

The Enzian Research Platform for Systems Software Research

David Cock, Abishek Ramdas, Daniel Schwyn, Michael Giardino, Adam Turowski,
Zhenhao He, Dario Korolija, Gustavo Alonso, Timothy Roscoe
ETH Zurich
Zurich, Switzerland

Enzian is an open-source research computer developed at ETH Zurich, consisting of a server-class processor closely coupled to a large FPGA. There are 14 operational Enzian systems in use at ETH by the Systems Group and external users at other institutions, and the cluster is available for research use.

Enzian was originally developed to address the lack of a suitable platform for systems software research into heterogeneous computing systems, both those directly using FPGAs and those using custom ASICs.

The rapid development and deployment of custom hardware devices (itself an effect of the end of Dennard scaling) poses a challenge to system software researchers, particularly in academia: either they follow the short-term product decisions of hardware vendors, in which case they are artificially constrained by cost-optimized hardware designs which generally target a narrow class of use-cases, or they are reliant on commodity hardware which is less and less like the systems deployed in, e.g., modern datacenters.

Enzian addresses this problem by being a fundamentally *over-engineered* platform that provides great flexibility in research use-cases, and has been shown to provide functionality and performance strictly greater than the convex hull of FPGA accelerator platforms available recently. It can be viewed as a future computing platform design in its own right, or a high-performance emulator for hardware accelerator designs, or a zero-overhead instrumentation platform for servers, or a combination of all three.

A single Enzian board consists of a ThunderX-1 CPU SoC (48 ARMv8-A cores at 2GHz) connected to a Xilinx XCVU9P FPGA by the CPU's native inter-socket cache coherence protocol, essentially forming an asymmetric heterogeneous 2-node NUMA system. The CPU has 128GiB of DDR4 RAM, a pair of 40Gb/s smart Ethernet NICs, and several other on-board accelerators. The FPGA has up to 1TiB of DDR4, plus 400Gb/s of network bandwidth brought out in 4 100Gb/s QSFP-28 cages. Each side also has PCIe and NVMe, and the low-latency cache coherence link between the two nodes has a bandwidth of 30GiB/s. This capacity makes Enzian performance-competitive with commercial PCIe-based

boards, while offering lower latency and more flexible access. The available network bandwidth allows multiple Enzians to be connected together to form a multiprocessor system with custom cache coherence.

Moreover, opening the cache coherence protocol on the FPGA allows use cases beyond that possible with PCIe-based interconnects. Direct access to coherence messages from Verilog or VHDL allows custom memory controller designs to be prototyped (for example, providing controlled persistence), logical views over main memory to be implemented, or individual cache interactions leveraged for message passing using few interconnect round-trips than can be achieved with conventional cache coherence.

Finally, Enzian has a large (and over-engineered) Board Management Controller which governs power and clock sequencing and can provide detailed telemetry information like energy consumption of individual I2C chips on the board. Enzian therefore offers unparalleled visibility into the hidden workings of a modern server board, and exciting opportunities for research into high-assurance firmware, as well as remote access to powerful debugging functionality.

The Enzian cluster at ETH Zurich can be used by external researchers, and consists of 12 machines fully connected on all network ports using 3 "Tofino" programmable network switches. The main CPU runs stock Ubuntu Linux, with a kernel module facilitating interactions with the FPGA. The FPGA runs our custom shell providing a convenient programming interface to the Enzian interconnect and other hardware, and also supports the popular Coyote FPGA "operating system". All software and hardware is available open-source.