Scale4Edge – Scaling RISC-V for Edge Applications*

Wolfgang Ecker[1], Milos Krstic[2,3], Markus Ulbricht[2], Andreas Mauderer[4], Eyck Jentzsch[5], Andreas Koch[6], Bastian Koppelmann[7], Wolfgang Mueller[7], Babak Sadiye[7], Niklas Bruns[8], Rolf Drechsler[8], Daniel Mueller-Gritschneder[9], Jan Schlamelcher[10], Kim Grüttner[10], Jörg Bormann[11], Wolfgang Kunz[12], Reinhold Heckmann[13], Gerhard Angst[14], Ralf Wimmer[14], Bernd Becker[15], Tobias Faller[15], Paul Palomero Bernardo[16], Oliver Bringmann[16], Johannes Partzsch[17], Christian Mayr[17]

[1]Infineon Technologies AG, [2]IHP –Leibniz Institut für innovative Mikroelektronik, [3]University Potsdam
[4]Robert Bosch GmbH, [5]MINRES Technologies GmbH, [6]Technical University of Darmstadt,
[7]Heinz Nixdorf Institute/Paderborn University, [8]University of Bremen / DFKI GmbH,
[9]Technical University of Munich, [10]German Aerospace Center (DLR), [11]Siemens EDA,
[12]Technische Universität Kaiserslautern, [13]AbsInt Angewandte Informatik GmbH, [14]Concept Engineering GmbH,
[15]Albert-Ludwigs-Universität Freiburg, [16]Universität Tübingen, [17]Technische Universität Dresden

Introduction

The Scale4Edge project is focused on enabling an effective RISC-V ecosystem for optimization of edge applications. The main objective of the project is the development of an ecosystem, based on a platform concept, to supply efficient and cost-effective application-specific edge devices and value-added services addressing different market segments. This is achieved through the automatic and very fine-grained adaptation of highly generic components to the application. The Scale4Edge ecosystem covers highly scalable components and tools and extends them for application-specific edge components at three levels: (1) CPU instruction level defined by the RISC-V Instruction Set Architecture (ISA), (2) software level defined by the C programming language standard C11 with compilers and libraries open to complementary standards like MISRA-C, and (3) operating system and firmware level through system services, configuration interfaces, and drivers.

The ecosystem platform is developed as customizable to the individual application, through the RISC-V ISA, which may define optional custom instructions, e.g., to support specialized computations, for individual applications. As such, Scale4Edge is based on a broadly scalable hardware addressing different pipeline architectures, multi-core architectures, co-processors, and hardware accelerators, e.g., for AI and DSP applications. In terms of hardware scalability, additions for non-functional properties such as energy efficiency, fault tolerance, reliability, safety, and security are important contributions of the Scale4Edge ecosystem. Various measures such as hardening of registers by error correction bits, scaling of clock frequency or protection of memory areas by an MPU are examples. Targeted main applications are classical edge components at the interface of the cloud or fog, respectively, with also specific components, such as sensor interfaces for automobiles and high-reliable (HiRel) electronics for reliability-critical applications.

The Ecosystem

The Scale4Edge ecosystem is composed of a large set of interacting tools to customize, design, verify, and produce application-specific RISC-V based microprocessors. Figure 1 gives an overview of the different components, their interaction and how they are aligned to the different phases of the design and manufacturing process. The CoreDSL language is the key means to define core variants and ISA extensions and allows the generation and customization of hardware blocks and verification/analysis tools. Tools in the ecosystem thus support the hardware/software design step, followed by hardware/software verification and debugging and finally tapeout and chip test. For early software analysis and validation, the ecosystem is based on virtual prototypes, which are applied for ISA specification, simulation, and software development. As such, the ecosystem smoothly complements and extends the industrial synthesis and Place & Route design flow by the following tools from different areas.

High-Level Synthesis

In the Scale4Edge ecosystem, we aim to make RISC-V ISA extensions accessible to non-experts by providing tools

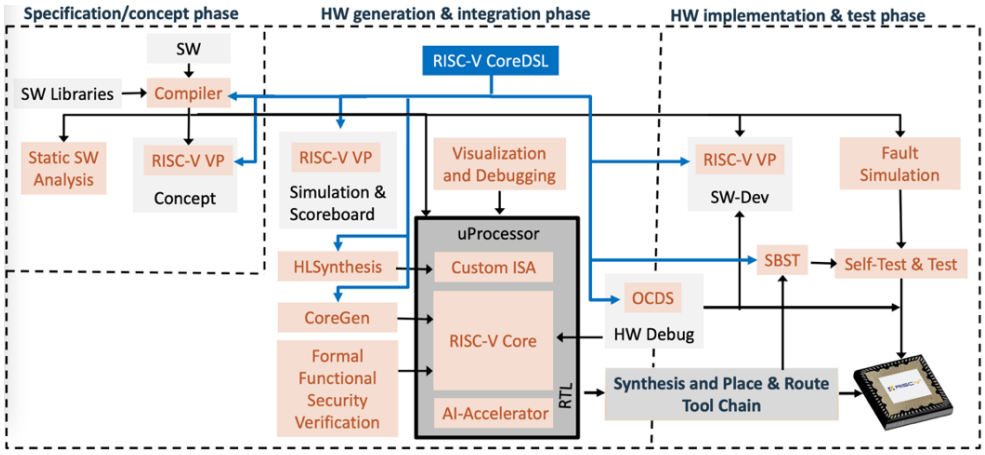

Figure 1: Scale4Edge ecosystem overview

* Funded by the German Federal Ministry of Education and Research (BMBF) under no.16ME0122K-16ME0140+16ME0465.

to automate most of the process using high-level synthesis (HLS) techniques. Specifically, we synthesize custom hardware blocks from the descriptions of instruction behavior in the CoreDSL source file and use the supplied metadata to automatically integrate these into the targeted host core's pipeline.

Software Analysis and Synthesis

Static Software Analysis. AbsInt offers three static analysis tools for inclusion into the Scale4Edge ecosystem. Astrée analyzes C source code with the objective of finding possible runtime errors or proving their absence. StackAnalyzer analyzes binary executables for determining their worst-case stack usage. aiT analyzes binary executables for determining their worst-case execution time.

Compiler Synthesis and SDK Generation. To allow an efficient usage of the RISC-V Instruction Set Architecture Extensions (ISAX), tool support for custom instructions throughout the whole software toolchain (compiler, linker debugger) is required. Based on a CoreDSL description of the ISAX, the project implements a compiler generator to provide the defined custom instructions as assembler code, intrinsic function or as part of an SDK.

Verification, Simulation, and Debugging

Formal Verification for Functional Correctness and Security. The Scale4Edge safety and security assurance rests on two new pillars to address the risk of backdoors and trojans in open-source hardware and software. Siemens EDA automates its formal GapFree verification approach for processor cores[1], which guarantees for a given ISA specification the complete coverage of all relevant processor behaviors. It is complementary with the integration of Unique Program Execution Checking (UPEC), which detects security violations by comparing two instances of the same SoC running any but the same program.

Virtual Prototype Simulation. The project develops two VP-driven approaches for RISC-V software and RTL processor verification. The first approach integrates concolic testing with a VP-based simulation environment (SymEx-VP). The second approach integrates a RISC-V instruction stream generator / fuzzer / symbolic execution engine in a co-simulation setting with an RTL processor under test and an ISS reference model (RISC-V VP).

Fault Coverage and Fault Simulation. Scale4Edge applies chip tests by the execution of compiled C programs. Prior to their application for the test, compiled C programs are analyzed and simulated on a QEMU based VP, which has been extended by fault coverage analysis and fault simulation for different RISC-V extensions.

SBST Generation. The project develops an approach for automatic generation of Software-Based Self-Test (SBST) programs for RISC-V architectures via bounded model checking (BMC). Research comprises end of production and burn-in tests to detect manufacturing defects, online-tests that can be run during idle times of the processor to detect degradation, and advanced fault models (cell-aware testing). Custom visualizations in StarVision PRO accelerate the SBST development.

Visualisation and Debugging. The Scale4Edge Ecosystem includes the powerful visualization and debugging engine StarVision PRO[2] that is adapted to the specific needs of the ecosystem. The hardware debug infrastructure is proven to work with Lauterbach's debug solution.

Applications and Demonstrators

Ecosystem Validation Demonstrators

The Scale4Edge ecosystem is validated by applying it to two different automotive applications. The first application consists of sensor-signal processing 130nm ASICs, where on-chip processors execute typical control and self-test tasks. The second application is an audio event detection system, where audio event signals are processed by ML algorithms running on an on-chip processor. For demonstrating the audio event detection, an ecosystem demonstrator chip with a PLL, an AI accelerator, and an ISA extended core was developed and verified in 2022 with the Scale4Edge ecosystem and standard industrial tools from Cadence and Siemens EDA with 22FDX technology from Global Foundries.

AI Edge Processing Demonstrators

The ecosystem supports a complete end-to-end TinyML flow that transforms a given machine learning model into deployable machine code with VP-VIBES support. TinyML deployment uses and extends two frameworks: (i) *TensorFlow Lite for Microcontrollers (TFLM)* can be used in an automated flow using their interpreter-based approach; (ii) the *TVM* framework is supported by a static code generation approach like the *TFLM* one, and the *TVM Ahead-of-Time (AOT)* backend. The Scale4Edge ecosystem includes two hardware accelerators: the *UltraTrail* and the *SpiNNedge* accelerator. The scalable *UltraTrail* accelerator is developed for application-specific ultralow-power edge AI processing and uses distributed memories to store the features, parameters, and local results. An array of MAC units calculates the convolutional and fully connected layers. Post-processing operations are handled by a separate output unit. The *SpiNNedge* accelerator allows to offload signal processing and recurrent neural network (RNN) workloads from the RISC-V processor with windowing, filtering, and frequency transforms (FFT, DCT) support.

HiRel Demonstrator

To enable dynamicity and reconfiguration of redundant resources for High Reliability (HiRel) applications, the Scale4Edge project employs a RISC-V multi-processor architecture that can be dynamically configured in core redundancy modes, enabling core level error detection and correction. Different operating modes support lifetime extension of the system like high performance, multi-processing, and destress mode. The HiRel demonstrator chip was designed and manufactured in 2022 based on IHP130 technology.

[1] www.onespin.com/solutions/risc-v
[2] www.concept.de/StarVision.html