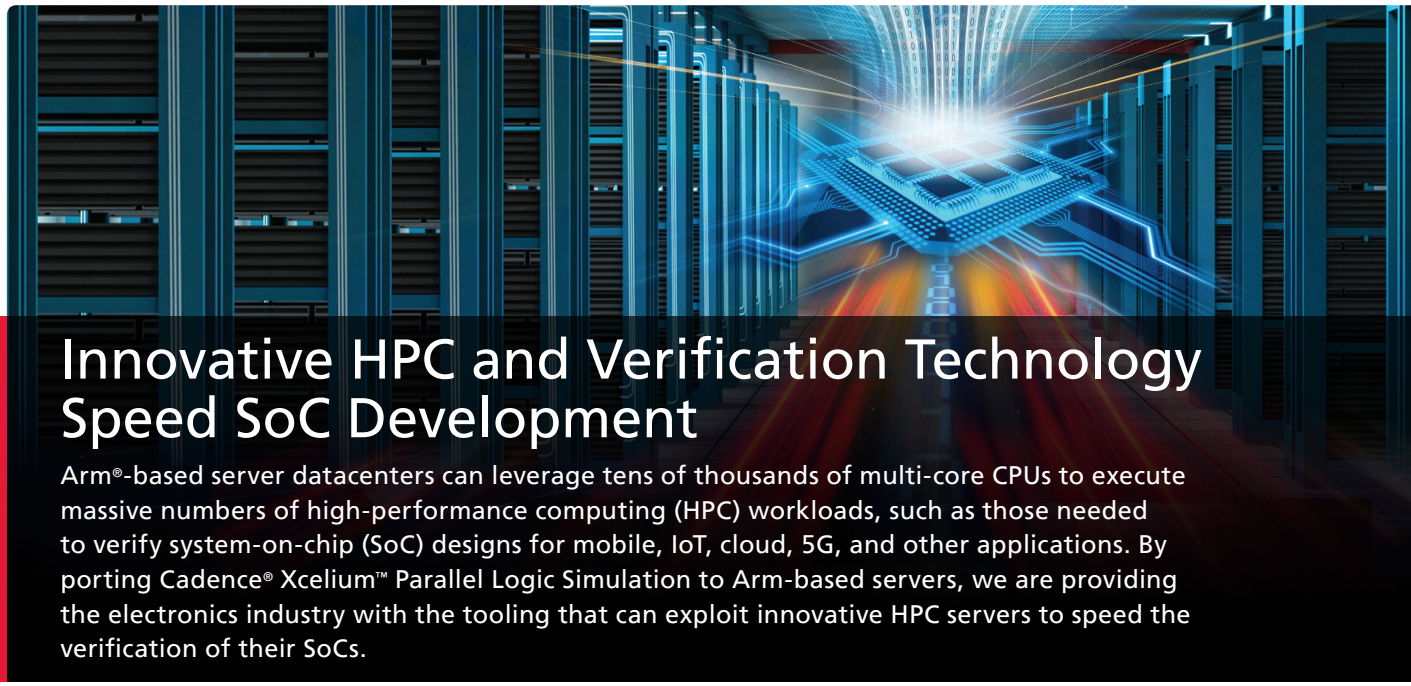

Innovative HPC and Verification Technology Speed SoC Development

Arm®-based server datacenters can leverage tens of thousands of multi-core CPUs to execute massive numbers of high-performance computing (HPC) workloads, such as those needed to verify system-on-chip (SoC) designs for mobile, IoT, cloud, 5G, and other applications. By porting Cadence® Xcelium™ Parallel Logic Simulation to Arm-based servers, we are providing the electronics industry with the tooling that can exploit innovative HPC servers to speed the verification of their SoCs.

Overview

Verifying that SoC designs function correctly prior to manufacturing is a massive task requiring HPC. Accounting for over 70% of the EDA compute workload, SoC verification is a key driver for growth and transformation of the datacenter. Datacenters need energy-efficient platforms optimized for improved performance of a variety of different workloads that can be deployed and managed as cost-effectively as possible. Arm IP empowers a broad and varied range of Armv8-based solutions that leverage a common software ecosystem to deliver compelling compute density, flexibility, and efficiency versus legacy architectures. Arm-based server datacenters can leverage tens of thousands of multi-core CPUs to execute massive numbers of HPC workloads, such as those needed to verify SoCs for mobile, IoT, cloud, 5G, and other applications.

A typical SoC will have a flip-flop or latch for every five to ten combinatorial gates. In a 100 million gate design, that means 10 million bits or 2^10,000,000 potential logic states. Engineers have multiple means to address this problem because verifying every state isn't reasonable, but the most pervasive technology underlying these approaches is logic simulation. By porting Xcelium simulation to Arm-based servers, Cadence is providing the electronics industry with the tooling that can exploit innovative HPC servers to speed the verification of their SoCs.

SoC Verification Evolution

The verification requirements for SoC designs have been growing since the late 1990s, but have become significantly greater with modern applications for mobile, IoT, automotive, medical, industrial automation, and aerospace/defense electronics. Early verification techniques applied relatively few directed tests (hundreds) with each one hard-coded to stimulate specific aspects of a design prior to sampling, lab testing, and re-spinning the design until it was ready for market. This process rapidly became too expensive in terms of project length, non-recurring engineering (NRE) costs, and quality misses. Verification engineers turned to more comprehensive methods involving both thousands to millions of high-throughput randomized tests and hundreds to thousands of long-latency tests. In this model, randomized tests tend to run in seconds to hours but the long-latency tests typically run in days to weeks. For a given SoC, the verification workloads vary across positive (functional mode) testing, negative (failure mode) testing, low-power, analog, and more continuously consuming HPC resources throughout the project.

HPC Workloads in SoC Projects

Figure 1 shows the high-throughput and long-latency workloads changing throughout a project cycle. Early in the project cycle, the number of short-duration workloads increases rapidly. Project teams address this increase by consuming more cores on a given server because each workload typically occupies a single core

as it executes. Users may choose to use a second core to offload less compute-intensive applications from the workload such as data (waveform) dumping. As the SoC design components are integrated into subsystems and the full chip, the number and length of the long-latency tests grow. These tests have traditionally run on a single core even though they consume up to the complete memory available on a given server. Toward the middle of the project, the regression cycle time also grows due to the high number of workloads running and grows again toward the end of the project due to the long runtime of the full SoC workloads. Project teams reduce the mid-project regression cycle time by accessing more servers, provided they have the funding. However, the logic simulators evolving in the electronic design automation (EDA) industry for the past 20 years have not been able to substantially address the long-latency workload runtime.

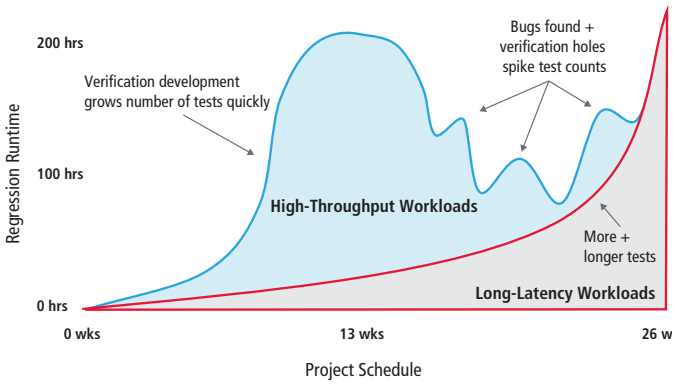

Figure 1: Changing workload use through the project cycle

Applying Xcelium Simulation and Arm-Based Servers

Projects like the one in Figure 1 are fully loading verification datacenters, but would use more HPC resources if they were available. That's where Arm-based servers are ideal. With a greater number of cores in the same rack, datacenters can add workload capacity without affecting the physical space of the datacenter. Moreover, they can do so with reduced power consumption.

Having Xcelium simulation running natively on Arm-based servers further extends the capacity and power benefits by more efficiently executing both high-throughput and long-latency workloads to reduce overall SoC verification time and costs. Launched in February 2017, Xcelium simulation provides both a faster single-core engine, to speed up each workload in the high-throughput test group, and a multi-core engine, to reduce the runtime of long-latency workloads. As a result, Xcelium simulation is ideally suited to the high core count typical of Arm-based servers. For high-throughput tests, the high core count means that projects can run more tests in parallel on a given set of servers, so project teams can achieve better overall SoC quality and faster turnaround time to quality bug fixes. For long-latency tests, where Xcelium simulation is able to scale to all of the cores available on a server, this means 3X to 10X faster regression cycle time in the critical final days and weeks of a project cycle.

Preparing Your SoC Verification Methodology for Arm-Based Servers

Getting ready for verification using Arm-based servers includes both general methodology improvements and specific improvements suited for the high number of available cores. The first step is to audit your current verification methodology. Doing so can identify methods to automate the verification planning and execution process, identify holes in your existing methodology, and assess the current performance requirements for your existing suite of workloads. Applying the audit will immediately improve your SoC verification on any server. If you have not done so, adopting a UVM methodology for comprehensive IP verification will improve SoC quality and creates a large volume of short, high-throughput tests that are well-suited for Arm-based servers. Furthermore, adopting Xcelium Multi-Core simulation for long-latency tests will reduce turnaround time in the critical final phase of your project because the performance scalability of the technology is well suited for the high number of cores available in Arm-based servers. Taken together, this work will enable you to achieve the productivity boost that will come from moving these HPC verification workloads to the combination of Xcelium simulation running on high-core-count Arm-based servers.

For Further Information

Learn more about Xcelium simulation on Arm-based servers at www.cadence.com/go/arm-based-servers

cadence®

Cadence software, hardware and semiconductor IP enable electronic systems and semiconductor companies to create the innovative end products that are transforming the way people live, work, and play. The company's System Design Enablement strategy helps customers develop differentiated products—from chips to boards to systems. **www.cadence.com**