This project has received funding from the Electronic Components and Systems for European Leadership Joint
Undertaking under grant agreement No 876925

ECSEL Joint Undertaking
Electronic Components and Systems for European Leadership

ANDANTE

AI for New Devices And Technologies at the Edge

D5.9 Demonstrations of the domain "Transport and Smart Mobility"

Deliverable No.	D5.9	**Due Date**	*31-Jan-2024*
Type	Demonstrator	**Dissemination Level**	*Confidential*
Version	1.0	**Status**	Final
Description	This document describes the final implementation of the demonstrators related to the Use Case Domain 3 "Transport and Smart Mobility", i.e. the ANDANTE Use Cases UC3.1 to UC3.5.		
Work Package	WP5 – Application Integration and Evaluation.		

Abstract (Published Summary)

This document serves as proof of the achievement of deliverable 5.9 “Demonstrations in the “Transport and Smart mobility” domain” for future solutions at Edge AI. Furthermore, it presents and discusses the final setup of the demonstrators as well as the components used for the implementation of these setups which have not been developed in other WPs for use cases (UC) selected in the context of Task 5.3 of WP5. of the ANDANTE project.

Within this task, five UCs are included: UC 3.1 “Drones/USV”, UC 3.2 “Underwater acoustic signal classification”, UC 3.3 “Detection of 3D objects and classification of road users based on LiDAR and camera”, UC 3.4 “Robust autonomous landing”, and UC 3.5 “Path planning and autonomous steering”. This document only provides an overview of the demonstrator configuration. D5.12 documents the evaluation and testing of these five use cases.

UC 3.1, Thales: Drones embedding neural network technologies are essential tools used in a large variety of applications from environment monitoring and wildfire detection to traffic jam monitoring, facilities, and border surveillance or even agriculture or resource management. These applications share same objectives as the detection or classification of objects or areas, they involve different timing constraints; for example, border surveillance requires small latency and inference time in the order of milliseconds, while environment monitoring may tolerate inference times of up to a few seconds.This use case aims to implement three different demonstrators to run the test campaigns and measure the KPIs (accuracy, inference performance and power consumption) and, finally, to analyze the executions to deduce the best configuration for a drone application. The results obtained show that a full formal approach for accelerating the inference on a drone, either using an ASIC or an FPGA, is a better option than the hybrid approach on an FPGA accelerator like the Thales hybrid one. The main reason is the power consumption of the ASIC solution is much lower in the order of mW instead of FPGA solutions in the order of ten watts.

UC3.2, Alseamar, CEA and SynSense The hardware/software solutions developed in ANDANTE aim to be integrated into Alseamar SeaExplorer glider, which is a long-endurance underwater drone with low power consumption. The main goal is to provide a solution that calculates in real time an accurate classification of underwater sound events with very low power consumption (of the order of magnitude of 1W). This deliverable presents two distinct solutions for this detection/classification application. The first approach is based on a vision processing framework where the audio is pre-processed online into images (called spectrograms) followed by a CNN (Convolutional Neural Network) architecture. This approach was implemented with the ASIC 2.1 (NeuroCorgi) and 4.1a Platform. The second approach was based on time series where the audio is converted into a sequence of events. It uses a SNN (Spiking Neural Network) architecture. This approach was implemented with the ASIC 1.3 and 4.1b Platform. It uses a SNN (Spiking Neural Network) architecture. This approach was implemented with the ASIC 1.3 and 4.1b Platform. The dataset used is called ocean dataset containing data recorded and labialized by Alseamar with many examples of 5 seconds audio samples belonging to 9 classes. The results obtained show high accuracy and very low power consumption of the order of mW.

UC3.3, Valeo: This use case aims to detect and classify objects from a 3D LiDAR point cloud and camera data will be explored for environmental modeling in the context of autonomous driving. LiDAR sensors are used for object detection, especially in high-speed driving scenarios such as highways and in busy environments, due to the high accuracy especially for long ranges imaging. However, details of the detected objects such as colors, textures, and other features related to pixel coloration cannot be used to perceive by the LiDAR. Therefore, cameras can be used to perceive object details. The idea is to fuse the camera’s semantic segmentation output of a CNN running on ASIC 2.1 and 4.1a platform with the 3D

LiDAR point cloud data processing unit to improve the overall environmental perception of the use-case. In this use case, the baseline model for camera semantic segmentation was trained on Valeo's internal dataset as well as the final segmentation model that consisted of the NeuroCorgi backbone (ASIC2.1) and custom segmentation head (in 4.1a Platform) performed well but could not outperform the baseline. The main reason for this is that the network backbone was fixed and only the segmentation head was trained on the dataset. However, the semantic segmentation with the NeuroCorgi backbone showed promising results and the use of neuromorphic chips in the perception pipeline for autonomous driving could be of great interest and benefit in the field of transport and mobility domain.

UC3.4, BRTE, GRADIANT, TVES and Cartogalicia: The main objective was to create a use case for robust autonomous landing. To achieve this goal, the following three independent designs were implemented and adapted to make inference on FPGAs: 1) Vision-based navigation for relative runway localization (BRTE), 2) Vision-based aerial and ground object detection (Gradiant), and 3) Communication quality analysis technologies (Televes).

BRTE collected a complete dataset called 'Skydata' containing images of runways during real landing maneuvers. This was the starting point for training the runway detection and obstacle detection models based on YoloV4 network. Then, BRTE implemented and adapted a neural network for runway detection and deployed it on the ZCU102 FPGA board. Finally, BRTE conducted a complete analysis on the inference accuracy, speed and power consumption of the runway detection model. The results obtained represent an improvement with respect to traditional GPU. Specifically, one-third reduction in power consumption, only 5% loss in accuracy, and the inference speed required for real-time operations.

GRADIANT adapted and optimized the RetinaNet (baseline) and YOLOv5 CNN models for hazard detection on runways and image recording for visual navigation to be deployed on a UAV platform using the ZCU102 FPGA-based board. After adapting the YOLOv5s model execution, this led to slightly lower maximum power consumption than a previously developed model based on the RetinaNet architecture on the ZCU102.Furthermore, there was an improvement in both peak processing speed (FPS) and accuracy metric (mAP@0.5) for object detection. Additionally, GRADIANT also deployed and evaluated these networks on a KRIA KV260 board. Performance and power consumption were measured in both cases.

TVES implemented ANN models for path loss estimation on the KRIA KV260 as an intermediate platform. These models were also ported in NeuroCorgi (ASIC 2.1). Test the performance of the deployed application and get more information about battery drain with different settings. TVES and Cartogalicia integrated the FPGA (used in Platform 4.1a) into the UAV in the laboratory and carried out energy consumption measurements.

Two different test flights were carried out, during which the demonstrator remained stationary. One of the test flights was carried out (with the KRIA board used in Platform 4.1a) for inference inferencing, while the other was carried out with platform 4.1a turned off, as a payload. The goal of these tests was to characterize the difference between the battery drain in these two scenarios. In addition to consumption, KPIs for accuracy and speed of execution were measured on the KRIA board. The result of the KPI measurements were favorable.

The final benchmarking of the flight platform built by Cartogalicia proved that the solutions implemented by BRTE, Gradiant and Televes can be deployed on a flight platform and perform controlled flight while performance inference tasks, essential to enable robust autonomous landing.

UC3.5, GML: This use case deals with dynamic path planning of vehicles, in the presence of other users, unforeseen obstacles, and changes in trajectory conditions. It can be compared to a highly dynamic and detailed short-range navigation system, but with constant updates to its local map of the environment, and constant updating the predicted trajectory of the vehicle among other vehicles and objects. Although the target of this type of application is typically found in e.g. warehouse robots, during the project we used the example of steering a car through a city.

This use case implements an Autonomous steering application. Datasets were generated by using the CARLA inner-city driving simulator and used to train a Resnet-50 network and converted into an event based SPARnet network. By working on this use case, we not only investigated, determined, and evaluated how to train GrAI Core on autonomous steering networks, but we also gained valuable insights regarding what it takes to integrate such networks into real-time operations, including some aspects of vehicle telemetry (steering, etc.). Overall, it was proven that the setup met the stringent automotive target KPIs (accuracy, latency, and degradation) partially.