The NimbleAI project, funded by the European Union's Horizon Europe programme, is developing a sensing-processing neuromorphic 3D chip that is inspired by the detection of light in eyes and the processing of visual information in brains. Looking to biological systems for inspiration, the project aims to deliver significant efficiency gains compared to mainstream processors. HiPEAC caught up with NimbleAI coordinator Xabier Iturbe (IKERLAN) to find out more.

Eye spy

NimbleAI's quest for a bio-inspired 3D chip for computer vision

Which biological systems provide inspiration for this project?

We are taking inspiration from both insect and vertebrate visual systems. We are designing a novel foveation mechanism – that is, a mechanism to boost resolution in selected regions or objects – for dynamic vision sensors (DVS). This will allow to dynamically adjust DVS resolution based on the value of visual information each sensor region brings to the application: a visual scene is sensed in low resolution to detect regions of interest to be foveated, that is, sensed in high resolution for better accuracy. The foveation is driven by selective-attention algorithms inspired by central and peripheral vision in vertebrates, which together provide an extremely efficient visual information-gathering mechanism.

We have also built the first-ever light-field DVS that captures directional information of incoming light to enable event-driven, ultra-low-latency and energy-efficient depth perception. This is inspired by the structure of compound insect eyes, which allow for effective depth and motion detection using very simple neural networks in the insects' tiny brains. To some extent, NimbleAI's neuromorphic vision is inspired by the sort of foveated compound eyes in dragonflies and mantises, which allow them to combine high-resolution vision of their prey with 3D perception to coordinate precise, swift aerial attacks. The NeuroEdge workshop, co-located with HiPEAC 2024, will elaborate on these concepts.

How do you plan to translate this biological inspiration into a chip?

We are adopting a time- and cost-effective approach that relies on reusing and extending the capabilities of existing technology (for example, using Prophesee DVS to capture light-fields). We are producing standalone testchips to implement new capabilities that are not commercially available, especially those related to the novel sensing mechanisms envisioned in the project, such as foveation. In parallel, we will be demonstrating how the different NimbleAI sensing and processing components work together at

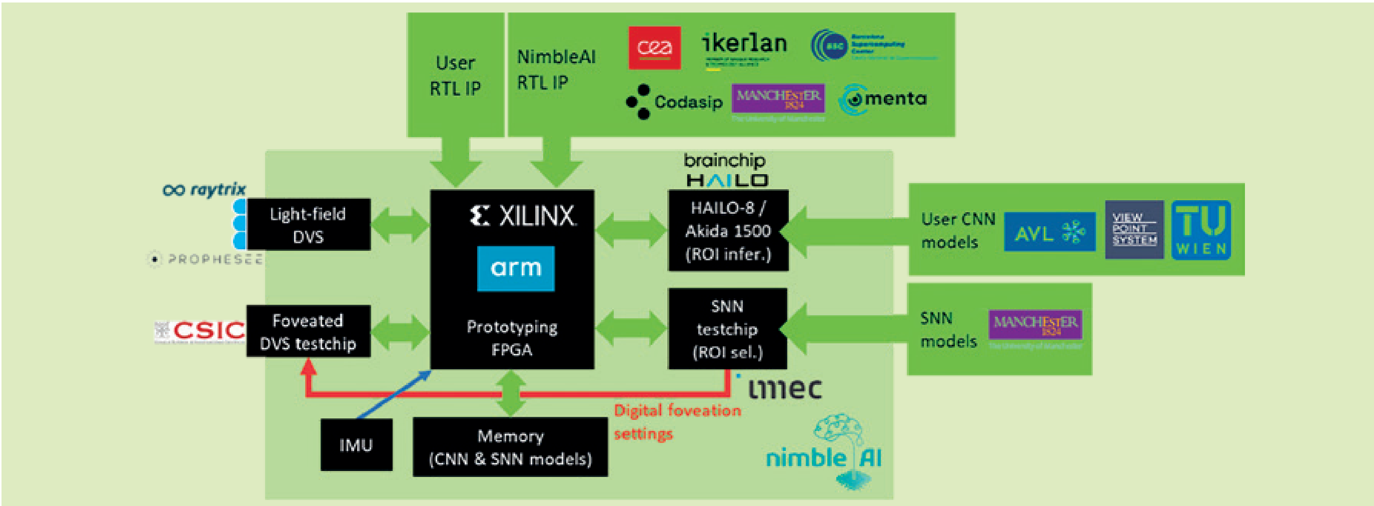

The NimbleAI PCB prototype to be delivered in Q4 2024 (some components will be available later) the logical and physical levels. We are advancing towards an actual 3D physical implementation of key components related to the novel NimbleAI vision modalities: the foveated DVS, light-field DVS and spiking neural network (SNN) engine that runs selective-attention algorithms. The latter components are to be partitioned and laid across at least three layers in the NimbleAI 3D silicon stack, allowing for the combination of different process technologies with specific support for features like low-noise image pixels and advanced high-performance nodes for neural-network processing. Although the project budget doesn't stretch to manufacturing and assembling such a complicated 3D chip, we will carry out a physical implementation including physical verification and sign-off to ensure technology feasibility.

In parallel, we are building a miniaturized printed circuit board (PCB) prototype that integrates NimbleAI testchips, commercial AI chips and a field-programmable gate array (FPGA), to emulate the functionality of the intended 3D chip. The expectation is that this prototype will help attract early adopters of the NimbleAI technology by allowing user applications to be tested on the emulated NimbleAI chip using live input data captured by NimbleAI DVS sensors and running users' convolutional neural networks (CNNs) on commercial AI chips, or on NimbleAI components prototyped on the FPGA. Both technology feasibility and adoption interest are key to continue developing the NimbleAI concept into an integrated 3D chip, and we expect to fulfil these two requirements within the project timeframe.

How is NimbleAI progressing in its objectives?

The design of the foveated DVS sensor is almost complete and will soon be sent to the foundry for manufacturing. The light-field DVS prototype has already been assembled and the first real-world datasets collected in an automotive setting. This allows us to validate the algorithms that we have been designing during the first year of the project using synthetic datasets. Likewise, CEA's computational memory has been optimized to run CNN inference on highly sparse DVS data.

What results can we expect over the next few months?

Once the light-field DVS algorithm is validated using real-world datasets, we will move from simulation to register transfer level (RTL) design and FPGA prototyping. This will be key to ensuring that the algorithm is not only useful in terms of efficiency and accuracy, but also delivers results with ultra-low latency. In fact, we expect ultra-low latency to be a key advantage of NimbleAI technology as, to the best of our knowledge, there is no passive 3D perception solution in the sub-millisecond range.

By the end of 2024 we will have a clearer view of the performance of our 3D perception algorithm using live real-world data. By the second quarter of 2024, Raytrix will have completed the adaptation of its software development kit (SDK) to support light-

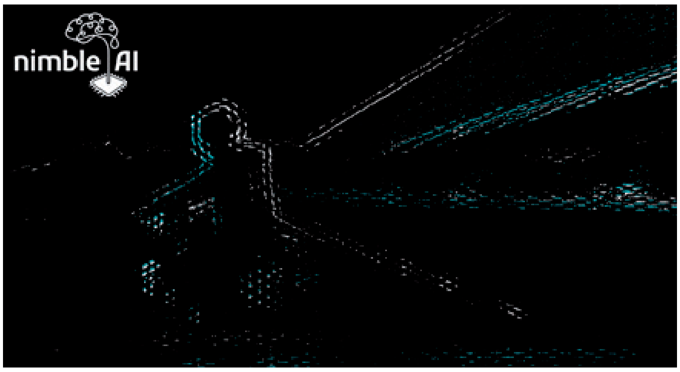

A motorcyclist captured in the first-ever light-field DVS dataset

field DVS data, allowing potential adopters of this technology to experiment with neuromorphic 3D perception. In early 2024, we will also launch the 3D silicon integration activities using the electronic design automation (EDA) tool and sensing and processing components designed in the first half of the project. Finally, the first functional prototype of the NimbleAI 3D chip, including platform software and hardware, will be available by the fourth quarter of 2024. We would encourage anyone wishing to test their vision pipelines on this prototype to contact us, so that they can harness the biological advantage of NimbleAI technology.

What kind of applications would you see this technology being used for?

Any application that requires energy-efficient and low-latency vision, especially those in which 3D perception is also important. Partner use cases in NimbleAI are good examples:

- automated and autonomous driving by AVL (where the focus is on ultra-low latency to react quickly to unexpected obstacles)
- eye tracking by Viewpointsystem (which focuses on energy efficiency for use in lightweight glasses)
- portable medical devices by ULMA Medical (where again the focus is on energy efficiency)

Beyond these, we target drone navigation and robotics, as well as space applications. In fact, we are discussing collaborations with the European Space Agency (ESA) and other key stakeholders in the space industry to explore the benefits of NimbleAI technology for space rendezvous and landing manoeuvres. Finally, we are also exploring the use of NimbleAI technology as an enabler for novel AI algorithms which have not yet made their way to industry, including neural circuit policy networks by TU Wien.

The **NeuroEdge** workshop at **HiPEAC 2024** will expand on topics discussed in this article, in particular the talk titled 'I spy with my little insect eyes: Combining DVS and light-fields' neuroedge.eu

NimbleAI has received funding from the EU's Horizon Europe research and innovation programme (grant agreement 101070679), and from UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (grant agreement 10039070).