Real Time Additions to the CVA6 RISC-V core

C. Allioux, O. Betschi, R. Hardy JC. Kircher, G. Miet, I. Schmid
N. Tribie

Robert Bosch Mobility Electronics, Engineering Integrated Circuits

The **CVA6** is an open-source RISC-V core, featuring a 6 stages pipeline, initially developed at ETH Zurich, and now maintained by **OpenHWGroup**. Its architecture, and future improvements deliver enough performance to support **automotive ADAS** (Advanced driver assistance system) applications, which makes it a good candidate to be integrated in our future designs.

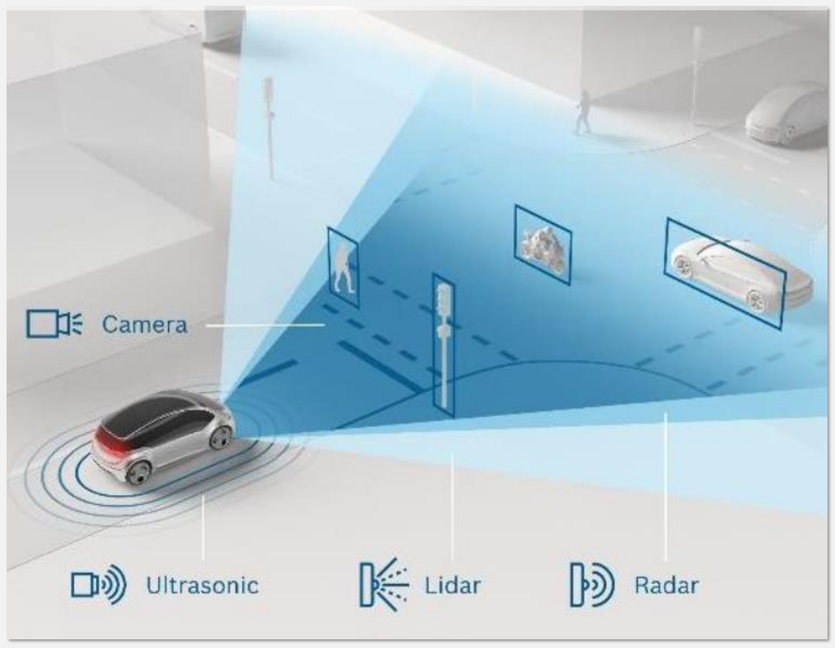

ADAS application:
Vehicles classification, using a CNN accelerator driven by a CVA6 core

However, it lacks some features dedicated to **real time applications,** like scratchpads (tightly coupled memories), or a dedicated peripheral bus. We added them to make the CVA6 more suitable for these workloads,.

We are validating this concept on a demonstrator, a System On Chip (Soc) dedicated to Convolutional Neural Networks (**CNNs**) inference, for road vehicles classification.

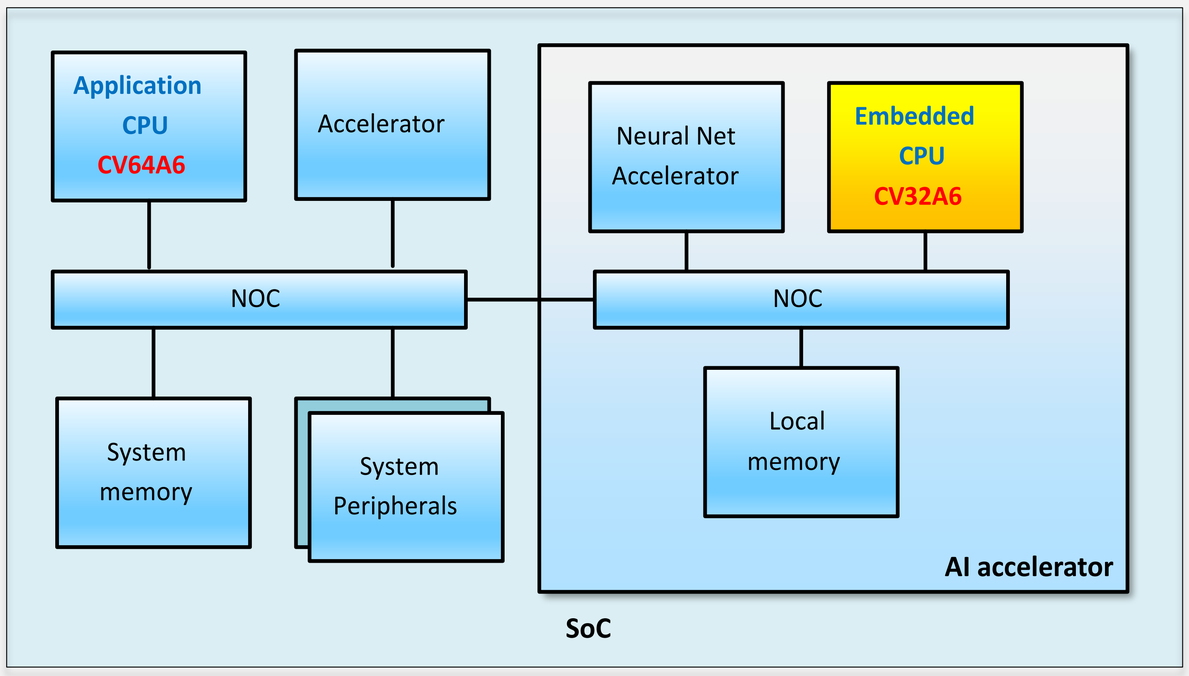

The highlighted CVA6 “embedded” handles the real time tasks : CNN accelerator control, non-linear operations, functional safety and interface with high level application.

Its internal architecture is shown on the next figure :

- **Instructions scratchpad** store the time critical routines: interrupt service routines, DSP-related code.
- **Data scratchpad** : holds the stack, DSP samples data, critical variables.

PMP is moved outside of the **MMU** : Keeping PMP only still provide memory protection, without the timing uncertainties associated with the address translation process (PTW).

- The instruction scratchpad allows to **replace the boot ROM:** instead**,** real-time code is preloaded by the application processor at startup.
- **Peripherals**: are connected to a new dedicated **AHB** interface: The main interconnect to the memory is simplified; Data/code side is split from the control side, for the benefit of functional safety.

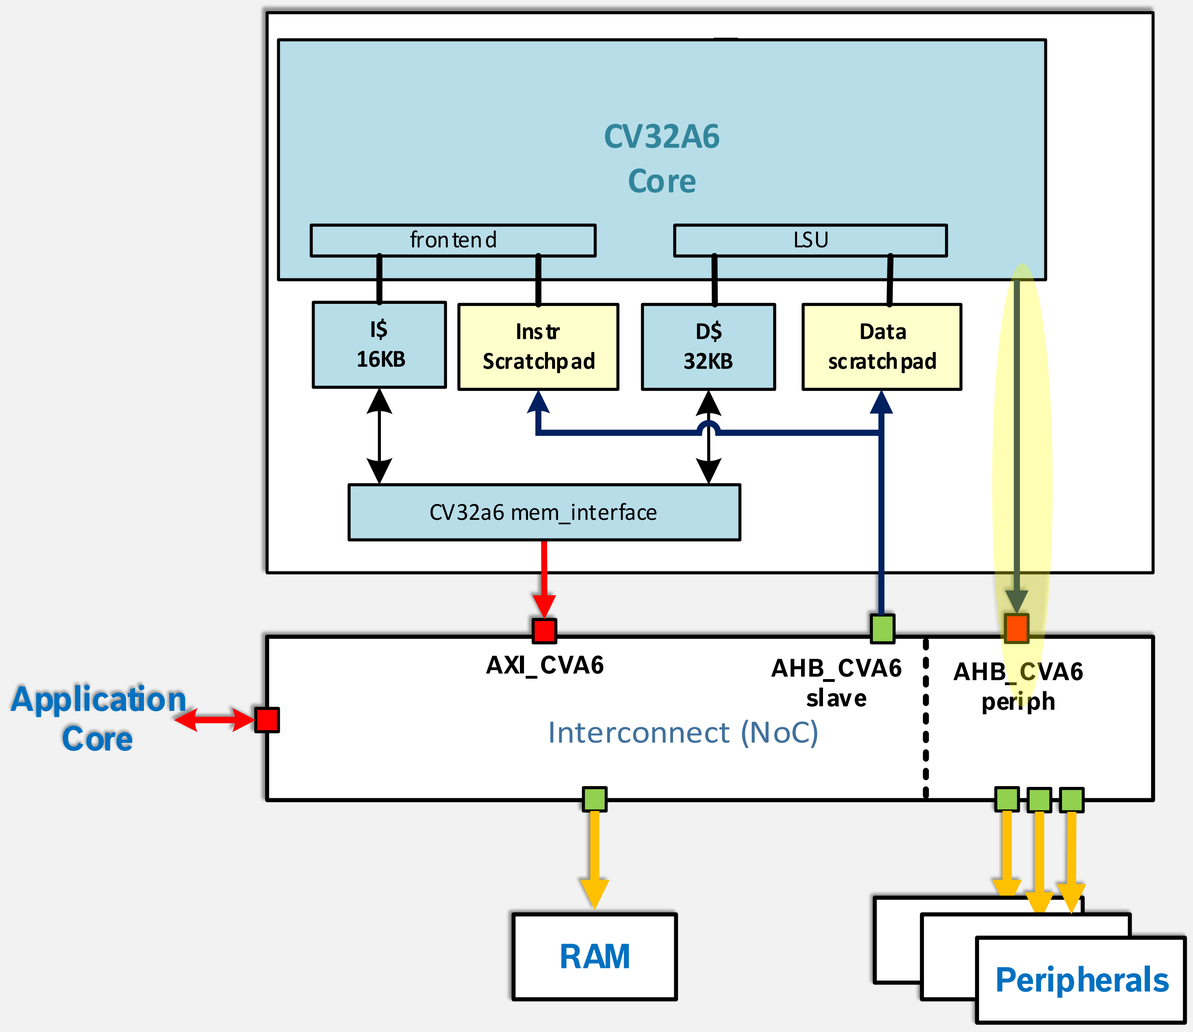

On the SPMV benchmark (matrix-vector product), this architecture reduced the execution loop timing jitter **from 7% to 0.3 %:**

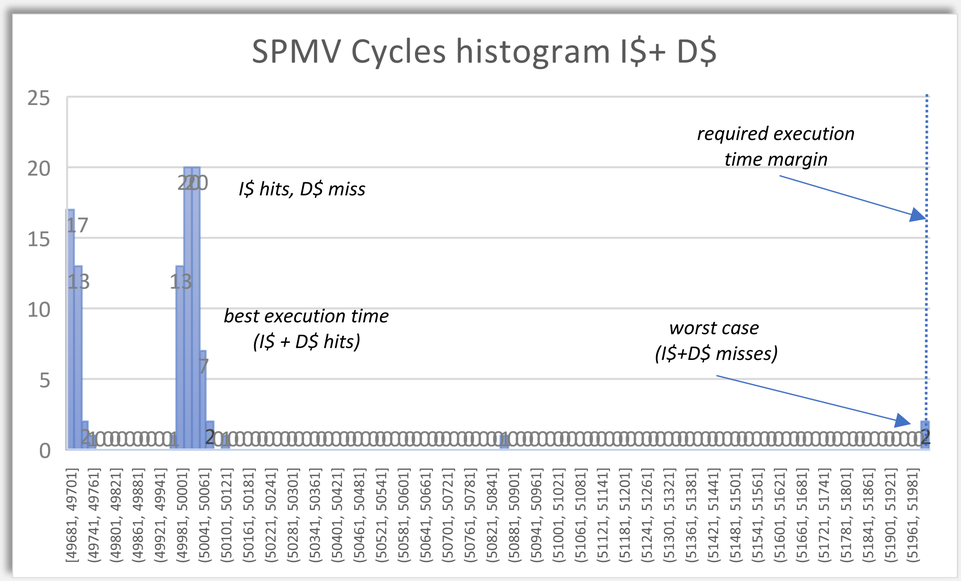

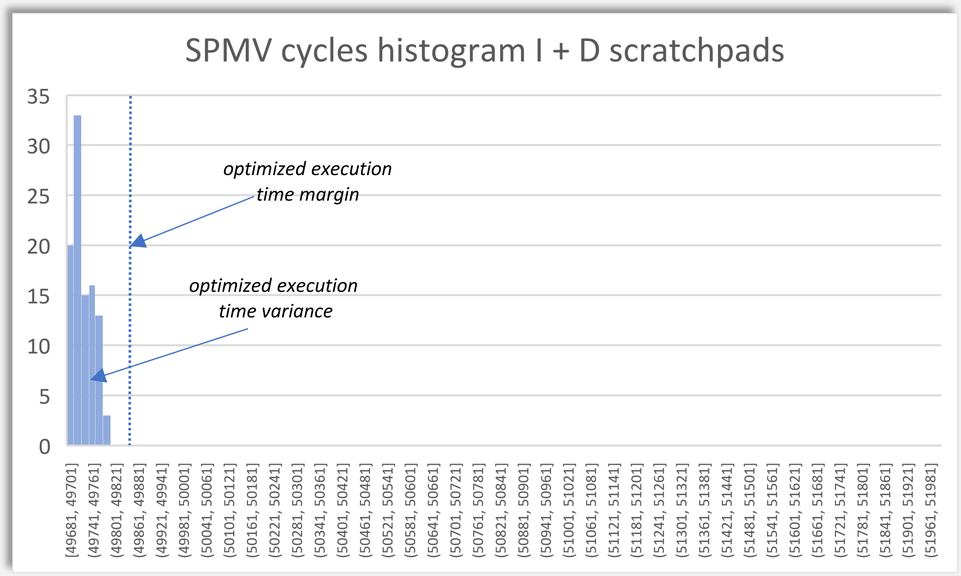

Conclusion

We demonstrated the required time determinism of this solution for real-time applications. Performance improvements brought by these features are still under development.

TRISTAN Project has received funding from the Chips Joint Undertaking (Chips-JU) under the grant agreement nr. 101095947. Chips-JU receives support from the European Union’s Horizon Europe’s research and innovation programme and Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Ireland, Israel, Iceland, Italy, Lithuania, Luxembourg, Latvia, Malta, Netherlands, Norway, Poland, Portugal, Romania, Sweden, Slovenia, Slovakia and Turkey.