02
Newsletter

October 2025

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 at a Glance

October 2025

About

The REBECCA project develops a RISC-V-based ASIC with integrated AI and security accelerators, for advanced edge-AI systems. Targeted at critical applications like automotive, healthcare, and smart cities, our ASIC connects to an external FPGA with application-specific AI accelerators and I/O, enhancing flexibility for complex, scalable AI tasks.

TECHNICAL GOALS

These goals aim to build a versatile, high-performance edge-AI platform.

- Scalable RISC-V ASIC;
- FPGA for Flexibility;
- Near-Memory Processing;
- Hardened Security Cores;
- Virtualized Software Stack;
- Proven Real-World Use.

Emmulation Flow

ASIC EMULATION

- A proper Linux image corresponding to the new system has been implemented
- Successfully booted Linux (first time CVA6 on ALINX) using 512MB HyperRAM as main memory
- Hardware acceleration: Intrusion Detection System (IDS), Near-Memory Processing,
- Neuromorphic Acceleration

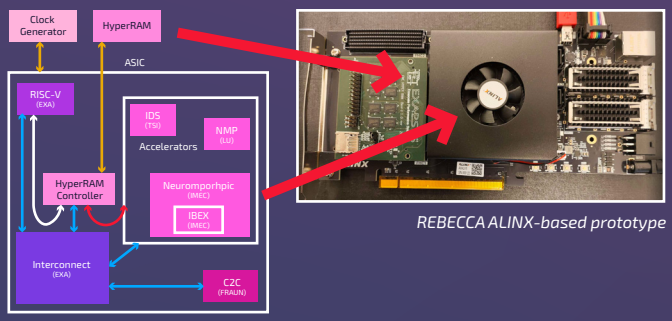

REBECCA ALINX-based prototype

Global Shared Address Space

HARDWARE GLOBAL ROUTING

- 512 MB HyperRAM
- Chiplet-to-Chiplet
- Shared memory model
- IO Coherency

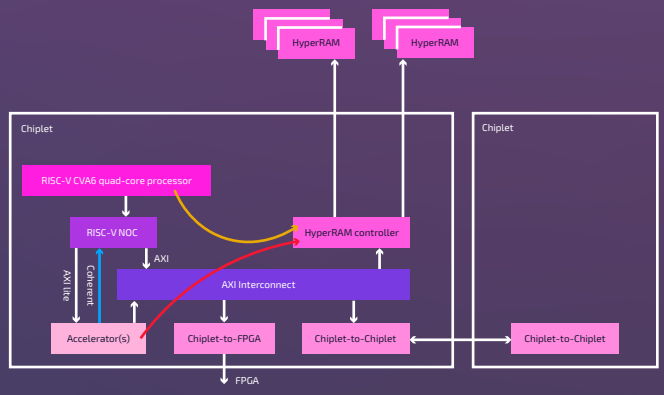

Software Development

SW ON ALINX

- A proper Linux image corresponding to the new system has been implemented
- Successfully booted Linux (first time CVA6 on ALINX) using 512MB HyperRAM as main memory
- IDS device driver (user space and uio driver)

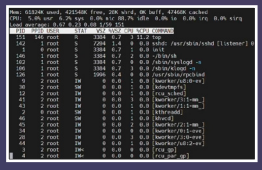

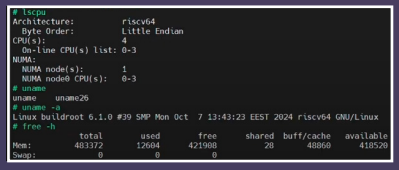

Short News

SML

EDDL library has been successfully deployed into REBECCA emulation platform. The deployed library was capable of performing inference of complex object recognition AI models, such as YOLOv8, a requirement for some of REBECCA Use Cases.

NKUA

The University of Athens team has progressed with the development of a novel full-system safety/reliability assessment simulation-based framework for heterogeneous SoCs which employs fault-injection to quantify the vulnerability of the RISC-V CPU and the REBECCA accelerators to multiple silicon defects (including transient, permanent, and marginal delay hardware faults) when the platform is deployed at the edge. Systolic-array based and other types of accelerators can be analyzed. Total failure rates as well as silent data corruption (SDC) rates can be measure for large scale deployments.

FENTISS

- FENTISS has finished the porting of his XtratuM hypervisor for the RISC-V architecture (without H extension support) to the gem5 RISC-V simulator.
- FENTISS, in collaboration with EXA, have finished the porting of his XtratuM hypervisor to the REBECCA emulator platform based on the CVA6 processor (RISC-V architecture, without H extension support).
- FENTISS, in collaboration with FORTH, has implemented a first approach to resolve the orchestration with dynamic resource allocation using Kubernetes. A new approach with new ideas to be developed in the forthcoming months.
- FENTISS, in collaboration with SML, has successfully tested the EDDL library in the XtratuM hypervisor for the RISC-V architecture (executed on the QEMU-virt emulator).

CSEM

Good progress has been made on CSEM's custom CNN accelerator. CSEM has completed the hardware integration and on-system validation of the accelerator on the prototype FPGA, successfully testing individual neural network layers. The final step of running a full neural network model on the hardware is now underway, in collaboration with use-case partner Intecs in particular.

SYSGO

We have defined an FGPA-based emulator of the REBECCA platform using the OpenPiton system with dual-core configuration. For this setup, we have created bitstreams and a PikeOS PSP (platform support package) supporting the FPGA-based emulator.

Real-time defect detection in PV panels on unmanned aerial vehicles (UAV) devices

www.intecs.it

The Real-time Fault Detection in PV Panels on UAVs use case, led by Intecs within the REBECCA project, aims to automate the inspection of photovoltaic plants using aerial images captured by infrared (IR) and visible (VIS) cameras. The core objective is to process data directly on-board the UAV, enabling immediate detection of defects and reducing the need for post-flight analysis, which in turn lowers O&M costs and accelerates maintenance actions. Over the last months, Intecs has continued the development of the real-time defect detection pipeline with a focus on improving efficiency and robustness for UAV deployment. A key milestone has been the adoption of YOLOX-Tiny, a modern lightweight deep learning architecture specifically designed for real-time applications and edge deployment. Compared to the previously used YOLO-S model, YOLOX-Tiny achieves higher accuracy with a lower parameter count and significantly reduced computational complexity, making it more suitable for continuous on-board inference during flight.

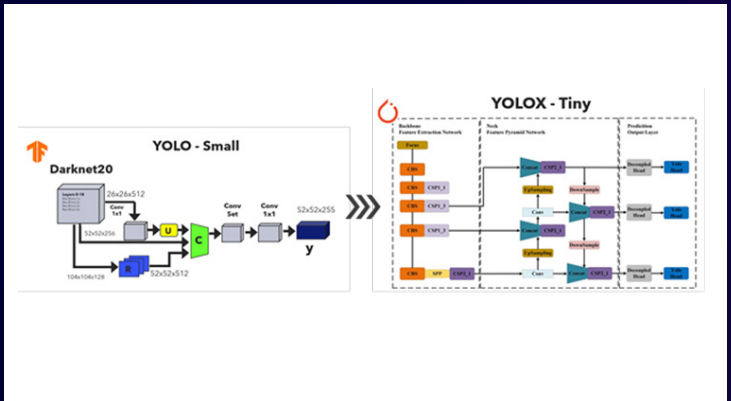

To ensure compatibility with constrained hardware platforms, the model has been quantized to INT8, while preserving accuracy and inference speed. Initial results show that YOLOX-Tiny maintains stable performance before and after quantization, achieving accuracy values around 90% mAP@0.5 with inference speeds above 1 FPS, thus fulfilling the baseline KPIs of the use case. Beyond hotspot detection in IR images, the use case also addresses anomaly detection from visible (VIS) images and the precise localization of defects on the ground.

While the first two stages are currently operational and under evaluation, the third stage — integrating AI outputs with GNSS and UAV telemetry data to compute defect positions with sub-meter accuracy — is still in development. Progress has been slowed by issues in the data acquisition pipeline, but solutions are being actively explored to enable reliable testing in the coming phases.

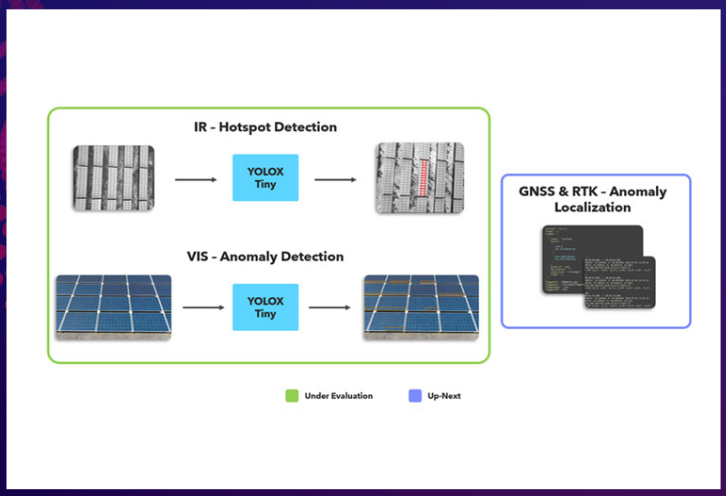

Finally, collaboration with REBECCA's subcomponent providers is continuing in parallel. Partners such as POLITO, Klepsydra, CSEM, Synthara and Sysgo are advancing work on hypervisors, frameworks, and hardware accelerators that will be integrated with Intecs' UAV inspection pipeline. These contributions will be fundamental in ensuring that the full solution meets real-time constraints and can be effectively deployed in the field.

AI-powered fridges that recognise food with image recognition

www.arcelik.com.tr

Arçelik's AI-powered fridge advances along two complementary tracks: Human Classification for privacy (binary classifier with a MobileNetV2 backbone) and Fridge Inventory for item recognition (YOLOv7). This update consolidates the model scope, current maturity, and collaborative testing plans under WP5. On deployment readiness, Arçelik reviewed quantization paths across PyTorch and TensorFlow/TFLite, noting backend trade-offs (e.g., FBGEMM/QNNPACK vs. TFLite symmetric/asymmetric schemes). Framework selection is guided by target hardware. The inference results of the quantized models executed on the desktop system are presented below.

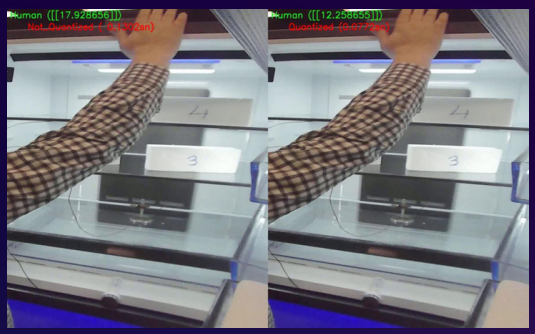

We are collaborating with three partners, Klepsydra, Lund University, and Synthara, across three acceleration scenarios: (1) Klepsydra's AI inference engine, (2) Lund's Near-Memory Processing (NMP) accelerator, and (3) Synthara's FPGA/AI accelerator.

Synthara

Synthara path, the Human Detection model has successfully run in RTL simulation; next steps are mapping the remaining pipeline (incl. YOLOv7) and aligning on a full use-case demo plan.

ULUND

In collaboration with Lund University (ULUND) on the Near-Memory Processing (NMP) accelerator, the Human Detection model was simplified (depthwise / pointwise replaced by a single Conv2D, ~88 KB) while maintaining ~93% accuracy. On a 50 MHz REBECCA emulation, the NMP accelerator delivered ~26× speed-up and ~0.7 s latency; with the ASIC expected at ~500 MHz, the cumulative improvement could approach ~100× vs. the original baseline. Given effort/benefit, the near-term focus remains on Human Detection rather than refactoring YOLOv7.

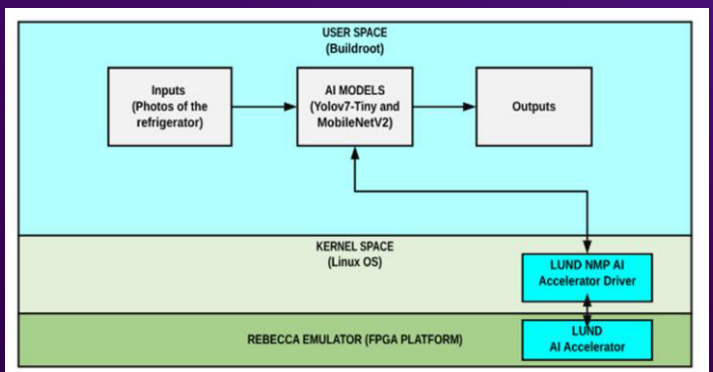

Arcelik's Lund NMP AI Accelerator Based Scenario

Klepsydra

With Klepsydra, Arçelik provided Human Detection models in TFLite/ONNX, FP32 and INT8, already supported by the Klepsydra engine; initial benchmarking is promising but end-to-end UX is not yet production-grade due to image transfer time. YOLOv7-Tiny is supported by Klepsydra and runs on some RISC-V boards (e.g., PolarFire Icicle) though it does not fit the REBECCA emulator's memory budget.

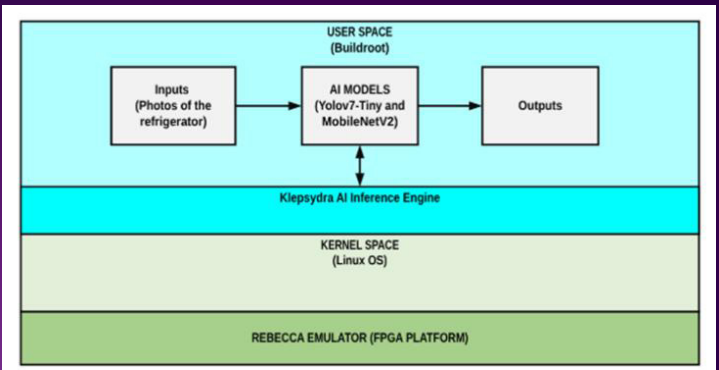

Arcelik's Klepsydra AI Inference Engine Based Scenario

Use case

October 2025

High-speed damage inspection in semiconductor equipment

www.itecequipment.com

ITEC pick-and-place equipment (ADAT 3 XF)

ITEC assembly equipment is used in semiconductor manufacturing to perform pick-and-place operations, handling extremely small chips of tens of micrometers up to medium sized chips of several millimeters. To achieve an accurate placement (1~5µm) the chip size and position need to be determined in real time. Currently, a frame-based camera is used which requires a motion-stop, limiting the machine speed to ~10 chips/second. The next generation ITEC machines will be designed at much higher speeds (>100 chips/second). To eliminate the motion-stop ITEC is investigating the use of an event-based camera.

ASPECT	Frame-based Full frames at fixed intervals 		Event-based Changes-per-pixel in real time 	
Latency	High (limited by frame rate)	--	Extremely low	++
Motion blur	Common in fast motion	--	Virtually non	++
Data volume	Large (includes redundant info)	--	Sparse (only relevant changes)	++
Frame rate	Limited (typically 10-200 fps)	--	Very high (up to 10.000 fps)	++
Technical Readiness Level	Broad, mature ecosystem	++	Emerging, requires new algorithms	--

Pros and cons of different camera types

The high-level goal for ITEC is to track the position of the object (chip). In the first step, the vision information (event-based) is calibrated with the motion of the hardware (encoder information). More steps are needed in a full edge-AI implementation on the Rebecca chip.

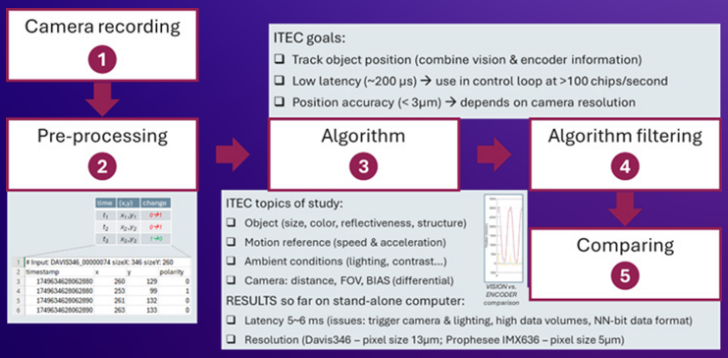

Multiple steps are needed to implement event-based position inspection on the Rebecca chip

In this use case implementation, the event-based camera selection, the ambient conditions and the hardware setup are all of great importance, impacting the data volume, latency and position accuracy. Data pre-processing and algorithm execution have been successfully deployed on a stand-alone computer. In collaboration with Rebecca partners, these data-processing steps will be transferred to the Rebecca emulation environment. Different algorithms will be evaluated to run the position calculation – classical algorithms whether or not in combination with Rebecca accelerators (e.g., Klepsydra, Sysgo) and possibly neuromorphic network algorithms (IMEC).

Past Events

EEAI 2024
21-23 of October, 2024

On the 21-23 of October, 2024, the REBECCA project was represented in the European Conference on Edge AI Technologies and Applications, in Cagliari, Sardinia, Italy. At the conference, the project coordinator, Iakovos Mavroidis, and Konstantinos Georgopoulos delivered a presentation, "A reconfigurable-based Heterogenous and Highly Parallel Processing Platform for safe and secure AI: RISC-V subsystem". The integration of AI/ML components into computing solutions is advancing rapidly while unlocking capabilities that were previously unattainable. However, this progress pushes the performance demands on computing devices, often necessitating the offloading of all but the simplest AI tasks onto cloud-based processing. In critical domains such as automotive and healthcare, relying on the cloud is often unreliable or impractical, driving the need for AI computation directly at the edge. To address this challenge, we are developing an innovative System-on-Chip (SoC) designed for edge AI applications.

This solution incorporates a CVA6 RISC-V-based processor subsystem, which utilizes 512 MB of HyperRAM as its main memory and an SD card for storage. The processor subsystem is tightly connected to advanced AI and security accelerators, enabling direct memory sharing between the accelerators and the processor. This architecture is designed to deliver an edge AI solution with significantly improved performance, energy efficiency, and security compared to existing systems. Additionally, the support for reconfigurable hardware allows the system to adapt to specific application requirements, extending its functionality without compromising energy efficiency. Alongside these hardware advancements, we will also develop the necessary software infrastructure, including operating systems, hypervisors, and libraries, to fully exploit the hardware's capabilities.

EFECS 2024
5-6 December, 2024

On December 5th and 6th, 2024, the REBECCA project was presented at the European Forum for Electronic Components and Systems (EFECS) conference and exhibition in Ghent. This year, the conference gathered over 800 participants and offered an excellent platform to explore Europe's technological and economic priorities through insightful presentations, strategic discussions, and valuable networking opportunities. The REBECCA KDT JU project team highlighted cutting-edge developments in RISC-V-based systems, chiplet technology, and specialised tools for safety and security, sparking engaging discussions and valuable insights.

RISC-V Summit Europe 2025
12-15 May 2025

On 12-15 May 2025, the REBECCA project was represented in the prestigious RISC-V Summit Europe 2025, held in Paris. This high-profile event brought together experts and innovators from across the globe to explore the future of open-standard computing and its applications across various technology domains. Representing REBECCA at the summit were Iakovos Mavroidis, Konstantinos Georgopoulos, and Joachim Rodrigues, who skillfully led the poster sessions. Their participation helped showcase the project's work and opened avenues for engaging discussions with the broader RISC-V and edge AI communities. The REBECCA project is pioneering advancements in edge AI systems using RISC-V technology, emphasising power efficiency, scalability, and open-source accessibility. It integrates a multicore RISC-V-based architecture with AI-specific accelerators, neuromorphic computing, and security features to deliver a high-performance, cost effective AI platform. The core of REBECCA is the CVA6 processor, leveraging a chiplet-based design and shared memory architecture to optimize real-time AI processing. The platform incorporates HyperRAM for efficient data access and a custom software stack to maximize efficiency and security. Initial prototypes using U55C development boards and FPGA-based testing validate the feasibility of RISC-V for AI-driven applications.

Future research will enhance neuromorphic computing, AI framework integration, and real-time performance optimization. With strong industry and academic collaboration, REBECCA is shaping the future of AI at the edge, positioning RISC-V as a compelling alternative to proprietary AI solutions.

EdgeAI Academy Summer School 2025
7-8 July, 2025

Full lecture material is available **online**

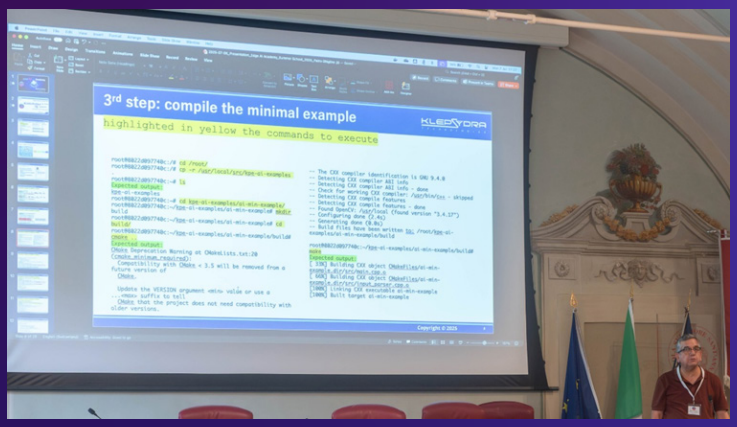

On July 7-8, 2025, REBECCA project participated in the EdgeAI Academy Summer School held in Pisa, Italy. The event brought together leading researchers, innovators, and projects in the Edge AI ecosystem to share knowledge and foster collaboration. At the summer school, our partner, Pablo Ghiglino from Klepsydra Technologies, delivered the course:"Deploying DNNs in RISC-V with Klepsydra AI: A step-by-step, end-to-end workshop." This hands-on session guided participants through the process of building AI applications using Klepsydra AI, simplifying the journey from model to deployment. Attendees gained practical experience in deploying, testing, and optimizing a pre-trained Deep Neural Network on target boards—showcasing Klepsydra AI's versatile and efficient inference capabilities.

The Summer School was organized by Chips JU EdgeAI in collaboration with multiple projects, including: CLEVER-project, REBECCA KDT JU, TRISTAN, NeuroKit2E, SMARTY Chips JU Project, dAIEDGE, SmartEdgeProject (HE). This collective effort emphasised the importance of cross-project synergies in advancing the state of Edge AI research and applications in Europe.

SafeComp 2025
9 September, 2025

On the 9th of September, 2025, Klepsydra took part in the 12th International Workshop on Next Generation of System Assurance Approaches for Critical Systems SASSUR 2025, part of the international SafeComp 2025 conference held in Stockholm. During the presentation, Klepsydra shared progress in two key areas of their work:

- Making artificial intelligence run ultra-fast on multi-core RISC-V processors including REBECCA's;
- Advancing towards the safety certification of our AI inference software.

This reflects REBECCA project mission to bring next-generation AI into real-world applications, combining top performance with absolute reliability.