

Q&A: A Deeper Look at Marvell's MoChi Technology

[Electronic Design](#)

[William Wong](#)

Mon, 2016-06-06 09:09



This past year, [Marvell Technology](#) introduced its MODular CHIP (MoChi) technology as the answer to the costly, complex system-on-chip (SoC) implementation. I talked with Dr. John Shin, Vice President at Marvell, about how the company perceives its LEGO-inspired technology will impact the SoC industry.

Wong: What was the impetus behind the development of Marvell's Modular Chip technology?

Shin: Modern SoCs are very complicated—it often takes more than two iterations to develop an SoC for production. After decades of technology scaling, Marvell realized that the SoC mask set cost was reaching the multi-million dollar mark. At that point, we recognized that even if a company could afford the astronomically high point of entry, the cost of development would eventually take a toll on the bottom line. Also, the ROI of these SoCs in advanced technology nodes (such as FinFET) would be very low, sometimes negative. Additionally, the SoC end market is extremely dynamic with requirements changing all the time.

Related

[Confabbing on the Fabless Fad](#)

[What's the Difference Between SPICE and FastSPICE Circuit Simulators?](#)

[11 Myths About Hardware Emulation](#)

However, as design and process complexity increases drastically, design and debug cycles and time to production get longer, and it is enormously challenging to react. All of these factors led us to explore how to build SoCs differently.

Wong: Can you expand on how the market has changed and how it relates to MoChi?

Shin: Vertically integrated powerhouses, like Apple and Samsung, jumped into the semiconductor market for internal consumption, significantly lowering volume opportunity for the rest of the industry. You may be wondering, what's the big deal? Every industry has its own competition. But the problem with other companies doing SoC business in an already shrunken market is that the industry still competes to make more complicated SoCs just because it's possible, believing that the more we can integrate, the better the chip.

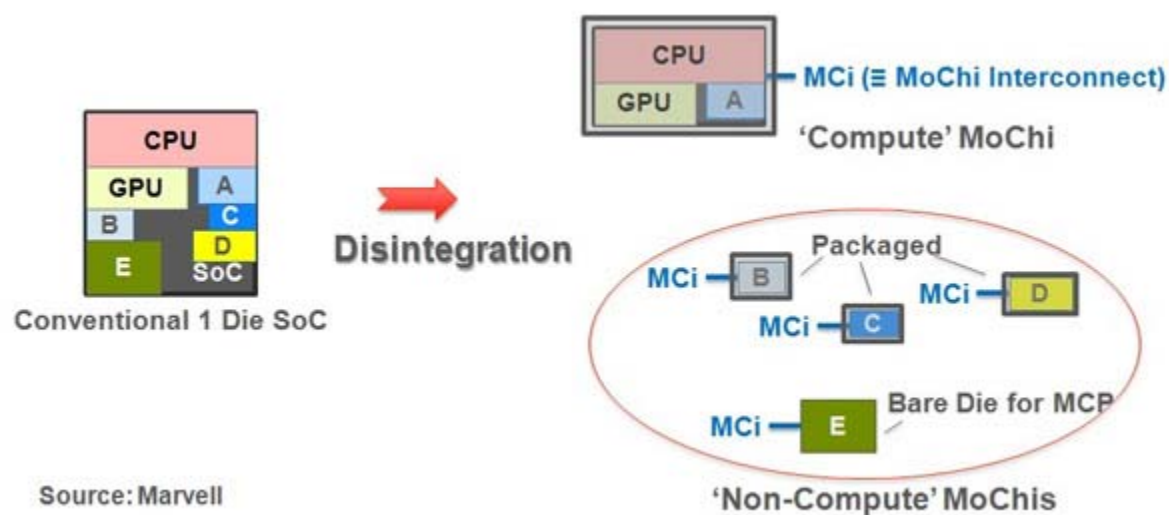
In the smartphone industry, for example, we integrate not just the CPU and GPU functionality, but also the 3G/4G modem, video/image/audio processor, display controller, GPS, and so on. Meanwhile, the smartphone market is constantly changing and, even if we develop the SoC with huge resources ahead of time in anticipation of the market, there is no certainty that the results will come to fruition.

these reasons, we had to shift to a new way of implementing equivalent SoCs, which is more flexible, quicker to respond, and less costly. MoChi technology and architecture is poised to revolutionize the way next-generation SoCs are designed, built, and brought to market.

Wong: How does MoChi technology address these issues?

Shin: In thinking about how to solve this challenge of developing modern SoCs, we needed to accept that the optimal integration is not necessarily building one super SoC with every function required. Rather, it's a handful of simpler and modular chips (MoChis), each optimized and fabricated in the most cost-effective technology.

The Marvell MoChi solution is a radical departure from the current convention of integrating as much as possible onto one IC, helping solve the toll that the cost and complexity of building billions of transistors on a single device is taking on the ecosystem. Much like LEGO blocks that we all likely played with as kids, the MoChi approach focuses on splitting the SoC into cost-optimized modular chips of specific functions that can easily and seamlessly be connected with each other without compromising overall system performance.



MoChis are split into “compute” and “non-compute” categories. Compute MoChis usually require high performance, thus are fabricated in most advanced technology nodes, and non-compute MoChis are typically standard fixed-speed interface functions built with older process generations (*see figure*). With a high-bandwidth, low-latency, point-to-point MoChi Interconnect common to all of these MoChis, this approach provides the ability to connect and combine different MoChis that “talk” to each other as if they are integrated on a single piece of silicon. This leads to the creation of virtual SoCs (VSoCs) in countless configurations with great flexibility. The MoChi Interconnect is essentially a packetized and serialized extension of the widely used on-chip bus, and its area and power overhead is minimal.

Wong: Marvell has also introduced Marvell Final-Level Cache (FLC) technology? How does this fit in?

Shin: On top of the challenge that we identified in designing SoCs, we also looked at existing computer systems that have been slowly evolving. The problem is that even if SoCs were free, the cost of building a computer system is still quite high because of the DRAM cost. That said, if you look at the task manager of your PC, you'll see that most processes are idle, but take up a lot of expensive DRAM space. This gave us a clue on how to solve the problem.

Marvell's FLC innovation introduces a new way to build computer systems that redefines the main memory hierarchy by substantially reducing the amount of expensive DRAM. By offloading the storage onto

;-expensive SSD memory and using a smaller amount of DRAM to cache the active processes, the FLC architecture provides unprecedented reductions in system cost, form factor, operating power, and battery size, or a significant increase in standby battery life. FLC can enable all computers to use a fraction of the energy they use today. This opens up a host of new applications and products, and provides a tremendous cost savings for higher-end devices.

Wong: How will MoChi impact the future of the SoC industry?

Shin: At an industry-level, the Marvell MoChi architecture will enable significant cost savings and flexibility as well as faster development cycles for the entire silicon market from end to end. Minor die overhead of MoChi Interconnect will be easily compensated by the reduction of upfront engineering cost. This comes as a result of the drastic reduction in the number of super SoC combinations that otherwise must be built to serve the various markets and corresponding customer requirements.

To date, Marvell has launched multiple MoChi processors (“compute”) incorporating the FLC technology and MoChi South Bridge products (“non-compute”) across a variety of applications, including networking, mobile computing, multimedia, printer, and storage, among others. The variety of VSoCs using these MoChis demonstrates the key advantages of MoChi, such as flexibly and quickly configuring VSoCs with optimal functions.

The days of the single-die SoC are numbered. The MoChi architecture provides a significant paradigm shift that delivers the return-on-investment required to advance the semiconductor industry.

Looking for parts? Go to [SourceESB](#).

Source URL: <http://electronicdesign.com/digital-ics/qa-deeper-look-marvell-s-mochi-technology>