

03

Newsletter

October 2026

REBECCA

Reconfigurable Heterogeneous Highly Parallel Processing Platform for safe and secure AI

Website *rebecca-chip.eu*
LinkedIn *@rebecca-kdt-ju*
Twitter @kdt_rebecca
YouTube @RebeccaProject_KDT



Co-funded by the European Union

Chips JU

REBECCA receives funding from KDT JU and national agencies under Grant Agreement No.101097224. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the granting authority. Neither the European Union nor the granting authority can be held responsible for them.

Rebecca Chips Successfully Passed Tapeout

About

Rebecca has passed every signoff check and the netlist is frozen. After months of work, the design is now with the foundry. Rebecca is an edge AI system-on-chip (SoC). It is built as a chiplet, meaning a single die does not have to carry every workload on its own. Two of them will sit in the same package and share the work over a Die2Die link. From a software perspective, this looks like one large system.

COLLABORATIVE CHIP DEVELOPMENT

The design was a joint effort with multiple partners. The IPs were developed independently and integrated together into the Rebecca die. Exapsys developed the RISC-V based processing system and several IPs along with top-level integration. IMEC developed the neuromorphic accelerator and the University of Lund developed Near Memory Processing (NMP) accelerator. Fraunhofer implemented Bunch-of-Wires Die2Die interface and was responsible for the physical implementation with synthesis and Place & Route (PnR).

- Exapsys
- IMEC
- Lund University
- Fraunhofer

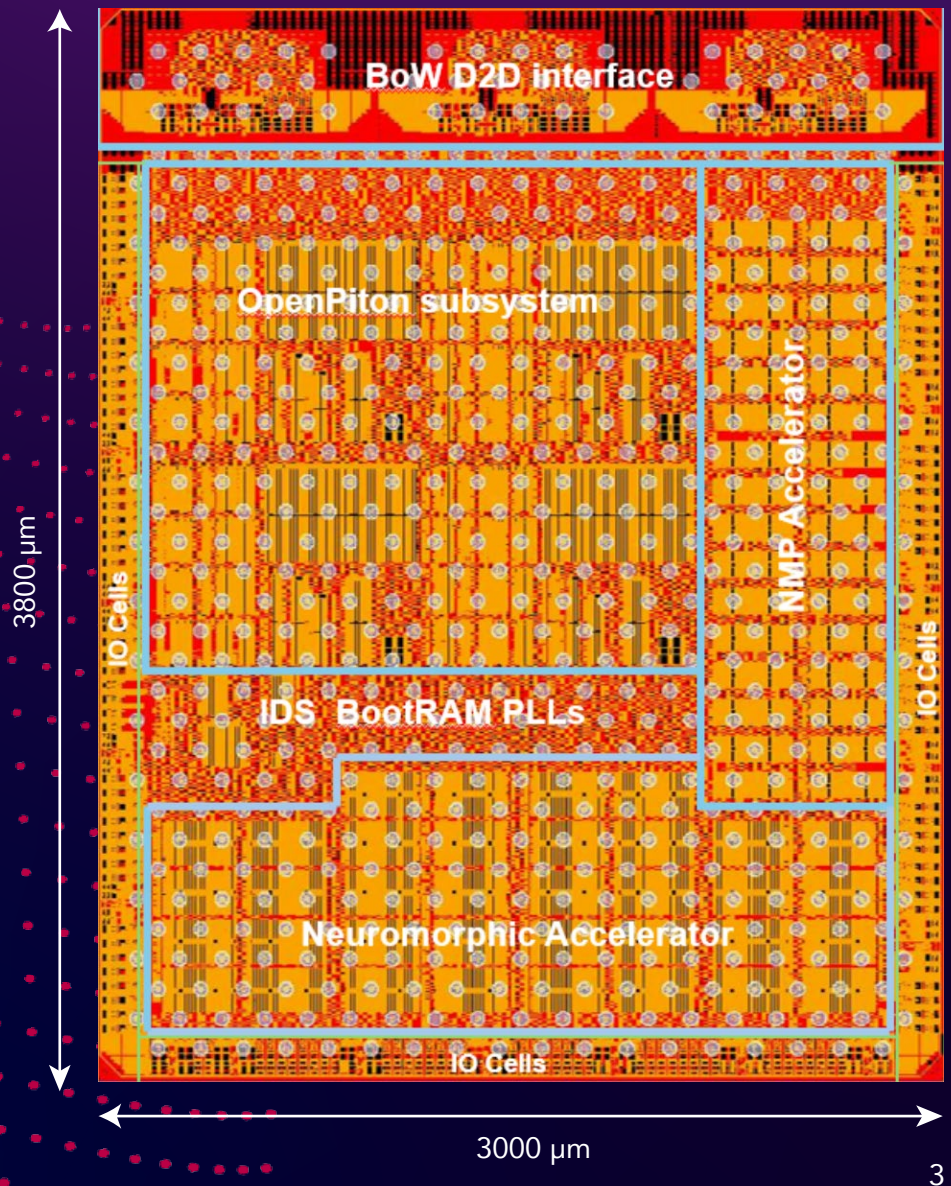
The controller/processing side is a four-core OpenPiton subsystem running at 500 MHz. The cores communicate over an on-chip network and reach the rest of the system through an AXI crossbar. The specialized accelerators sit beside the cores. The neuromorphic accelerator for spiking neural networks which is highly power efficient and reprogrammable for learning on the device. The NMP performs its multiply accumulate operations inside its own memory, which reduces internal data traffic during convolutional neural network inference. Two HyperRAM channels with four memories each running at 150 MHz DDR provide the accelerators the bandwidth they need. Apart from these, the chip has a set of interfaces that includes Die2Die, chip to FPGA, SD card, QSPI flash, two UARTs, GPIO, JTAG and control inputs providing flexibility to interact with the chip.

The Rebecca die is highly memory-dominated. There are 223 SRAM instances across 12 configurations. The chip is built in GlobalFoundries 22FDX, a 22 nm FD-SOI process on a ten-metal stack. The lower metal layers are for standard cells, intermediate layers for routing and the top layer for power distribution with an aluminium layer on top for the pads. The die is 3.0 × 3.8 mm and is assembled flip-chip on copper pillar bumps rather than wire bonds, which keeps the connections short. This configuration spreads the supply current across the whole surface of the die. The figure shows the layout of the chip and provides and overview of where each major IP is located.

Implementation was run through the Cadence tools: Genus for synthesis, Innovus for PnR, Tempus for timing signoff, Conformal for equivalence checking, Xcelium for simulation and Calibre for physical verification. Synthesis turned the RTL into roughly 2.6 million standard cells, which were then placed and routed around the 223 memory macros and the PLLs.

We expect silicon back at the start of next year. The two dies will then be packaged together, connected over the Die2Die link and will then be fitted to the test board that is being finished in the meantime.

The applications for the Rebecca chip can be huge. They include smart appliances, energy generation, infrastructure inspection, avionics, automotive and healthcare domains. Within the project scope, the chip will be used for finding defects in solar panels from a drone, smart fridges, inspecting semiconductor equipment and surveying water infrastructure. Each one runs a vision or AI workload on the device itself rather than sending data somewhere else to be processed and each does so under tight limits on power, latency and safety.



Short News

SYSGO

SYSGO has successfully ported the PikeOS real-time operating system and hypervisor to the REBECCA platform, targeting its quad-core CVA6 RISC-V configuration. The achievement brings together REBECCA's open RISC-V-based hardware with robust virtualization, real-time execution, safety, and security capabilities. Using SYSGO's CODEO development environment and the REBECCA SDK, the team established a complete build and deployment workflow and demonstrated a dual-partition setup running ELinOS alongside a PikeOS native application.

This separation allows demanding AI workloads and critical real-time applications to operate securely on the same hardware while remaining isolated from one another—an essential capability for trustworthy Edge AI. The milestone brings REBECCA another step closer to real-world applications ranging from autonomous infrastructure inspection and photovoltaic panel monitoring to semiconductor manufacturing and intelligent home appliances.

More

FENTISS

- FENTISS has adapted XNG 1.5 (an upgrade from the previous XNG 1.4 adaption) to the REBECCA emulation platform. This new adaptation will streamline the software validation process and aligns with the best practice of keeping systems up to date.
- Building upon the XNG 1.5 port, FENTISS has also successfully adapted the Linux partitions to the REBECCA emulation platform. Leveraging one of these partitions, the necessary adjustments have been made to run SML's EDDL inference library.
- In collaboration with SML, FENTISS has developed a communication protocol designed to transmit images from the water drone to the REBECCA emulation platform. The platform then executes the inference process and relays the results back to the drone.

UCLM

UCLM has defined a set of safety cases that offer a clear explanation of how REBECCA examines and manages potential risks in its system. They help to understand the practical steps taken to ensure that REBECCA system behaves safely, from identifying possible issues to showing how they are addressed.

SML

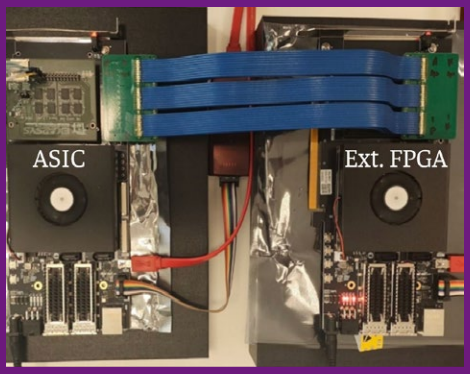
SML has optimized EDDL's to better adapt it to REBECCA's emulation platform hardware requirements, reducing memory consumption and inference times.

In collaboration with FENTISS, SML developed a communication protocol aimed at receiving images from the underwater drone of UC4 into the REBECCA platform. The images will be used as input for the YOLOv8 crack detection models, and the results will be sent back to the drone using the same protocol.

Short News

EXA

EXAPSYS has reached a major milestone in the development of the REBECCA platform, successfully progressing from a fully functional FPGA-based prototype to an ASIC-synthesizable implementation. The FPGA prototype, built using two Alinx AXKU15 boards, provided a complete environment for validating the REBECCA system architecture, including RISC-V CPUs, AXI4 communication, peripherals, memory interfaces, bootloaders, firmware, and operating system bring-up. It also enabled hardware and software developers to test their components on a realistic platform well ahead of silicon availability.



A key part of this work included the validation of the external HyperRAM memory interface using a dedicated FMC module developed by EXAPSYS.

Following successful FPGA validation, the team transitioned the design toward ASIC implementation, replacing FPGA-specific components and refining the RTL through several rounds of simulation, synthesis, and verification.

More

FORTH

FORTH has made significant progress in developing full-custom FPGA-based AI accelerators for the REBECCA platform, targeting AI-on-the-Edge applications for automated crack detection in harbor structures. Working with neural network models developed and trained by REBECCA partners ALM and ASXL, FORTH has developed three accelerator architectures, exploring different CNN models, FPGA platforms, levels of parallelism, memory configurations, and optimization techniques. This work demonstrates how accelerator architectures can be adapted to balance computational performance, resource constraints, and AI model accuracy.

The latest accelerator, based on YOLOv8, is being integrated with the RISC-V software implementation developed by SML. In this setup, the RISC-V processor interacts with the FPGA accelerator layer by layer, demonstrating close hardware–software integration on the REBECCA emulation platform.

Together with POLITO's framework-based approach, FORTH's full-custom accelerators strengthen REBECCA's expertise in FPGA-based Edge AI, providing valuable know-how that can be transferred beyond harbor inspection to a wide range of future AI-at-the-Edge applications.

More

Use case

Underwater Robot for infrastructure inspection with edge AI REBECCA Chip

www.aquasmartxl.com
www.almende.com

Use case: Infrastructure inspection with a surface/underwater robot
One of the use cases of the REBECCA project is a surface/underwater robot supporting both under- and above-water inspection of harbour infrastructure, developed by AquaSmart Engineering (ASXL) and Almende.

Automated inspection for infrastructure damage
The goal of this surface/underwater robot is to inspect harbour infrastructure like bridges, quay walls, sheet pile walls, and pillars efficiently and regularly to avoid maintenance backlogs. Traditional marine inspections often rely on manual visual checks by human divers, which can be time-consuming, costly, and hazardous. To automate and streamline this process, a novel AI crack detection model has been developed. This model analyzes visual feed during navigation to automatically identify structural flaws and generate actionable maintenance insights.

Testing across environments and initial results
The initial version has undergone rigorous testing across multiple stages: in a simulation, on a workbench, inside a controlled test container, and finally in the river Schie in Rotterdam. The testing phase yielded highly promising results and provided valuable real-world data under true field conditions.



AI model for crack detection
For defect detection and segmentation, Almende and AquaSmart Engineering trained a dedicated YOLO (You Only Look Once) detector model. Operating in maritime environments presents unique computer vision challenges: varying water clarity, fluctuating surface reflections, and organic growth on submerged concrete walls. The developed YOLO model is specifically trained to filter out these environment-specific artifacts and isolate structural defects: Defect detection and segmentation: Precision detection of critical structural issues, including cracks, corrosion, and spalling/missing concrete.



Driving smarter harbour maintenance
By combining AquaSmart's robotic platform with the developed AI models, routine inspections can be conducted more frequently and safely. Identifying early signs of concrete degradation allows port operators to schedule targeted preventative repairs long before structural safety is compromised, significantly extending the lifespan of critical harbour assets.

Use case

AI-powered fridges that recognise food with image recognition

Arçelik

www.arcelik.com.tr

Arcelik's CamFridge use case continues with two complementary AI models: a custom MobileNetV2-based Human Detection model for privacy and an Object Detection model for fridge inventory recognition.



Not human **Human**



The Human Detection model is approximately 88 KB with approximately 93% accuracy and has successfully run in partner emulator environments. On Lund's U55C-based REBECCA emulation platform, it achieved approximately 0.7 s inference latency and approximately 26x speed-up compared with CPU-only execution. The Object Detection work has moved from YOLOv7-Tiny toward YOLOv4-Tiny because of operation support, memory, and architectural complexity constraints.

The Human Detection model is approximately 88 KB with approximately 93% accuracy and has successfully run in partner emulator environments. On Lund's U55C-based REBECCA emulation platform, it achieved approximately 0.7 s inference latency and approximately 26x speed-up compared with CPU-only execution. The Object Detection work has moved from YOLOv7-Tiny toward YOLOv4-Tiny because of operation support, memory, and architectural complexity constraints.

An updated YOLOv4-Tiny INT8 TFLite model was shared with Lund. Lund's initial review found the quantization appropriate, while per-channel verification and quantitative accelerator benchmarking remain pending. To improve hardware mapping and performance, Lund recommended replacing model architecture. These changes require retraining and validation, and the accuracy of YOLOv4-Tiny must be compared with YOLOv7-Tiny before the final model is selected.

Klepsydra successfully tested the updated object-detection network in its AI engine using a workaround for the unsupported Transpose operation: the network is cut before Transpose and converted to floating point through a Dequantize layer. Latency and memory benchmarking can now proceed, but final REBECCA chip emulator results are not yet available. The deployment paths have different requirements, with Lund targeting full INT8 quantization and Klepsydra evaluating a UINT8 QOperator/Opset 12 representation, so model adaptation remains necessary.

The last confirmed Synthara result is the successful execution of the Human Detection model in RTL simulation; no new technical update was received during the latest reporting period. The next priorities are to apply and validate the Lund-recommended YOLOv4-Tiny changes, complete Lund's per-channel verification and benchmarks, continue Klepsydra's latency and memory measurements, compare YOLOv4-Tiny with YOLOv7-Tiny.

Past Events

RISC-V Summit Europe 2026
8-12 June, 2026

In June 2026, REBECCA partners: Alexandros Ntyrkai, Vasileios Karakostas, and Dimitris Gizopoulos from the National & Kapodistrian University of Athens participated in the RISC-V Summit and presented the poster - Evaluating the Vulnerability of RISC-CPUs Against Cache Timing Attacks.

Assessing the vulnerability of caches against side-channel attacks is of critical importance when enhanced microarchitectural security is a must-have feature for a multicore CPU implementation. Previous works have proposed various metrics and methodologies to assess such vulnerabilities. However, those works suffer from limitations regarding either the range of target cache attacks, the support for the RISC-V ISA, or rely on in-house simulators.

The goal of this research is to provide support for systematically evaluating RISC-V multicore CPUs against a wide range of cache timing attacks. This support should allow the assessment of both real and simulated systems, enabling early security evaluation in the design phase of the processor using the open-source gem5 microarchitectural simulator.





The 52th European Solid-State Electronic Research Conference: ESSERC 2026

7-10 September, 2026







In September 2026, Joachim Rodrigues (Lund University, FAU Erlangen-Nürnberg, System-on-Chip at FAU), Sergio Castillo Mohedano (Lund University), and Arturo Prieto (Lund University) participated in the ESSERC 2026 Conference, where they presented two research papers supported by the REBECCA KDT JU project:

- Scalable and Power Efficient Approximate Multipliers for Error-Resilient Applications in 22 nm
- LUCIA: A Scalable and Versatile 0.82 TOPS/W/mm² Chiplet Architecture for CNN in 22 nm FDSOI.

The team also presented measurement results from a chip incorporating IP developed within the REBECCA project, marking an important step from research towards validation. These contributions showcase REBECCA's progress in developing power-efficient, scalable, and high-performance computing technologies for next-generation Edge AI.

REBECCA at EEAI2026

19 to 21 October, 2026

From 19 to 21 October 2026, the REBECCA project participated in the European Conference on Edge AI Technologies and Applications (EEAI 2026), bringing together leading experts from academia, research, and industry to explore the latest developments and future directions in Edge Artificial Intelligence.



This year's conference provided an excellent opportunity to present REBECCA's progress, exchange knowledge, and engage with the European Edge AI community.

Iakovos Mavroidis, REBECCA Project Coordinator, presented the project, highlighting its progress, key achievements, and contributions to advancing Edge AI technologies. Matthias Rosenthal (ZHAW) moderated the session "Adaptive Allocatiing approaches. Christos Kozantis (FORTH) delivered two research presentations. These contributions highlighted the expertise of REBECCA's partners and their active involvement in advancing research and innovation across the Edge AI ecosystem.

EEAI is an important European forum for exchanging knowledge, showcasing innovative research, and fostering collaboration between academia and industry. Covering topics from hardware and software technologies to emerging developments in generative and agentic AI, the conference offered valuable insights into the evolving Edge AI landscape and opportunities to connect with researchers and innovators shaping its future.



