


Case Study


Push-Button Ease of Arteris NoC IP Freed Up the Team to Focus on Designing a Software-Centric Machine Learning SoC



Customer Overview

SiMa.ai develops system-on-chip (SoC) devices that provide high-performance machine learning (ML) capabilities on the embedded edge where the internet meets the real world.

To address problems with existing computer vision solutions, SiMa.ai has developed and released the world’s first software-centric, purpose-built machine learning system-on-chip (MLSoC™) platform that delivers an astounding 10X better performance per watt than the nearest competitive solution. In addition to running any computer vision application, any network, any model, any framework, any sensor and any resolution, this MLSoC platform provides the developers of computer vision systems with push-button performance.



SiMa.ai MLSoC Evaluation Board

Business Challenges

- Form a new company and grow a world-leading SoC design and verification team.
- Develop a multi-billion-transistor SoC and bring it to market in only four years.

Design Challenges

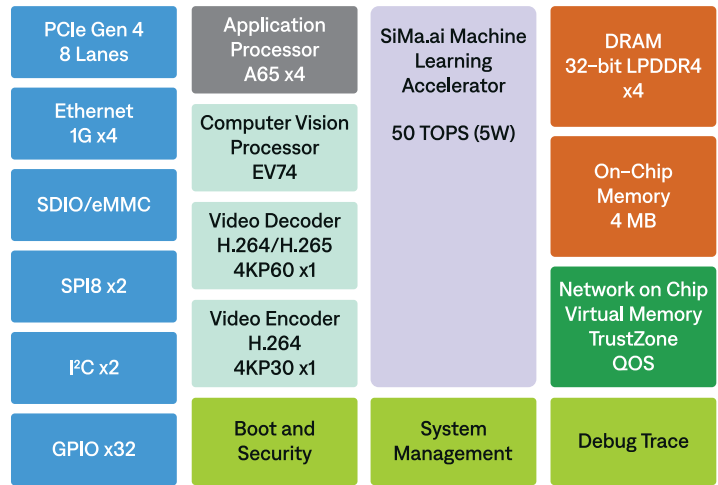
- Design a state-of-the-art machine learning accelerator (MLA).
- Connect large numbers of complex IP blocks with a variety of interfaces from different third-party vendors while ensuring the requisite bandwidths are achieved.

Arteris Solution

- FlexNoC® interconnect IP

Results

- Developed a very complicated SoC in just over three years.
- Generated a NoC with push-button ease leveraging Arteris IP saving years to the project timeline.
- Quickly honed in on an optimal implementation given the ability to explore the solution space with a variety of NoC architectures.
- Delivered the 10X advantage promised to customers.



High-level block diagram of the MLSoC.

The Challenge

The uses for computer vision systems that employ ML to perform tasks like object detection and recognition at the embedded edge without having to transport data to the cloud are boundless. Markets and applications include robots, industry 4.0, drones, autonomous vehicles, medical systems, smart cities and the list goes on. Unfortunately, in addition to consuming a lot of power, traditional solutions involve a steep learning curve that negatively impacts development costs and time-to-market.

Implemented at the 16nm process technology node, SiMa.ai’s MLSoC comprises of billions of transistors. The “secret sauce” that differentiates this device from its competition is the internally developed machine learning accelerator (MLA) intellectual property (IP) block and provides 50 trillion operations per second (TOPS) while consuming a minuscule 5 watts of power.

The many challenges involved in creating this device included designing the MLA IP and combining it with IP from multiple third-party vendors, including on-chip memory, encoders and decoders, peripheral and communications functions, and a quad-core Cortex-A65 application processor (AP) from Arm.

With devices like the MLSoC that involve large numbers of IP blocks from multiple sources, it is necessary to connect them as efficiently as possible without losing out on performance. The only realistic solution for a device of this complexity is to use a network-on-chip (NoC).

There are many aspects involved in building a NoC. Its creators must worry about things like bandwidths, quality of service (QoS) and deadlocks. There is also the fact that different IP blocks may have distinct interfaces (data widths, clock frequencies, protocols, etc.). For example, one block may use 16-bit words of data while another employs 32-bit words; one block may be running at 500 MHz while another is running at 650 MHz; some blocks may use AXI3 or AXI4, while others might use APB, and so on.

Designing and verifying a NoC from the ground up can consume large amounts of time and resources. For example, Srivi Dhruvanarayan, who is VP of hardware engineering at SiMa.ai, was involved in developing such a NoC at his previous company. That project involved six to seven people architecting, designing and verifying the NoC, and this portion of the design took close to two years.

“It was push-button easy to generate the NoC with Arteris. It would have taken months or years without Arteris IP”

Srivi Dhruvanarayan
VP of Hardware Engineering
SiMa.ai

The Solution

At SiMa.ai, the main thing the team wanted to focus on was the design of their MLA. Contrariwise, the last thing they wanted to do was to expend their valuable time and resources developing their own NoC. They needed a powerful, flexible and robust NoC that would allow them to seamlessly connect IP blocks with a variety of interfaces from many different sources. Srivi evaluated the market and, as he says, “quickly determined that FlexNoC from Arteris was the best solution out there.”

Architecting a device as complex as the MLSoC is a non-trivial task. Even when using a NoC, there are many possibilities with respect to how the IPs end up talking to each other. Should everything be presented as a single block? Should there be multiple blocks? Should the design be hierarchical? How should things be split up logically? How should things be organized physically? And, even when the functional design team thinks it has arrived at a solution, what happens when the physical design team comes back and says, “No, we can’t route this!” How does the functional design team set about trying something else? Fortunately, FlexNoC lived up to its name by providing the flexibility SiMa.ai required to allow them to explore the solution space and perform multiple design iterations. The ability to build multiple NoCs in a relatively short time, taking each of them through the entire design flow, was a huge advantage to homing in on an optimal solution.

Another important consideration is the ability to perform performance analysis. It is one thing to connect all of the IP blocks, but it is also important to ensure the requisite bandwidths have been achieved. FlexNoC comes equipped with a performance analysis tool designers can leverage to determine if their various pipes and highways are big enough and fast enough.

“We anticipate ML at the embedded edge is going to be a 25 to 30 billion-dollar market in the next five to six years, so it was important

for us to get things right the first time,” says Srivi. “I have to say that the customer support from Arteris has been excellent. The Arteris support team was always there to hold our hands when we needed them.”

Results and Future Plans

And what was the end result of all this effort? “Let me put it this way,” says Srivi, “ We designed a really complicated SoCs in just over three years; we taped out in February 2022, and we received our first silicon in May 2022, all thanks to our amazing partners like Arm and Arteris. With zero bugs, we continued our march towards first-time-right success by going straight into production. Just six weeks after receiving first silicon, we were working in our customers’ labs running their end-to-end application pipelines and knocking their socks off with the 10X advantage we had promised to deliver. We’ve already started work on our next-generation device, and—with respect to the NoC—we didn’t even think of looking elsewhere because FlexNoC from Arteris was an automatic and obvious choice!”

“We’ve already started work on our next-generation device, and—with respect to the NoC—we didn’t even think of looking elsewhere because FlexNoC from Arteris was an automatic and obvious choice!”

Srivi Dhruvanarayan
VP of Hardware Engineering
SiMa.ai

About Arteris

Arteris is a leading provider of system IP for the acceleration of system-on-chip (SoC) development across today’s electronic systems. Arteris network-on-chip (NoC) interconnect IP and SoC integration technology enable higher product performance with lower power consumption and faster time to market, delivering better SoC economics so its customers can focus on dreaming up what comes next. [Learn more at arteris.com](https://arteris.com).