tapeout chips that can only be justified with high volume demand. More importantly, adding functionality or creating variations to support multiple end products increases die size and costs. Chiplets are becoming an alternative solution. Fortunately, the ecosystem for chiplets is quickly developing to provide companies a new tool to create highly optimized and cost effective solutions for their various end markets.

Intel presented details about its embedded multi-die interconnect bridge (EMIB) chiplet packaging technology at the annual Hot Chips Conference in August, 2018. According to Intel, its EMIB technology "facilitates high-speed communication between multiple die in-package." Intel positions EMIB as a key component of the company's mix-and-match, heterogeneous computing strategy.

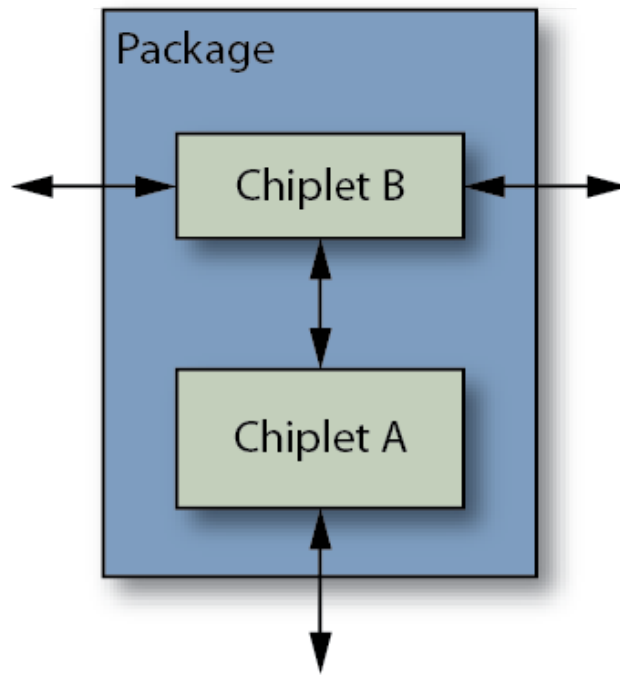
In November, 2018, AMD quickly followed suit by announcing that it will employ a chiplet design approach for fabricating their EPYC server processors using multiple Zen 2 processor chiplets to integrate as many as 64 x86 processor cores in a single package. According to AMD, these microprocessor packages will employ microprocessor die made with TSMC's 7nm manufacturing technology combined with I/O die made using a 14nm process technology. The I/O die will use AMD's Infinity Fabrics to interconnect the chiplets and will incorporate eight SDRAM interfaces.

---

## MCMs and Chiplets

It is not overly difficult to assemble modules from chiplets. The enabling MCM technologies have been available for decades, literally, but the required ecosystem and the infrastructure needed to make chiplet-based SoC assembly and testing economically attractive are not as well developed as they are for the creation of monolithic ICs. However, the death of Dennard Scaling and the end of Moore's Law are making chiplet-based device manufacturing increasingly attractive.

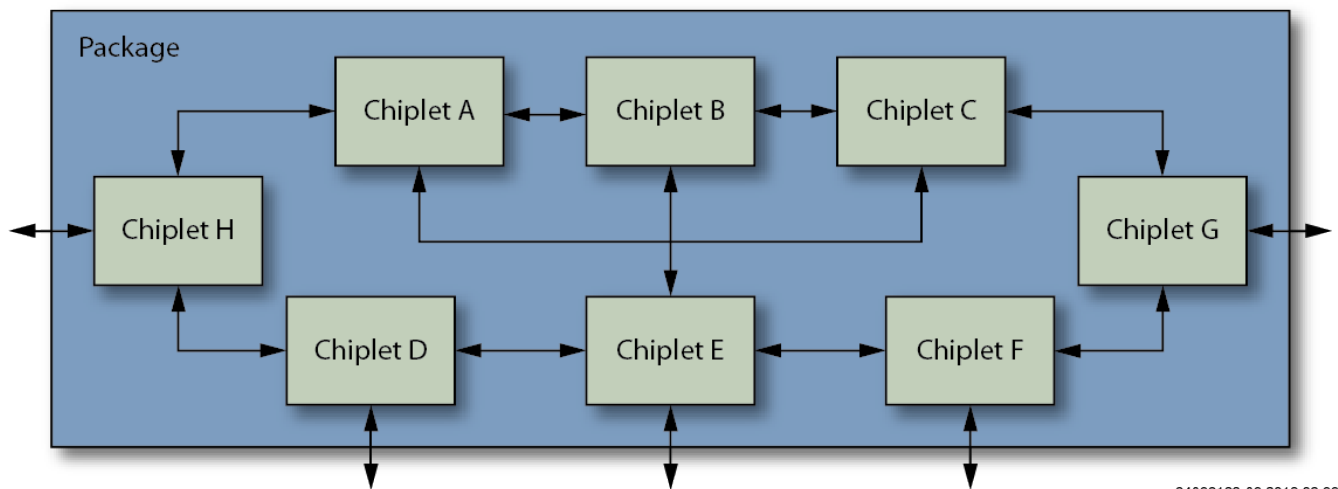
Chiplet-based MCMs can be as simple as two die in a package, as shown in below. Even this simple module provides several advantages including lower power, less board space, and better performance when compared to placing two packaged ICs on a PCB.



**Figure 1: A Simple, Chiplet-Based IC Module Incorporating Only Two Semiconductor Die**

The next figure shows a more complex example, where each chiplet in the package may have a different purpose or there may be multiple copies of the same chiplet in a package along with dissimilar chiplets. Chiplet functions can include digital logic, memories and 3D memory stacks, FPGAs, high-speed SerDes ports, high-performance ADCs and DACs, and even optical devices. Placing several chiplets in a common package allows for shorter, faster interconnections compared to individually packaged die on a PCB.

The semiconductor material used to make each chiplet is not limited to silicon, which is another chiplet advantage. For example, specialized chiplets could be made from a variety of composite semiconductor materials including SiGe (silicon germanium), GaAs (gallium arsenide), GaN (gallium nitride), or InP (indium phosphide) to exploit the unique electronic properties of these semiconductor materials.



**Figure 2: Several Chiplets in an MCM**

Each chiplet in an MCM is designed, taped out, manufactured, tested, and debugged separately — which permits a divide-and-conquer design strategy that can be more efficient and less risky than a monolithic design approach. Chiplets for an MCM can be designed all at once or, more likely, developed and used over time. Once all of the required chiplets are designed, fabricated, and debugged, then they are assembled into one MCM. New chiplets can also be designed at a later date to create new variations for enhanced functionality or to target new applications.

In the earlier days of MCM design, engineers used *ad hoc* approaches for die-to-die interconnect. The connections between and among die in an MCM conformed to no industry standards, because such standards did not exist (and still do not). Connections were generally parallel in nature because inter-die interconnect standards were lacking and a parallel connection is relatively simple.

For chiplet-based design to take off, *ad hoc* connections cannot be the norm. Standards are needed just as they were for computer peripherals (RS-232, GPIB, SCSI, PCIe, etc.) and for networking (Ethernet).

Along with being cost competitive with monolithic silicon solutions, a chiplet must meet the following requirements:

- High bandwidth
- Low latency
- Low power consumption

In addition, chiplet-based design requires substrate standards to become practical. Further, the issue of known-good die must be handled.

The big advantage of chiplet-based design is related to the way that multigenerational SoCs have been designed for many years. Often, as an SoC design evolved from one generation to the next, much of the SoC's design did not change. For example, more on-chip memory may be needed in a next-generation design. Certain interfaces will change from one generation to the next; others will not. It makes good economic sense to redesign, lay out, and tape out only those portions of the design that change. However, it is not possible to redesign just a portion of a monolithic IC design. If anything changes on a monolithic chip, the entire chip must go through redesign and re-verification to accommodate the changes and to ensure that nothing has been broken by the chip's redesign.

Similarly, if a monolithic chip is merely shrunk to take advantage of a new process node, the entire chip must be redesigned, re-simulated, and then taped out according to the design rules for that new process node. This holds true even if most of the functions on the chip remain unchanged.

A chiplet-based approach allows a design team to redesign only the parts of a design that must be redesigned. The remaining parts can be left as is. This method is an excellent, low-cost, low-risk design approach for creating multiple product variants of a basic design. It also simplifies the addition or deletion of options to an IC product family.

## FPGAs as Chiplets

---

Achronix provides the electronics industry with FPGA resources in three forms:

- Speedster® FPGAs
- Speedcore™ eFPGA (embedded FPGA) IP cores
- Speedchip™ FPGA chiplets

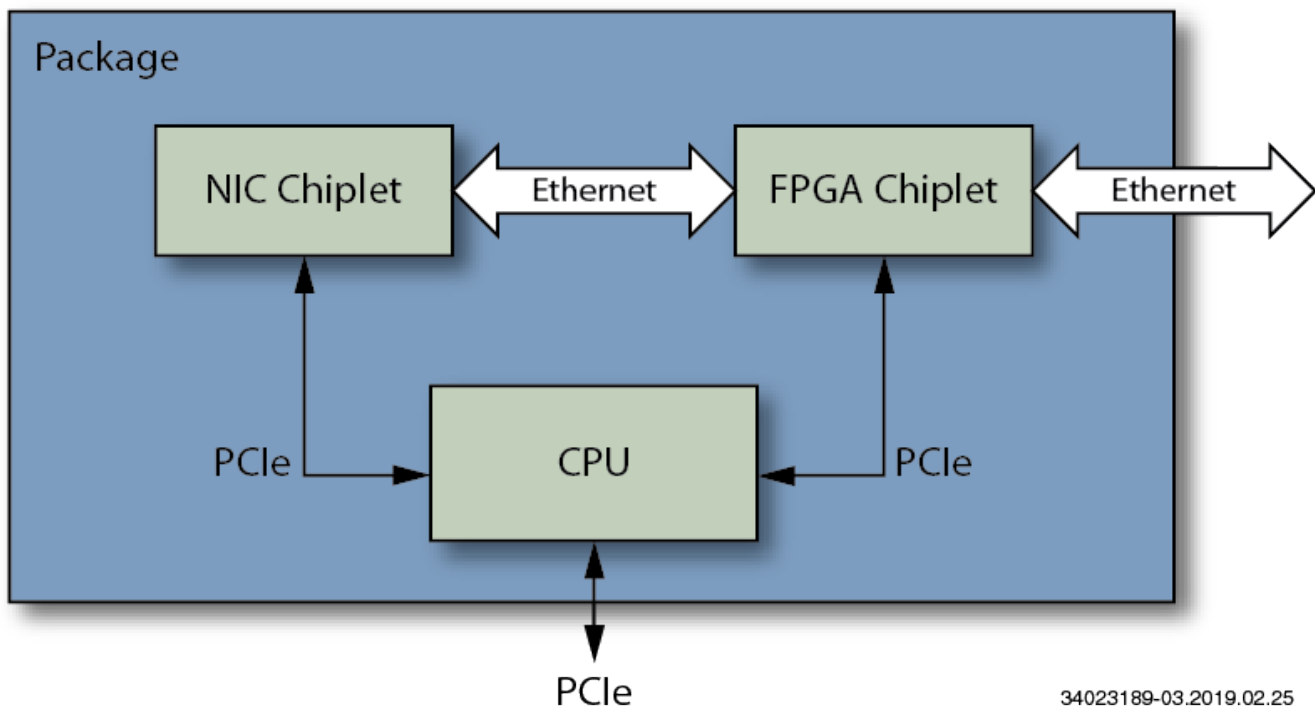
Assuming that a design requires programmable logic, Speedchip chiplets are a good way to add an FPGA to a semiconductor product. FPGA chiplets and MCM assembly complement the use of Speedcore eFPGA IP in monolithic ASIC and SoC designs. The two alternatives (Speedchip FPGA chiplets and Speedcore eFPGA IP) are not mutually exclusive either. Both design approaches can provide much tighter integration and a superior form factor compared to two packaged devices. In addition, both of these design approaches reduce interface power – relative to two packaged devices – and can eliminate 5W to 10W of power consumption just by avoiding the need for on-board connections between two devices.

For example, a 1 Tbps interface between two packaged chips over PCB traces might typically consume 10 pJ per bit. A similar, chipllet-based connection on an MCM might consume just 1 pJ per transition. That power-consumption level is roughly ten times less. In addition, chipllets deliver much better I/O bandwidth and much lower latency compared to connecting separately packaged ICs over a PCB.

All of these chipllet advantages create completely new usage models for FPGAs in systems. For FPGAs, a chipllet-based approach achieves much tighter integration between the SoC and the FPGA, permitting a **flow-through model** (see page 6) that can eliminate the throughput-killing, ping-pong effect of a memory buffer. The availability of Speedchip chipllets and Speedcore eFPGAs represents the first opportunity for semiconductor companies to add FPGA technology to their chip-level design solutions.

As an example, the following figure illustrates a flow-through architecture for a SmartNIC that might use two or three chipllets. In this design (a TOR switch for example), data from an external Ethernet connection flows through the FPGA chipllet to the NIC chipllet. The FPGA chipllet provides reprogrammable resources that can achieve programmable, wire-speed protocol and exception handling that the NIC chipllet is not equipped to handle. Meanwhile, the CPU on the MCM manages both the FPGA and the NIC chip over PCIe connections. The CPU need not be fast enough to manage the high-speed Ethernet traffic directly.

This flow-through model is a superior design solution for using FPGAs to accelerate certain datapath functions such as cryptography or data compression, where a higher bandwidth connection is required.



**Figure 3: A Chipllet-Based SmartNIC Design Where the FPGA Chipllet Serves as a Bump in the Wire**

## Chiplet Interface Standards

---

The figure above illustrates three chiplets interconnected using two different interface standards – PCIe and Ethernet – to connect the chiplets. Although the use of chiplets (or more accurately, bare die) dates back to at least the 1970s, there are still no interconnect standards for chiplet-based MCM design. To date, bare die have been interconnected using proprietary, usually *ad hoc* interface choices. However, there are plenty of candidates for chiplet interconnect standards.

On the physical-interface level, candidates include:

- Simple parallel connections using single-ended digital or LVDS signaling (as employed by most proprietary approaches).
- Pin multiplexing (for example, the use of TDM for address/data multiplexing).
- Standard SerDes-based connections including PCIe and Ethernet.
- Ultra-short reach (USR) SerDes designed specifically for chiplet-to-chiplet connections, such as the Kandou Glasswing SerDes or the OIF's CEI-56G interface.
- Extra-short reach (XSR) SerDes, which are designed for chiplet-to-chiplet interconnect like USR SerDes, but with the additional ability to drive short PCB traces to permit the chiplet to connect with an optical module.

Except for simple parallel connections, these physical interfaces need higher-layer protocol standards. Again, there are candidates. For example, chiplets might adopt Arm's AXI bus protocol, which was really intended to be used for on-chip interconnect. Because it was not designed for chiplet-to-chiplet interfacing, the AXI bus needs to be augmented with a way to handle transmission errors to be used as a chiplet interconnection protocol. Transmission errors might be handled by adding a forward error correction (FEC) protocol block or by retransmitting the packet received in error. However, these error-handling extensions are not part of the existing AXI specification and would need to be added.

The PCI Express (PCIe) protocol, which appears in [figure above \(see page 4\)](#), is a ubiquitous interface protocol, used in PCs, servers, and myriad embedded designs. Designers understand PCIe and it has been constantly upgraded since it first appeared in 2004. PCIe 5.0, announced by the PCI-SIG in 2017, runs at 32 GTTransfers/sec, which is 64 GBytes/sec in each direction for a 16-lane configuration, so PCIe offers bandwidth to burn. Because it is ubiquitous, PCIe is relatively inexpensive. Consequently, PCIe is certainly a candidate for use as a chiplet-to-chiplet interface at the link layer and transaction layer above alternative physical layers.

Ethernet, which also appears in the figure above, is a fast, standard interface protocol that offers built-in transmission-error handling (similar to PCIe). Both of these protocols are well-understood standards; however, they are heavyweight protocols. Both of these interface protocols introduce a non-trivial amount of latency to the interface connection. In addition, these interfaces consume a significant amount of power and die area, so they are not ideal interconnection protocols for many chiplet-based applications.

The cache-coherent interface for accelerators (CCIX — pronounced "see-six") is a chip-to-chip interface, specifically for hardware acceleration applications. CCIX is designed to interconnect individually packaged chips and is currently built on top of the PCIe physical-layer protocol as a PCIe extension. It might be possible to adopt the CCIX specification for use as a chiplet interconnect by employing another, lighter-weight protocol instead of PCIe. However, a CCIX implementation that does not use PCIe as a physical-layer protocol is not considered a standard at the moment.

Some chiplet applications will require a lightweight protocol that introduces virtually no latency. Such low-latency, chiplet-to-chiplet protocols do not yet exist. However, the Interlaken protocol, invented by Cisco Systems and Cortina Systems in 2006, is one possibility. It is already in use as a chip-to-chip interface standard. It introduces less overhead than PCIe. Interlaken bandwidth scales nicely. Interlaken is flexible and somewhat ubiquitous. However, Interlaken is still considered a heavyweight protocol with far more overhead and complexity than needed for chiplet interconnections.

There are other, proprietary chip-to-chip protocols similar to Interlaken that are in use, but none of them have become an interface standard for chiplets either. Protocols rely on acceptance to become standards and die-to-die interface protocols are still seeking standards.


In a quest to define that standard, several interested companies including Achronix, Netronome, GLOBALFOUNDRIES, Kandou Bus, NXP, Sarcina, and SiFive have formed the Open Domain-Specific Accelerator (ODSA) Workgroup to develop open specifications for a chiplet-specific, standard interface that encompasses a complete protocol stack with an application layer, memory-management layer, link layer, multiple PHY layer interfaces, and a substrate layer. The ODSA Workgroup is leveraging existing industry standards where applicable and is developing new, open IP and specifications where needed.

## System Architectures for FPGA Chiplets

At least two alternative system architectures, flow-through and sidecar, are possible when incorporating FPGA chiplets in an MCM. Although mentioned above, it is worthwhile to discuss these two architectural implementations in more depth because the choice will have a significant impact on system performance depending on the tasks assigned to the FPGA.

### Flow-Through Architecture

A flow-through architecture allows high-speed data to flow through the FPGA chiplet, which processes the data as it flows through the device. The figure below illustrates the use of an FPGA chiplet in a flow-through implementation.



**Figure 4: A Flow-Through MCM Configuration**

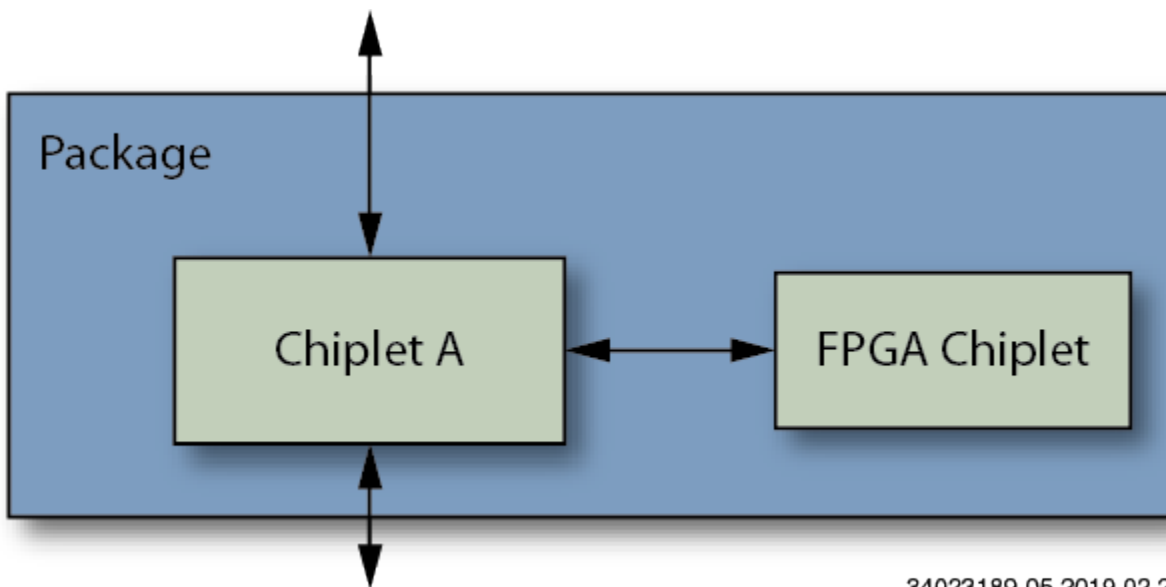
A flow-through architecture is commonly used when the FPGA resides in a data plane and needs to process data at a very high rate. The flow-through architecture reduces data movement between memory and processing resources, which improves performance and reduces power consumption. The flow-through model also supports simultaneous protocol conversion and interface adaptation.

In the example shown in the figure above, the FPGA chiplet has direct connections to package pins. Therefore, an FPGA chiplet used in this manner must either have appropriate I/O structures and capabilities to handle board-level interface requirements, or it must be connected to other I/O chiplets that can provide such facilities.

For example, the FPGA chiplet can be flanked by high-performance ADCs and DACs to implement an analog-digital-analog signal-processing chain for high-performance RF applications. Using multiple chiplets this way decouples the performance and power binning of the analog and digital portions of the RF signal-processing chain, which can significantly reduce costs when dealing with high-performance analog converters.

## Sidecar Architecture

Conversely, a sidecar architectural implementation pairs the FPGA chiplet to one other device in the MCM, as shown in the figure below. This architectural choice is commonly used when the FPGA serves as a hardware accelerator. In such cases, the other chip in the multichip package will likely be a processor.



**Figure 5: A Sidecar MCM Configuration**

In sidecar architectural implementations the FPGA chiplet only communicates with one other chiplet, so it does not require large I/O structures that can handle off-module I/O requirements. One architectural advantage of the sidecar implementation is that the FPGA chiplet can be a delete option for the multichip device, which can be offered with and without hardware acceleration at different price points. Omitting the FPGA chiplet from the MCM reduces the MCMs cost by eliminating the cost of the chiplet die, the incremental assembly cost of adding the chiplet to the package, and the associated test costs.

Hybrid configurations that combine flow-through and sidecar properties are certainly possible as well.

## FPGA Chiplet versus eFPGA IP Design Considerations

---

IC designers have pushed device density for decades due to the constant pressure to pack more functionality into every new generation. With the slowing of Moore's Law, that task becomes progressively harder. Different semiconductor players are adopting different approaches as the difficulty increases. Some large IC vendors are staying monolithic and making huge, expensive efforts to make their largest chips yield. There is significant risk in this approach as some IC vendors have discovered. The delays resulting from the failure of some efforts are painfully public.

Other IC vendors are adopting the chiplet approach for certain devices to reduce manufacturing costs and risks. Quite simply, two or more chiplets may now deliver better economy than one large die. In other words, two chiplets measuring 400 square millimeters will yield better, and therefore, cost less, than one chip measuring 800 square millimeters. Examples of this approach include Intel's EMIB technology, mentioned above, and AMD's new Ryzen Threadripper processors built with chiplets and MCM technology.

Even though the MCM concept is decades old, MCM manufacturing has been an *ad hoc* manufacturing process during most of those years. As a result, chiplet technology has been rather underutilized. Although some semiconductor foundries and outsourced semiconductor assembly and test vendors (OSATs) have started to offer advanced MCM manufacturing capabilities, the business models for industry-wide chiplet use are nascent and are still evolving as discussed above.

The decision to use FPGA chiplets versus the choice to incorporate eFPGA IP in an ASIC or SoC design is not clear cut and there are many considerations to be examined. As a supplier of both eFPGA IP and FPGA chiplets, Achronix has already discussed chiplet technology and its possibilities with many clients, and welcomes the opportunity to have such discussions with other companies interested in being on the forefront of this promising technology.



# Achronix<sup>®</sup>

## Data Acceleration

Achronix Semiconductor Corporation

2903 Bunker Hill Lane  
Santa Clara, CA 95054  
USA

Website: [www.achronix.com](http://www.achronix.com)  
E-mail : [info@achronix.com](mailto:info@achronix.com)

---

Copyright © 2019 Achronix Semiconductor Corporation. All rights reserved. Achronix, Speedcore, Speedster, and ACE are trademarks of Achronix Semiconductor Corporation in the U.S. and/or other countries All other trademarks are the property of their respective owners. All specifications subject to change without notice.

NOTICE of DISCLAIMER: The information given in this document is believed to be accurate and reliable. However, Achronix Semiconductor Corporation does not give any representations or warranties as to the completeness or accuracy of such information and shall have no liability for the use of the information contained herein. Achronix Semiconductor Corporation reserves the right to make changes to this document and the information contained herein at any time and without notice. All Achronix trademarks, registered trademarks, disclaimers and patents are listed at <http://www.achronix.com/legal>.