



UNICA

UNIVERSITÀ  
DEGLI STUDI  
DI CAGLIARI



Università di Cagliari

UNICA IRIS Institutional Research Information System

**This is the Author's *accepted* manuscript version of the following contribution:**

G. Bellocchi, D. Madronal, A. Capotondi, F. Palumbo, A. Marongiu “An FPGA-based accelerator design methodology for smart UAVs in precision agriculture: A case study”, in Journal of Systems Architecture, 2025

**The publisher's version is available at:**

<https://doi.org/10.1016/j.sysarc.2025.103592>

**When citing, please refer to the published version (once available).**

**Currently use the following format:**

Gianluca Bellocchi, Daniel Madronal, Alessandro Capotondi, Francesca Palumbo, Andrea Marongiu, An FPGA-based accelerator design methodology for smart UAVs in precision agriculture: A case study, Journal of Systems Architecture, 2025, 103592, ISSN 1383-7621, <https://doi.org/10.1016/j.sysarc.2025.103592>.

(<https://www.sciencedirect.com/science/article/pii/S1383762125002644>  
)

This full text was downloaded from UNICA IRIS <https://iris.unica.it/>

# An FPGA-based accelerator design methodology for smart UAVs in precision agriculture: a case study

Omitted for blind review

---

## Abstract

Smart and Precision Agriculture (SPA) methods and technologies, such as autonomous robots, AI/ML, sensors, and actuators, enhance farming productivity by automating the retrieval of environmental parameters and the decision-making process, while Fog- and Edge-based paradigms enable more informed and responsive practices. Unmanned Aerial Vehicles (UAVs) can autonomously inspect crops and promptly cooperate with terrestrial vehicles to perform treatments, as recently demonstrated by the EU-funded COMP4DRONES (C4D) research project, focused on the provisioning of innovative UAV technologies for civilian applications. Modern companion-equipped UAV leverage Heterogeneous Systems-on-Chip (HeSoCs) to execute complex on-board tasks. HeSoCs generally combine a general-purpose, multi-core processor with a domain-specific accelerator-rich subsystem, massively integrating application-specific accelerators. Field Programmable Gate Array (FPGA)-based HeSoC are ideal fabrics to attain high performance and energy efficiency because of their massively parallel, deeply pipelined, non-Von-Neumann processing logic and custom memory hierarchies. Automated hardware-software co-design methodologies, e.g., FPGA overlays and toolflows, largely simplify the design phases, including the optimization of the accelerator interfaces, such as the merging of redundant components to reduce area usage. In this context, our contribution consists of a System-Level Design (SLD) methodology for the design of overlay-based UAV companion computers, including a modular and scalable accelerator-rich RISC-V HeSoC, a heterogeneous software stack and an automation toolchain to generate and integrate application-specific accelerators into our overlay. Our results show three optimized overlay variants targeting an UAV-based system employed in a SPA context. Experimental results denote improvements in performance and area usage, up to 18.5% on a FPGA-based HeSoC with respect to traditional design flows.

*Keywords:* UAV, Smart and Precision Agriculture, Companion Computer, FPGA Overlay, Accelerator-Rich Architecture

---

## 1. Introduction

*Smart and Precision Agriculture (SPA)* is revolutionizing the agricultural industry with innovative methods to improve operational efficiency, farming yield and crop health [1, 2]. These approaches largely evolved with the advent of Internet of Things (IoT) solutions, which allowed farmers to monitor many environmental parameters, perform analytics and observe the status of farms and machinery, including soil and crop health, usage of agrochemicals, fertilizers, water and seeds, livestock conditions, temperature and energy consumption, damage assessment caused by wildlife or other natural forces, and more [3, 4, 5]. This is achieved employing cutting-edge technologies—including autonomous robots, AI/ML algorithms, advanced sensors and actuators—to automate and enhance traditional human-centered farming activities [6, 7, 8, 9].

Moreover, agriculture demands spot-on remedies at exact space and time to reduce wastage of resources and ensure land and plant productivity [10, 11]. State-of-the-Art (SoA) computing paradigms—such as *Cloud*, *Fog* and *Edge*—are thus key to enable SPA [12]. These show large advantages compared to traditional *Human-Centered* or *Offline* practices, that generally postpone the application of treatments and/or farmer intervention with a negative impact on the crops' health status, costs and productivity. Indeed, *Offline* approaches include an initial on-field data collection and recordings, such as Light Detection and Ranging (LiDAR) point-clouds and multi-spectral images, and then the disjoint post-processing of raw measurements, which is performed by offline servers extracting insightful parameters and analyzing and classifying activities [13, 14, 15]. However, this implies significant delays in the decision-making process and a negative impact in the plantation foreseeable future [16]. *Cloud*-based solutions process data online on dedicated cloud servers. However, these retain major challenges, including privacy, scalability, high latency, network congestion and Quality of Service (QoS) degradation when huge amounts of heterogeneous farming data are transmitted [11, 17, 18]. Moreover, certain applications can suffer from the high energy consumption due to the data transfer to (and from) the cloud [19]. *Fog*- and *Edge*-based solutions commonly find similar definitions, as both are based on the concept of bringing cloud services and

35 resources closer to the edge data sources [12]. This is becoming computa-  
36 tionally practical with the advent of modern computing platforms, thus en-  
37 abling approaches which expedite data processing and enable more informed  
38 and responsive practices by moving the decision-making process closer to the  
39 crops [20, 21].

40 *Autonomous Unmanned Aerial Vehicles (UAVs)* will be pivotal in SPA  
41 activities, constituting reasonable human substitutes to examine the health  
42 status and growth evolution of the plantations, thus rapidly applying proper  
43 treatments [22]. The adoption of UAVs is mainly driven by the increased  
44 intelligence aboard and the advent of *companion-equipped drones* capable of  
45 executing complex tasks, e.g., real-time, data processing and analytics, or  
46 complex actuation [23, 24, 25, 26]. Another driver is the value generated by  
47 collected data, sampled closer to the crops thus guaranteeing more informed  
48 decisions [27]. Moreover, UAVs are expected to extensively cooperate with  
49 other robots, including Unmanned Ground Vehicles (UGVs), for synergistic  
50 and collaborative aerial and ground operations, aimed at enhancing oper-  
51 ational outcomes and providing greater flexibility in adapting to dynamic  
52 environments [28, 29]. Given the heterogeneity of robots, the cooperation  
53 may happen in different ways to exploit their complementarities [30].

54 Besides, the EU’s Common Agricultural Policy noted that greener farm-  
55 ing practices will be fundamental to increase crop yield and meet the expected  
56 doubling of world food production by 2050, given no additional land will be  
57 available to meet these needs [31]. It is widely acknowledged that the fu-  
58 ture of farming relies on research, innovation, and capacity building in the  
59 agri-food sector leveraging on cutting-edge technologies, including advanced  
60 embedded systems, robotics, and digital platforms, which help to build a  
61 more sustainable and efficient food system. A recent example is provided  
62 by the EU-funded COMP4DRONES (C4D) research project<sup>1</sup> that focused  
63 on innovating UAV technologies for civilian applications [32]. One of the  
64 use-cases was specifically devoted to the improvement of the autonomous  
65 capabilities of UAV-based systems in the SPA context<sup>2</sup>. Meeting this re-  
66 quirement complies with the ad-hoc re-design of the drone system and its  
67 components, where the *Flight Controller (FC)* represents the most critical.

68 The FC is the core component of any UAV and several Commercial Off-

---

<sup>1</sup><https://cordis.europa.eu/project/id/826610>

<sup>2</sup><https://www.youtube.com/watch?v=1NkCDg6HjiQ>

69 The-Shelf (COTS) solutions are typically based on  $\mu$ -controllers [33, 34].  
70 Modern UAVs also integrate cutting-edge, *on-board companion computers*,  
71 delivering SoA computing, memory and network characteristics, which en-  
72 able advanced on-board data processing and real-time decisions. Such com-  
73 panion computers commonly adopt embedded Heterogeneous Systems-on-  
74 Chip (HeSoCs), of which several COTS-based solutions are available [35].  
75 These combine a general-purpose, multi-core processor with a domain-specific  
76 subsystem for advanced processing, e.g., a custom ASIC, GP-GPU or *Field*  
77 *Programmable Gate Array (FPGA)*.

78 In particular, massively integrating application-specific accelerators in a  
79 single FPGA-based HeSoC—following a so-called *accelerator-rich paradigm*—  
80 offer unique opportunities for autonomous UAVs and SPA applications [36,  
81 37]. Indeed, such a companion computer can flexibly map and accelerate  
82 sophisticated workloads on its massively parallel, deeply pipelined, non-Von-  
83 Neumann processing logic and custom memory hierarchies, thus attaining  
84 high performance and low energy consumption, which is ideal for SPA appli-  
85 cations [38, 39]. However, heterogeneity hardens the hardware-software co-  
86 design process associated with bringing up the companion computer, hence  
87 requiring approaches to mitigate design costs [40, 41].

88 For instance, *FPGA overlays* consist of Hardware (HW) abstraction lay-  
89 ers that facilitate the mapping and scheduling of applications tasks to the  
90 FPGA fabric [42]. Besides, *toolflows* are also crucial for the optimization  
91 of the accelerators communication and control interfaces, including I/O data  
92 ports, register files, etc. In this context, Coarse-Grain Reconfigurable (CGR)  
93 accelerators represent a valuable option because they comprise low-overhead,  
94 decoupled, and inexpensive control interfaces. These features enable various  
95 optimization strategies, including merging redundant accelerator parts. This  
96 optimization would reduce area usage without additional overheads from  
97 complete, or partial, FPGA reconfiguration.

98 Under this premise, our contribution consists of a System-Level Design  
99 (SLD) methodology<sup>3</sup> for the design of overlay-based UAV companion com-  
100 puters, including:

- 101 • *On-board Overlay Compute Platform (OOCF)*, consisting of a modular  
102 and scalable accelerator-rich RISC-V HeSoC;

---

<sup>3</sup>Our contributions—OOCF, OODK and MDC—are open-source, as parts of the C4D research project.

- 103 • *On-board Overlay Development Kit (OODK)*, a heterogeneous Soft-  
104 ware (SW) stack to streamline the development of the heterogeneous  
105 application;
- 106 • *Multi-Dataflow Composer (MDC)*, an automation toolchain that allows  
107 automatic generation and integration of CGR accelerators inside our  
108 OOCF.

109 We show the results achieved with three alternative variants of the OOCF  
110 architecture. Experimental results were obtained by running the C4D use-  
111 case application on a COTS FPGA-based HeSoC from the AMD Zynq Ultra-  
112 scale+ (US+) family. We demonstrate improvements in both performance  
113 and area usage of up to 18.5% compared to traditional integration flows such  
114 as [43], thanks to the merging and optimization features of our proposed flow.

115 The manuscript is organized as follows. Section 2 describes background  
116 content, including the C4D’s project, the SPA use-case scenario and UAV  
117 components. Section 3 discusses the related works, while Section 4 intro-  
118 duces the components of the SLD methodology—OOCF, OODK and MDC.  
119 Then, in Section 5, we discuss the experimental results after the design of  
120 three variants of our overlay-based UAV companion computer, as well as an  
121 assessment of the methodology scalability for larger and more complex UAV  
122 systems. To demonstrate the broad applicability of the proposed approach,  
123 not limited to the SPA context, Section 6 describes additional projects en-  
124 abled with the proposed approach.

## 125 2. Background

### 126 2.1. The COMP4DRONES project

127 The ECSEL JU project C4D—Framework of Key Enabling Technologies  
128 for Safe and Autonomous Drones—terminated in 2023 under the coordina-  
129 tion of Indra and involvement of 50 partners. The Italian cluster included 12  
130 partners and was led by the University of Sassari (UNISS), also involved in  
131 providing one of the assessment test case for the C4D technologies. The main  
132 goal of the project was to offer a framework of key enabling technologies for  
133 the design and operation of drones, ranging from application to electronic  
134 components, realized as a tightly integrated multi-vendor and compositional  
135 embedded architecture solution, and a toolchain to assemble and safely op-  
136 erate drones [32]. The main idea of the assessment test case from the Italian

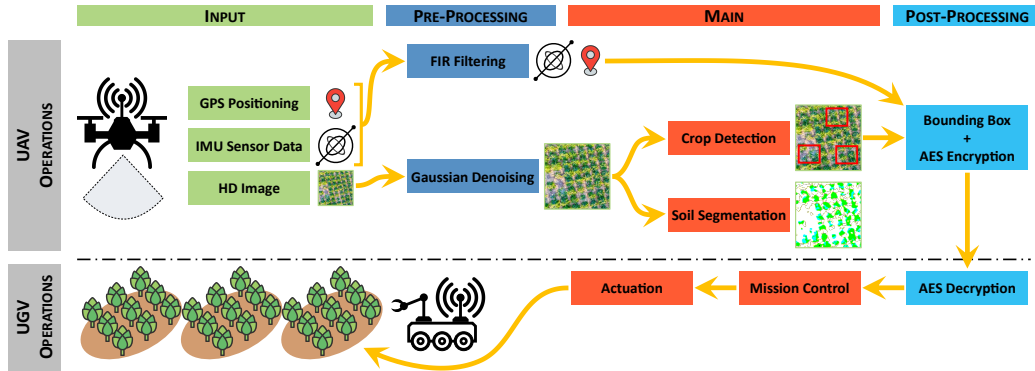


Figure 1: Representation of the C4D use-case application scenario.

137 cluster was to improve SPA technologies by providing more advanced ob-  
 138 servation and intervention methodologies through combined usage of a UAV  
 139 and a UGV.

## 140 2.2. UAV for smart and precision agriculture

141 In the SPA context, UAV may handle many tasks, including the mapping  
 142 and detection of crops and plants [44, 45, 46], monitoring sessions and in-  
 143 spection operations [47, 48, 49], disease detection [50, 51, 52] and application  
 144 of treatments [53, 54]. Besides, flight-oriented data processing—such as path  
 145 planning, obstacle detection and collision avoidance—are also crucial for safe  
 146 environment exploration [55, 56, 57].

147 The C4D SPA use-case scenario embodies a subset of the tasks above and,  
 148 specifically, the autonomous detection and monitoring of artichoke crops.  
 149 The final goal consists of preserving a healthy yield status by monitoring the  
 150 crop health status and searching for potential anomalies during the UAV’s  
 151 aerial traversal over a preset plantation route. During the flight, the drone  
 152 examines the plants and searches for crop discontinuities. Besides, the UAV  
 153 collaborates with a UGV in a synergistic effort for plant monitoring and  
 154 health and growth management. Closing the loop faster brings advantages in  
 155 the agriculture context, which largely benefits from fast reactions to anom-  
 156 alies, at exact space and time. In general, both the UAV and the UGV act as  
 157 mobile sensors: the drone acts as a *sentinel* ensuring rapid reconnaissance  
 158 and online decisions, while the terrestrial rover acts as a *mobile actuator*,  
 159 meant for detailed recognition, spot-on spraying of nutrients, water and pes-  
 160 ticides and support in case of anomalies or abnormal circumstances. Besides,




Figure 2: System-level design stack.

161 both extract and process complex plantation images to deduce relevant infor-  
 162 mation and communicate over a trusted and secure medium, which is ensured  
 163 with lightweight cryptographic modules for data-privacy, authentication and  
 164 communication integrity.

165 Figure 1 shows the C4D use-case application scenario, where the UAV and  
 166 UGV collaborate synergistically. In this manuscript, we focus our attention  
 167 to the portion of the use-case executed by the UAV, given the UGV employs  
 168 the same hardware-software technologies. Hence, the processing pipeline  
 169 implemented by the UAV consists of: (i) INPUT DATA ACQUISITION; (ii)  
 170 PRE-PROCESSING; (iii) MAIN; and (iv) POST-PROCESSING.

171 The UAV initially samples both *visual and text information*—including  
 172 a  $1440 \times 900$  HD image, IMU sensor data, and GPS position—and then  
 173 forwards them to the companion computer, which initializes computing re-  
 174 sources and starts executing the successive processing stages. As soon as the  
 175 pre-processing phase starts, the field image is smoothed by applying *Gaus-*  
 176 *sian Denoising (Conv)* to prune image noise for the subsequent processing  
 177 steps, while IMU and GPS data encompass a *Finite Impulse Response (FIR)*  
 178 filtering stage. Afterwards is the main processing phase, where *Crop De-*  
 179 *tection* and *Soil Segmentation* are executed. Crop detection is implemented  
 180 by a *Convolutional Neural Network (CNN)* model for artichoke plant detec-  
 181 tion and classification through a customized Feature Pyramid Network CNN  
 182 that works as a Single Shot Detector (SSD) [46]. Besides, image segmenta-  
 183 tion leverages at its heart a *Canny edge detector (Canny)* [58]. Finally, the  
 184 post-processing stage packs textual data in a *Bounding BOX (BBOX)*, which  
 185 is subsequently encrypted with an *Advanced Encryption Standard (AES)* for  
 186 trusted and secure communication toward the UGV.

187 *2.3. UAV components*

188 An UAV system comprises many components, including a multi-rotor  
189 or fixed-wing frame, Electronic Speed Control (ESC) modules, and a nav-  
190 igation system, relying on various sensors including Inertial Measurement  
191 Unit (IMU), barometers, Global Navigation Satellite System (GNSS), and  
192 radio transmitters [35]. Moreover, a crucial UAV component is the FC, of  
193 which many COTS alternatives are available [33]. The FC is generally based  
194 on  $\mu$ -controllers and supports flexible I/O interfaces to facilitate the integra-  
195 tion of additional sensors and mission-specific payloads, such as cameras and  
196 LiDARs. Besides, the FC handles various on-board functionalities, including  
197 flight control, navigation, data acquisition and processing, mission control  
198 and actuation [34].

199 Modern UAVs couple the FC with cutting-edge, *on-board companion com-*  
200 *puters* to empower the drone system with advanced processing capability, as  
201 well as SoA memory and network capability [16]. UAV companion computers  
202 typically adopt embedded HeSoCs, combining a general-purpose, multi-core  
203 processor with a domain-specific subsystem for advanced processing, such as  
204 a custom ASIC, GP-GPU or FPGA.

205 We believe that FPGA-based HeSoCs offer significant advantages for  
206 the design of complex UAV stacks, including the on-the-fly deployment of  
207 computing tasks and the availability of embedded soft-cores to simplify the  
208 HW/SW partitioning of the target application [38, 59]. However, existing  
209 tools are not mature enough to reduce the design strain to bring up a full-  
210 fledged FPGA-based companion computer and the integration of heteroge-  
211 neous components and accelerators [40]. Thus, in the following sections,  
212 we propose an innovative SLD methodology, based on a FPGA overlay, to  
213 streamline the design and optimization of accelerator-rich HeSoCs.

214 **3. Related work**

215 *3.1. About the design of UAV companion computers for SPA applications*

216 This section compares existing literature on companion-equipped drones  
217 targeting SPA applications, including our work. In particular, our analysis  
218 focuses on works where the UAV companion computer enables autonomous  
219 and informed decisions on crops, safe navigation within plantations, and  
220 communication with other system components, including ground stations or  
221 other robots [26, 35, 28]. Table 1 summarizes the features of each approach.

Table 1: Overview of UAV companion computers for SPA applications.

	TECHNOLOGY			MAIN SPA TASK	
	Platform	Type	Accelerator	Task	Algorithm
Alsalam et al. [60]	Odroid-U3+	CPU	—	Weed detection	Vision
Dai et al. [61]	Intel i5-6260	CPU	—	Pesticides spraying	Vision
Horstrand et al. [62]	NVIDIA Jetson TK1	HeSoC	GPU	Crop monitoring	VI <sup>1</sup>
Saddik et al. [63]	Intel Cyclone V	HeSoC	FPGA	Crop monitoring	VI <sup>1</sup>
<b>This Work</b>	<b>AMD Zynq US+</b>	<b>HeSoC</b>	<b>FPGA</b>	<b>Crop detection, Soil Segmentation</b>	<b>CNN, Edge Detection</b>

<sup>1</sup> VI: Vegetation Index

222 UAV companion computers can be designed using homogeneous plat-  
 223 forms, where SPA tasks are executed entirely on general-purpose CPUs. *Al-*  
 224 *salam et al.* use the Odroid-U3+ to detect invasive weeds and implement  
 225 a decision-making control loop [60]. The companion computer runs vision  
 226 algorithms and controls the UAV flight height. Other components handle  
 227 path planning, environment inspection, and herbicide application. Although  
 228 the Odroid-U3+ integrates an ARM Cortex-A9 CPU and a Mali-400 GPU,  
 229 it is only used to a limited extent and GPU acceleration is neither utilized  
 230 nor discussed. Similarly, *Dai et al.* present a vision-based UAV system  
 231 for spraying pesticides on fruit trees [61]. The system employs an Intel i5-  
 232 6260 CPU to process images, handle flight control, and communicate with a  
 233 ground station. In contrast, our UAV system communicates with a UGV in  
 234 a synergistic effort to improve response time to anomalies in the plantation.  
 235 Moreover, our approach and other works advocate heterogeneous platforms  
 236 to increase intelligence aboard UAVs by combining CPUs and accelerators,  
 237 such as FPGAs and GPUs, for more complex SPA tasks.

238 *Horstrand et al.* developed an autonomous UAV for hyperspectral mea-  
 239 surements that provides vegetation indices for vineyard analysis [62]. The on-  
 240 board computer includes an NVIDIA Jetson TK1 for autonomous flight con-  
 241 trol, data acquisition, processing, and communication with a ground station.  
 242 However, the main limitations of GPUs are the non-extendable HW com-  
 243 ponents, non-deterministic execution, and high power consumption, often  
 244 exceeding the availability of resource-constrained autonomous robots [38, 64].

245 To this end, FPGA-based HeSoCs are ideal solutions for UAV compan-  
 246 ion computers. *Saddik et al.* propose a vegetation monitoring system that  
 247 accelerates GRVI and GLI indices using an Intel Cyclone V FPGA [63]. How-  
 248 ever, their Digital Signal Processor (DSP) utilization is low ( $\text{DSP}_{\%,\text{ref}} = 6\%$ ),

Table 2: Overview of SLD methodologies for FPGA-based HeSoCs.

	ACCELERATOR					OVERLAY					
	Type	Host	Proxy Core <sup>2</sup>		Interconnect <sup>3</sup>						
	TI <sup>1</sup>	HLS	RTL	CGR	System-Level Integration	RV64	AArch64	RV32	MicroBlaze	Cluster-level	HeSoC-level
PYNQ [67]		✓			✓				✓	-	STREAM
RobotCore [68]	✓	✓			✓		✓			-	STREAM
OpenESP [69]	✓	✓	✓		✓	✓				-	NOC
TaPaSCo [70]	✓	✓	✓		✓		✓			XBAR	XBAR
<b>This Work</b>	✓	✓	✓	✓	✓	✓	✓	✓	✓	XBAR	XBAR

<sup>1</sup> TI: Technology-Independent.

<sup>2</sup> Note that both OpenESP and TaPaSCo support RV32, but they do not exploit it as a proxy core, i.e., for tightly-coupled runtime management and optimization of accelerator scheduling, memory transfers and other pivotal platform tasks, as discussed in Section 4.1.

<sup>3</sup> Interconnects are classified as: (i) XBAR: Crossbar; (ii) NOC: Network-on-Chip; and (iii) STREAM: Based on a streaming protocol.

249 whereas previous works show benefits of highly-pipelined DSP datapaths in  
 250 FPGA-based accelerators [65, 66]. In contrast, our companion computer uses  
 251 an accelerator-rich paradigm, where many application-specific accelerators  
 252 are integrated to fully exploit the advantages of FPGA fabrics.

### 253 3.2. About SLD methodologies to simplify the design of FPGA-based HeSoCs

254 In this section, we examine various techniques for streamlining the design  
 255 of FPGA-based HeSoCs. To this end, Table 2 summarizes characteristics of  
 256 several industrial and academic approaches, including our own.

257 *PYNQ* is an open-source overlay design methodology that improves us-  
 258 ability of AMD HeSoC products [67]. It is particularly suitable for non-FPGA  
 259 experts, as it greatly simplifies the application development for embedded  
 260 FPGA-based HeSoCs by offering support for Ubuntu Linux, Python, and  
 261 Jupyter notebooks [71]. PYNQ overlays enable the deployment of accelerator-  
 262 rich systems with support for dynamic partial reconfiguration [72]. However,  
 263 beyond provided examples, users must still undergo traditional platform de-  
 264 velopment using the AMD Vivado Suite. In contrast, our SLD methodol-  
 265 ogy fully automates the customization of accelerator-rich overlays. Users  
 266 need only supply accelerator sources, e.g., designed with High Level Synthe-  
 267 sis (HLS), and Python specification files to define: (i) how accelerators are  
 268 interfaced and optimized (e.g., through datapath merging), as described in  
 269 Section 4.2.2; and (ii) how accelerators are aggregated at the system level

270 along with their respective optimizations, as detailed in Section 4.3.

271 The PYNQ IP portfolio includes all AMD IPs, including: (i) a 64-bit  
272 ARM host processor, supporting legacy SW stacks; (ii) a MicroBlaze soft  
273 core, a flexible and deterministic interface to a wide range of peripheral I/Os.  
274 In comparison, our OSCP features a broad portfolio of open-source IPs from  
275 the Parallel Ultra Low Power (PULP) Platform. Unlike PYNQ, our overlay  
276 is technology-independent and our design flow can be extended to support  
277 FPGAs from different vendors.

278 Finally, the PYNQ soft core communicates with the host processor via  
279 shared memory, but supports only basic communication libraries [73]. To en-  
280 able more complex interactions, users must implement their own communica-  
281 tion SW, requiring development for two processors. However, complex UAV  
282 applications demand robust communication mechanisms between the host  
283 and the FPGA subsystem. Moreover, accelerator-rich systems require flexi-  
284 ble control and optimization of data transfers and accelerator management.  
285 To address these needs, our OODK supports the OpenMP programming  
286 model, which simplifies the development of complex heterogeneous software.  
287 With this model, users write a single application that covers host and proxy  
288 core. The heterogeneous application begins execution on the host, while  
289 compute-intensive tasks are offloaded to the accelerator-rich subsystem us-  
290 ing a dedicated computation offloading command [74]. Our FPGA overlay  
291 thus provides built-in support for flexible communication between the host  
292 and the accelerator-rich subsystem. This interaction is managed by the proxy  
293 core, which acts as a man-in-the-middle to coordinate data movement and  
294 task execution across subsystems.

295 *RobotCore* targets Robot Operating System (ROS) 2 computational graphs  
296 for robotic applications [68]. The approach can accelerate compute nodes and  
297 ROS 2 communication channels to avoid expensive interactions with the host  
298 processor. However, RobotCore’s FPGA subsystem is hardwired, making it  
299 less flexible than our programmable FPGA overlay.

300 *OpenESP* consists of a SLD framework for the seamless hardware-software  
301 integration of application-specific accelerators into HeSoCs [69]. Unlike our  
302 proposal, OpenESP uses the FPGA fabric as an emulation platform for ASIC  
303 prototyping and integrates coarse-grained accelerators that are loosely cou-  
304 pled with the processor. In contrast, our accelerator-rich overlay comprises  
305 numerous fine-grained accelerators, which are grouped into clusters and then  
306 integrated at the system level—for example, to form more sophisticated  
307 accelerator-rich pipelines. Additionally, our approach differs from OpenESP

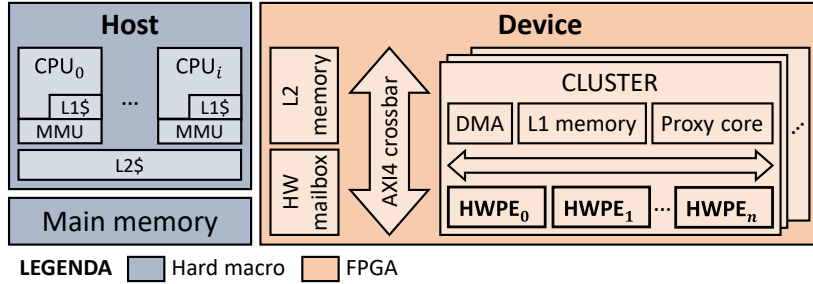
308 in terms of SW support. While ESP interfaces accelerators using a Linux  
309 driver, OOCp connects accelerators directly within the heterogeneous appli-  
310 cation or through user-space libraries. This is enabled by our OODK, which  
311 provides full support for OpenMP [75].

312 *Task Parallel System Composer (TaPaSCo)* is an open-source methodol-  
313 ogy designed to simplify the development of accelerator-rich platforms, sup-  
314 porting many FPGA fabrics—from embedded devices to High Performance  
315 Computing (HPC) systems [70]. In TaPaSCo, HLS accelerators of the same  
316 type are grouped into homogeneous clusters. Heterogeneous platforms are  
317 then built by interconnecting different clusters. In contrast, our approach  
318 introduces heterogeneity directly at the level of OOCp clusters, which inte-  
319 grate heterogeneous, fine-grained accelerators. This architectural choice is  
320 fundamental to efficiently supporting accelerator-rich optimizations, includ-  
321 ing: (i) area-oriented techniques, such as datapath merging enabled by MDC;  
322 and (ii) performance-oriented strategies, such as the use of a proxy core for  
323 double-buffering and memory latency hiding, as discussed in Section 6.

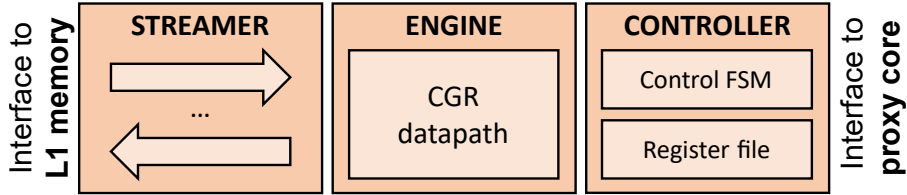
324 TaPaSCo also supports RISC-V soft cores, used as processing elements  
325 for many-core platforms. Our proxy core is also RISC-V compliant but differs  
326 from the TaPaSCo soft core in intent. Indeed, the proxy core is a key enabler  
327 of high-performance, accelerator-rich platforms by orchestrating their many,  
328 co-existing tasks. This approach is flexible because it is SW-defined. This  
329 also includes the task scheduling of fine-grained accelerators, with automatic  
330 generation of SW libraries and Application Programming Interfaces (APIs).  
331 Furthermore, it is low-cost since the proxy core is tightly-coupled to the  
332 components of the OOCp cluster, including accelerators. This supervisory  
333 role of the proxy core does not limit the SW developer, who can still choose  
334 to deploy computation on the proxy core, as demonstrated in Section 6.2.

## 335 4. Methodology

336 This section describes our SLD methodology, shown in Figure 2 and in-  
337 cluding: OOCp(i) OOCp, our FPGA overlay consisting of a modular and  
338 scalable accelerator-rich RISC-V HeSoC; (ii) OODK, our heterogeneous SW  
339 stack to streamline the development of the heterogeneous application; and  
340 (iii) MDC, our automation toolchain that allows automatic generation and  
341 integration of CGR accelerators inside OOCp. We start by describing these  
342 components and the complete workflow that enables the design of overlay-  
343 based UAV companion computers.



(a) HeSoC architecture deployed on Zynq US+ devices.



(b) HWPE-based wrapper.

Figure 3: Overview of the OOC components.

#### 344 4.1. Overlay architecture

345 Our FPGA overlay consists of two components: (i) *OOC*, a modular  
 346 and scalable accelerator-rich RISC-V HeSoC [43]; and (ii) *OODK*, a het-  
 347 erogeneous SW stack to streamline the development of the heterogeneous  
 348 application. The architectural template of the overlay builds on top of the  
 349 PULP Platform, an open and scalable hardware-software research and de-  
 350 velopment platform for highly parallel, ultra-low-power architectures, based  
 351 on the open-source *RISC-V* Instruction Set Architecture (ISA) [76].

352 The OOC architecture sticks to a standard HeSoC with two independent  
 353 subsystems—*host* and *device*—as shown in Figure 3a. The *host* is a Linux-  
 354 capable multi-core application processor. If the target HeSoC fabric supports  
 355 a physical host processor, then it can be leveraged with its HW/SW legacy.  
 356 For instance, this is the case for the Zynq US+ family of devices, where  
 357 an AArch64 application processor is integrated [77, 78]. Alternatively, our  
 358 architectural template also supports the OpenHW Group CVA6 based on the  
 359 RV64 ISA, which can implemented in FPGA [79].

360 The *device* subsystem consists of an accelerator-rich subsystem, including  
 361 *clusters* of processing elements sharing a multi-banked SW-programmable L2

362 Scratchpad Memory (SPM) via a fully-connected AXI4 CrossBAR (XBAR).  
363 The communication between the host-device subsystems happens via coarse-  
364 grained computational offloading, which leverages the off-chip main memory  
365 and a HW mailbox for synchronization.

366 The *cluster* integrates *accelerators* and additional support components,  
367 including a *proxy core*, L1 memory—implemented as a multi-banked Tightly-  
368 Coupled Data Memory (TCDM)—and a SW-programmable Direct Memory  
369 Access (DMA) engine for efficient bi-directional 64-bit data transfers between  
370 the L2 and L1 memory hierarchies. L1 memory accesses occur through the  
371 cluster Low-Latency Interconnect (LIC).

372 The *proxy core* is a pivotal component meant for tightly-coupled run-  
373 time management of various cluster tasks, such as the orchestration of the  
374 accelerator tasks, DMA bandwidth monitoring and control operations and  
375 many more [80, 81]. It consist of a low-cost soft-core, implementing the RV32  
376 ISA and fetching instructions from the L2 memory through its instruction  
377 cache [82, 83]. This approach is more flexible compared to a dedicated Finite  
378 State Machine (FSM) and more scalable and tightly-coupled compared to a  
379 costly invocation of the host processor for the same intent. Moreover, the  
380 proxy core implements a flexible and customizable host-device interface.

381 *Accelerators* leverage Hardware Processing Engine (HWPE) interfaces to  
382 ensure trustworthy and flexible communication between their datapaths and  
383 the cluster components [84, 85]. HWPE-based accelerators are integrated  
384 inside the OOC cluster and tightly-coupled to the L1 memory.

385 Figure 3b shows the key modules of a HWPE wrapper: (i) an *engine*,  
386 implementing the accelerator datapath; (ii) a *streamer*, providing an I/O  
387 data interface to L1 memory; and (iii) a *controller*, for coarse-grained accel-  
388 erator control/(re)configuration. The use of streaming interfaces decouples  
389 the CGR datapath from the outer system, with benefits in terms of latency  
390 tolerance. Besides, the adoption of the HWPE wrapper increases the MDC  
391 merging flexibility. Indeed, a typical concern of merging approaches is that  
392 they cannot usually cope with the merging of kernels with different commu-  
393 nication protocols because of incompatibilities at the I/O interface. With  
394 the adoption of HWPE, more uneven CGR datapaths can be merged in a  
395 single accelerator wrapper, thus integrated into our accelerator-rich OOC.

396 The OODK consists of a heterogeneous SW stack supporting the ap-  
397 plication developer. OODK supports the OpenMP programming model to  
398 facilitate the development of complex heterogeneous SW, as well as a SoA  
399 LLVM-based heterogeneous compiler, which allows for the development of




Figure 4: MDC: From dataflow specification to the generation of a CGR datapath.

400 single-source single-binary heterogeneous applications with OpenMP 4.5 of  
 401 offloading. The heterogeneous application always starts on the host processor  
 402 and then critical tasks can be transparently offloaded to the accelerator-rich  
 403 subsystem, simply taking advantage of OpenMP pragmas.

#### 404 4.2. Automatic generation of CGR accelerators

405 The automated design and integration of application-specific accelerators  
 406 is accomplished by *MDC*, an open-source automated tool for the generation  
 407 of CGR accelerators and their integration inside OACP [86, 87].

408 Generally speaking, according to the typical MDC flow, the generated  
 409 HW/SW components consist of:

- 410 • *accelerator datapaths*, merging different functionalities in a single ac-  
 411 celerator exploiting the principles of CGR reconfigurability;
- 412 • *accelerator wrappers*, implementing communication/control interfaces  
 413 meant for the interaction of the CGR accelerator datapaths with other  
 414 HW components that, in typical usage so far, were intended as com-  
 415 munication/control interface toward the processing system;
- 416 • *accelerator drivers*, to simplify and abstract the configuration and man-  
 417 agement of the CGR accelerator implemented within the wrapper.

##### 418 4.2.1. CGR datapath generation

419 The generation of CGR datapaths is in charge of the baseline MDC func-  
 420 tionality, which is composed of a front-end and a back-end. The front-end  
 421 is called Multi-Dataflow Generator (MDG) and is responsible of datapath  
 422 merging; whereas the back-end is called Platform Composer (PC), and is re-  
 423 sponsible of translating a high-level merged dataflow model into an Hardware  
 424 Description Language (HDL) specification.

425 The MDG mandates a *dataflow specification*, which comprises a set of  
426 tasks—or functions—to be accelerated defined as dataflow networks, that  
427 are directed graphs connecting actors through First-In First-Out queues (FI-  
428 FOs). In the example of Figure 4, three tasks are represented in the top-left  
429 corner. The MDG block analyzes the input networks, locates the shared  
430 actors—A and C are shared, in this specific case—and generates a *multi-*  
431 *dataflow network* combining all the functionalities. The combination is done  
432 by generating only one instance of each shared actor and including a set of  
433 switching boxes (SBoxes). These SBoxes allow the selection of the path that  
434 the data will follow when selecting a specific functionality to be executed.  
435 To do so, a *configuration table* gathering the specific values to configure the  
436 SBoxes is also automatically generated.

437 As a second step, the PC derives the HDL specification of the CGR  
438 datapath from the multi-dataflow network, requiring: (i) the XML definition  
439 of the communication protocol between actors; and (ii) the HDL description  
440 of the actors, e.g., SystemVerilog, Verilog or VHDL, contained in the *HDL*  
441 *components library*. This latter can be manually generated or automatically  
442 synthesized leveraging HLS techniques, e.g. Vitis HLS.

#### 443 4.2.2. Automatic integration inside the FPGA overlay

444 An automatic procedure also occurs for generating accelerator wrappers,  
445 which are necessary to integrate CGR accelerators within our OOCF.

446 MDC has been extended to support the HWPE interface to ensure trust-  
447 worthy and flexible communication between accelerators and the OOCF com-  
448 ponents. HWPE-based accelerators are integrated inside the OOCF cluster  
449 and tightly-coupled to the L1 memory.

450 The Co-Processor Generator (CPG) functionality of the MDC tool, shown  
451 in Figure 5, currently supports a novel feature to automatically generate the  
452 HWPE-compliant wrapper. Through this feature, the CPG became capable  
453 of encapsulating the PC outcome, the CGR datapath of a reconfigurable  
454 accelerator, with the necessary glue logic for integrating it inside the OOCF  
455 cluster. Moreover, as part of the new feature, MDC CPG also takes care of  
456 generating the necessary accelerator drivers to invoke the accelerator from  
457 the heterogeneous application, developed via the OODK.

458 To this end, the tool requires a *wrapper specification* along with the *CGR*  
459 *datapath specification*. The former describes the characteristics of the inte-  
460 grated CGR datapath for customizing the HWPE streamer and controller.

461 Multiple CGR datapaths cannot run concurrently when merged in the

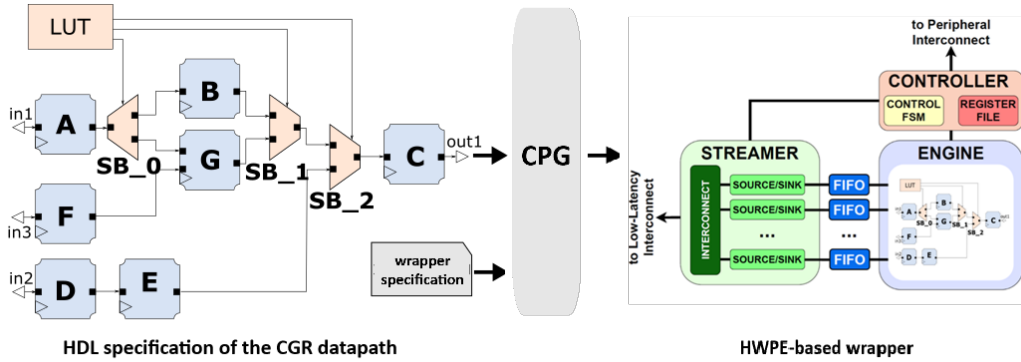


Figure 5: MDC: From multi-dataflow specification to the generation of the HWPE wrapper.

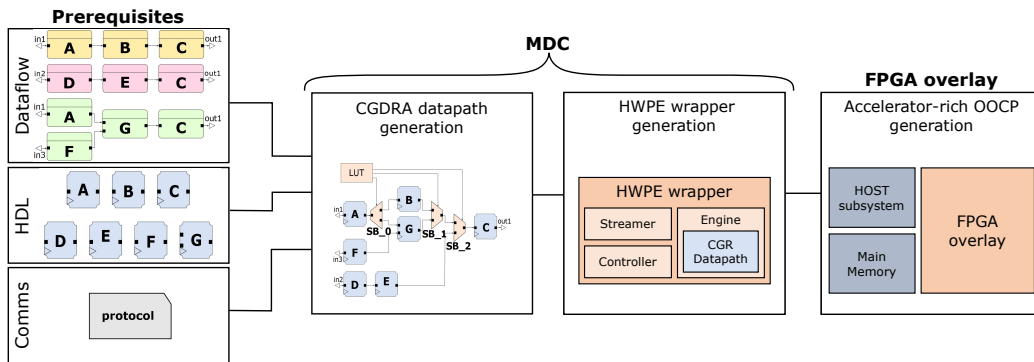


Figure 6: SLD methodology to design full-fledged UAV companion computers based on an accelerator-rich FPGA overlay.

462 same HWPE wrapper. Yet, to switch functionality requires a single clock cy-  
 463 cle, given task selection occurs via the simple SW configuration of a HWPE  
 464 register. Besides, the area usage of a multi-functional accelerator, where sev-  
 465 eral tasks co-exist in the same wrapper, is minimized thanks to the merging  
 466 process of actors, I/O resources and register files.

#### 467 4.3. SLD methodology for UAV companion computers

468 The combination of our FPGA overlay with the MDC framework results  
 469 in the SLD methodology, shown in Figure 6. This workflow is meant to  
 470 streamline the design, optimization and deployment of accelerator-rich UAV  
 471 companion computers. The flow offers an abstract HW/SW co-design inter-

```

1 class oocp_specs:
2
3     def OOCF(self):
4         self.name      = Accelerator-rich platform
5         self.target    = FPGA fabric
6         self.l2_mem    = [Number of ports, Size]
7         return self
8
9     def cluster_0(self):
10        self.acc       = [Accelerator name, ...]
11        self.proxy     = [IP, Number of cores, ...]
12        self.dma       = [IP, Job queue size, ...]
13        self.l1_mem    = [Number of ports, Size]
14        return self
15
16        ...
17
18    def cluster_N(self)
19        ...

```

Listing 1: Example of OOCF specification.

472 face to improve ease-of-use and supports datapath merging and automatic  
473 accelerator integration in an accelerator-rich context.

474 As prerequisites for using our SLD methodology, the user defines the ap-  
475 plication requirements, taking note of the application task(-s) to accelerate,  
476 leveraging the presented dataflow approach and providing the input speci-  
477 fications as *dataflow networks*. Moreover, the user must provide a library  
478 of actors, described as *HDL specifications*, and the *communication* protocol  
479 specifying their data exchange protocol.

480 At this point, MDC automatically generates the *CGR datapaths*, merging  
481 the input tasks' specification (as detailed in Section 4.2.1). Subsequently,  
482 MDC also generates the HWPE wrappers and related drivers (as detailed in  
483 Section 4.2.2). Please note that the merging process is completely under the  
484 control of the user. Indeed, the user should select the set of input dataflow  
485 specifications to be accelerated, and pass them in input to MDC. In this  
486 work, we opted for two different strategies to select the set of tasks to be  
487 merged within the HWPE, *application-aware*, that takes into consideration  
488 the application task scheduling, and *specification-aware*, taking into account  
489 the I/O characteristics of the input networks.

490 The integration of generated HWPE-based accelerators into the OOCF  
491 and the UAV companion computer is performed using a *toolchain* included  
492 in the OODK. A *OOCP specification file* defines the key attributes of the

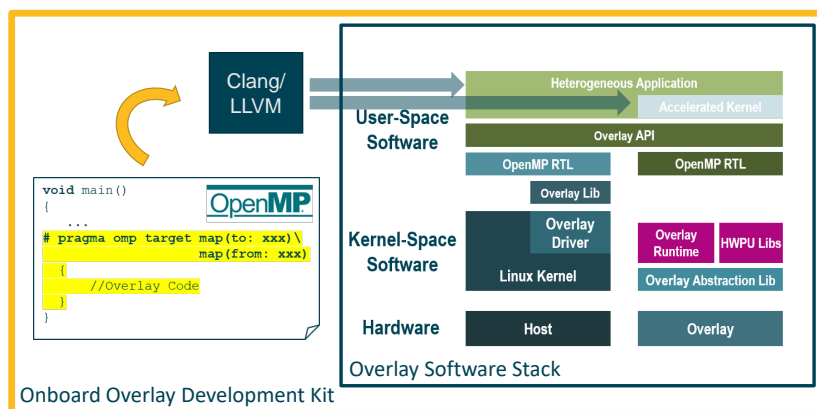


Figure 7: OODK software compilation flow using OpenMP runtime API.

493 accelerator-rich platform. These include the characteristics of the OOCF  
 494 and *cluster* fabrics, as well as the distribution of accelerators across the  
 495 platform. Listing 1 provides an example of the `oocf_specs` class. The OOCF  
 496 (rows 3-7 of Listing 1) method specifies the `name` and `target` for the FPGA  
 497 fabric. `l2_mem` defines instead size and number of ports for L2 memory.  
 498 The `cluster` (rows 9-19 of Listing 1) method specifies the attributes of the  
 499 *cluster fabrics*. An arbitrary number of heterogeneous clusters can be defined,  
 500 from `cluster_0` to `cluster_N`. The `acc` object specifies which accelerators to  
 501 integrate within an *cluster*. Additionally, the values assigned to `proxy`, `dma`,  
 502 and `l1_mem` enable specialization of other *cluster* components. Additional  
 503 parameters specify the number of *proxies* and allow flexible configuration of  
 504 the DMA design. The L1 memory is similarly defined in terms of its size and  
 505 number of ports. Based on these specifications, an accelerator-rich platform is  
 506 automatically generated, with system-level integration of accelerators using  
 507 the generated HWPE. Finally, an FPGA bitstream is built using vendor-  
 508 dependent tools (e.g., AMD Vivado Suite) to implement a specific *overlay*  
 509 instance and complete the SLD design process.

510 Users can execute code on the overlay through a high-level and established  
 511 programming model such as OpenMP. Figure 7 shows the compilation flow  
 512 included in the OODK. By using the OpenMP accelerator model and target  
 513 directives, users can offload portions of execution from the host processor  
 514 to the overlay. Consequently, the accelerator APIs can be used to access  
 515 accelerated functionalities. This high-level programming model proves to be  
 516 much more flexible and agile compared to custom, low-level, vendor-specific

Table 3: FPGA resource availability of the AMD-Xilinx ZU9EG HeSoC fabric.

Resource Type	Availability
Look-Up Tables (LUTs)	274080
Flip Flops (FFs)	548160
Block RAMs (BRAMs)	912
Digital Signal Processors (DSPs)	2520

517 models used in overlays like PYNQ and ESP.

## 518 5. Experimental Evaluation

519 In this section, our methodology is leveraged to design a UAV companion  
 520 computer for the C4D use-case application scenario. Initially, we characterize  
 521 the C4D use-case tasks and inspect those traits that are exploited at merge-  
 522 time. Then, we propose three variants of our FPGA overlay, each integrating  
 523 the C4D tasks; then, we profile performance and area usage to select the most  
 524 suitable architecture at serving the SPA application.

### 525 5.1. Experimental setup

526 We implement OSCP on an AMD ZCU102 board, which integrates a  
 527 Zynq US+ ZU9EG HeSoC. Table 3 reports the resource availability of the  
 528 FPGA fabric, including: (i) Look-Up Tables (LUTs), i.e., reconfigurable digi-  
 529 tal circuits to implement combinational logic functions; (ii) Flip Flops (FFs),  
 530 to save logical states, e.g., for processor and accelerator registers; (iii) Block  
 531 RAMs (BRAMs), that can be cascaded to create large and fast memory ar-  
 532 rays, e.g., SPMs and FIFOs; and (iv) DSPs (Digital Signal Processors), i.e.,  
 533 dedicated DSP slices.

534 Since the ZU9EG HeSoC includes a physical ARM Cortex-A running  
 535 at  $f_{\text{Host}} = 1.2\text{GHz}$ , we leverage it as a host processor. Besides, we deploy  
 536 the accelerator-rich subsystem to the FPGA fabric at a target frequency  
 537  $f_{\text{FPGA}} = 100\text{MHz}$ . The FPGA design is synthesized, implemented and de-  
 538 ployed leveraging the AMD Vivado Suite, including: (i) Vivado HLS v2018.2  
 539 to design the HW-accelerated use-case kernels; (ii) Vivado v2019.2 to syn-  
 540 thesize and implement the FPGA design; and (iii) Petalinux v2019.2 to boot  
 541 the OSCP. Experimental reports discussed in the following sections include  
 542 runtime performance measurements and FPGA resource utilization, retrieved  
 543 after the implementation phase.

Table 4: Overview of the C4D use-case kernels.

	Task	Kernel	Description	I/O
PRE-PROCESSING	<i>Low-pass filter</i>	FIR	Fully-parallel systolic implementation, configurable with 64 or 128 coefficients.	1/1
	<i>Blur</i>	Conv	Dataflow implementation of a Gaussian Convolution with internal stages optimized in a pipeline manner.	1/1
MAIN	<i>Crop detection</i>	CNN	CNN Accelerator, configurable in different modes to accommodate variegated convolutional layers. The implementation is based on previous C4D partners work [46].	2/1
	<i>Soil segmentation</i>	Canny	Canny edge detector, based on a Sobel implementation [58].	2/1
POST-PROCESSING	<i>Data annotation</i>	BBOX	Packing of textual metadata with image and IMU data.	3/1
	<i>Data encryption</i>	AES	AES data encryption.	1/1

544 5.2. Overview of the use-case SPA kernels

545 Acquiring details about the features of the SPA kernels is essential to  
 546 guide the merge process, which is then by different strategies, e.g., the *ap-*  
 547 *plication task scheduling* and the characteristics of the *I/O accelerator inter-*  
 548 *faces*.

549 Table 4 resumes the tasks composing the C4D’s SPA use-case scenario,  
 550 along with a brief functional description. The PRE-PROCESSING and MAIN  
 551 phases are HW-accelerated in the accelerator-rich subsystem, while the POST-  
 552 PROCESSING is run in SW on the host processor. Notably, a companion  
 553 computer targeting SPA applications is called to perform multiple acceler-  
 554 ated functions, during the mission. Moreover, it is worth pointing out that  
 555 MDC supports different design approaches, which have been employed to  
 556 develop these kernels. In particular, the CAPH language is used to develop  
 557 the **Canny** [88], while HLS C/C++ is leveraged to design **FIR**, **Conv** and **CNN**.  
 558 Finally, the last column, on the right, shows the number of I/O interfaces for  
 559 each kernel implementation. In the case of a HW accelerator, this consists  
 560 of the actual I/O ports exposed at the interface of the CGR datapath, while  
 561 for SW implementations, it consists of the number of I/O data arrays.

562 Figure 8 shows the latency of execution of the SPA kernels. Each bar

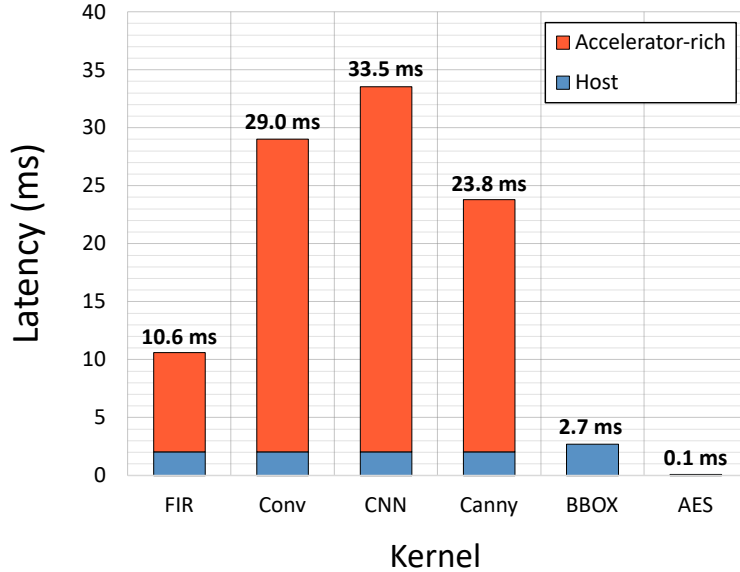


Figure 8: Latency profiling of the use-case kernels.

563 expresses a latency breakdown of the executions on: (i) the host, including  
 564 startup, allocation and offloading; and (ii) the accelerator-rich subsystem,  
 565 including the proxy core startup, DMA data transfers and accelerator exe-  
 566 cution.

### 567 5.3. Assembling an overlay-based UAV companion computer

568 We proceed by assembling *three variants* of our overlay-based UAV com-  
 569 panion computer, both leveraging the proposed methodology, as well as the  
 570 acquired information concerning the SPA kernels. The variants are meant to  
 571 fit the requirements of the C4D’s SPA use-case, described in Section 2, and  
 572 *differ by the adopted merge strategies*, as shown in Table 5: (i) the BASELINE  
 573 instantiates a dedicated HWPE wrapper for each CGR datapath; (ii) the  
 574 SCENARIO-AGNOSTIC strategy groups CGR datapaths by the *number of in-*  
 575 *put accelerator ports*, thus merging **Conv** and **FIR**, as well as **CNN** and **Canny**;  
 576 and (iii) the SCENARIO-AWARE strategy is driven by the *application task*  
 577 *scheduling*, where the goal is to maximize the number of parallel tasks. For  
 578 instance, **CNN** and **Canny** are instantiated in dedicated wrappers to run in par-  
 579 allel, as they retain no mutual dependency and are very latency-expensive.  
 580 Besides, **Conv** and **FIR** are merged, given the former has a dependency on

Table 5: Overview of the adopted merge strategies.

Strategy	Implementation <sup>1</sup>
BASELINE	[FIR] <sub>a</sub> , [Conv] <sub>b</sub> , [CNN] <sub>c</sub> , [Canny] <sub>d</sub>
SCENARIO-AGNOSTIC	[FIR + Conv] <sub>a</sub> , [CNN + Canny] <sub>b</sub>
SCENARIO-AWARE	[FIR + Conv] <sub>a</sub> , [CNN] <sub>b</sub> , [Canny] <sub>c</sub>

<sup>1</sup> SPA kernels are either instantiated in dedicated wrappers [kernel]<sub>i</sub> or merged in a single wrapper  $[\sum \text{kernels}]_i$ .

581 both **CNN** and **Canny**; indeed, the execution of **Canny** can be hidden by running  
 582 it in parallel to the other kernels.


583 The SW application starts on the host processor, which acquires data  
 584 from sensors and then offloads the PRE-PROCESSING and MAIN stages to the  
 585 accelerator-rich subsystem. Once offloading completes, the RISC-V proxy  
 586 core commands the DMA to move input data to local accelerator buffers.  
 587 Data are divided into tiles due to the limited dimension of the L1 memory.  
 588 As soon as the transfer completes, accelerators are launched, according to  
 589 the execution profile of the C4D application. Figure 9a consists of a time  
 590 diagram corresponding to the SCENARIO-AWARE strategy, which guides the  
 591 accelerator optimization given the information about the application task  
 592 scheduling, as well as the BASELINE, where no merging constraints are ap-  
 593 plied to accelerators. Differently, the SCENARIO-AGNOSTIC strategy forces  
 594 **CNN** and **Canny** to run sequentially, being merged in the same wrapper. The  
 595 diagram also includes colour-coded features that describe the number of I/O  
 596 interfaces of the accelerators. **BBOX** and **AES** are not colour-coded, as they  
 597 are executed in SW on the host processor.

### 598 5.3.1. Latency

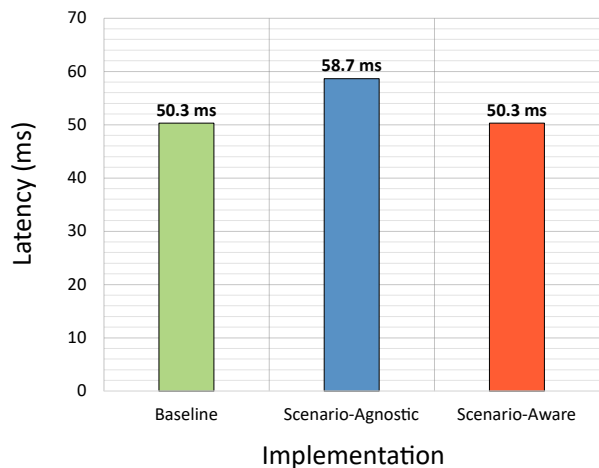
599 Figure 9b shows the resulting execution latency for each variant. The  
 600 SCENARIO-AWARE approach terminates in 50.3ms, while the SCENARIO-  
 601 AGNOSTIC strategy degrades of a factor of 13.7%, as of the serialization of  
 602 the **CNN** and **Canny** tasks. Additionally, the BASELINE does not constrain  
 603 kernels by merging their interfaces; thus, each can operate freely in parallel,  
 604 following the application’s profile. As a consequence, the former performs as  
 605 the SCENARIO-AWARE approach.

### 606 5.3.2. Area usage

607 Table 6 shows the FPGA resource utilization of the BASELINE variant,  
 608 normalized on the resource availability of the ZU9EG HeSoC fabric. Values



(a) Application task scheduling



(b) Execution latency

Figure 9: Application task scheduling and execution latency concerning the three architectural variants of the overlay-based UAV companion computer.

609 are organized following three hierarchical levels: (i) *overlay-level*, referring to  
 610 to an OOCPP instance with a single cluster, the host communication inter-  
 611 face, i.e., the HW mailbox, and the L2 memory subsystem, accounting for  
 612 resource usage of the AXI4 interconnect and the L2 SPM; (ii) *cluster-level*,  
 613 with a RISC-V proxy core, DMA, L1 memory and the cumulative occupation  
 614 of accelerator wrappers and datapaths; and (iii) *accelerator-level*, including  
 615 specific information concerning the HWPE wrappers and CGR datapaths.

616 6

617 The main benefit of the merge process consists of the area savings origi-  
 618 nating from the sharing of the HWPE wrapper, thus we expect a reduction  
 619 of their area usage with a higher employment of the optimization. With a re-  
 620 duction of the wrappers, fewer L1 memory ports are required, thus positively  
 621 affecting the occupation of the cluster LIC (cumulated with the L1 memory).

Table 6: FPGA utilization of the baseline C4D OOC variant.

Overlay-level		LUT%	FF%	BRAM%	DSP%
	Host communication	2.0	0.7	0	0
	L2 memory	5.0	1.4	7.0	0
	Cluster	59.7	23.8	17.4	36.9
Cluster-level		LUT%	FF%	BRAM%	DSP%
	RISC-V soft-core	8.6	1.5	9.6	0.4
	DMA	2.3	1.1	0	0
	L1 memory	11.1	3.7	2.9	0
	Accelerator wrappers	7.6	2.9	0	0
	Accelerator datapaths	30.3	14.7	4.9	36.6
Accelerator-level		LUT%	FF%	BRAM%	DSP%
<b>FIR</b>	Streamer	0.8	0.2	0	0
	Controller	0.4	0.4	0	0
	Datapath	8.6	8.5	0	5.1
<b>Conv</b>	Streamer	1.7	0.4	0	0
	Controller	0.9	0.4	0	0
	Datapath	1.6	0.8	2.4	0.5
<b>CNN</b>	Streamer	1.2	0.4	0	0
	Controller	0.6	0.4	0	0
	Datapath	10.5	1.6	0.1	30.5
<b>Canny</b>	Streamer	1.4	0.3	0	0
	Controller	0.6	0.4	0	0
	Datapath	9.6	3.8	2.4	0.5

NOTE: Normalized FPGA utilization (%) refers to the resource availability of the AMD-Xilinx ZU9EG HeSoC fabric used in our experiments.

622 Thus, LUT would decrease of up to  $\Delta_{\text{LUT,max}} = -18.7\%$ , consisting of about  
 623 28% of the BASELINE overlay. Besides, FF decrease is expected for up to  
 624  $\Delta_{\text{FF,max}} = -6.6\%$ , hence 26% of the BASELINE overlay.

625 Table 7 shows the occupation results of each merging strategy, to con-  
 626 firm these conjectures. Indeed, at the overlay-level, we denote a maximum  
 627 decrement in resource occupation of  $\Delta_{\text{LUT}} = -12.7\%$  and  $\Delta_{\text{FF}} = -4.4\%$ , in  
 628 correspondence with the SCENARIO-AGNOSTIC approach. No variation ap-  
 629 plies to BRAMs and DSPs because the merge does not apply to the CGR  
 630 datapaths or other overlay components making use of those macros. Zoom-  
 631 ing in at the cluster-level, the same decrement is found, as the merge only  
 632 influences the cluster components: the accelerator wrappers and the cluster  
 633 LIC.

Table 7: FPGA utilization of the C4D OOCF variants.

Overlay-level	LUT%	FF%	BRAM%	DSP%
BASILINE	66.7	25.9	24.4	36.9
SCENARIO-AGNOSTIC	54.0	21.5	24.4	36.9
SCENARIO-AWARE	55.7	22.3	24.4	36.9
Cluster-level	LUT%	FF%	BRAM%	DSP%
BASILINE	59.7	23.8	17.4	36.9
SCENARIO-AGNOSTIC	46.9	19.4	17.4	36.9
SCENARIO-AWARE	48.7	20.2	17.4	36.9
Accelerator-level	LUT%	FF%	BRAM%	DSP%
BASILINE	37.8	17.5	4.9	36.5
SCENARIO-AGNOSTIC	34.9	16.1	4.9	36.6
SCENARIO-AWARE	35.2	16.7	4.9	36.6

NOTE: Normalized FPGA utilization (%) refers to the resource availability of the AMD-Xilinx ZU9EG HeSoC fabric used in our experiments.

634 *5.3.3. Performance and area trade-off*

635 We dig deeper into the exploration of the three variants and investi-  
 636 gate aggregated trade-offs, derived from performance and FPGA resource  
 637 utilization—LUT and FF usage. The implementation that better responds  
 638 to trade-offs is determined by the larger value of the aggregated metric, which  
 639 is normalized on the results of the SCENARIO-AWARE strategy

640 Figures 10a and 10b shows the trade-off between performance and LUT  
 641 occupation. The SCENARIO-AWARE strategy better performs than the BASE-  
 642 LINE OOCF by a factor of 16.5% and 18.5%, respectively at overlay- and  
 643 cluster-level. Even though the SCENARIO-AGNOSTIC worsen performance  
 644 compared to the BASELINE—as described in Section 5.3.1—then it still pre-  
 645 vails on the BASELINE by factors of 5.0% and 7.4%. Besides, it is also worth  
 646 noting the magnification of the metric when zooming inside the cluster, again  
 647 explainable with the merge only applying to the cluster components.

648 Figures 10c and 10d shows instead the performance and FF occupation  
 649 trade-off. In this case, the SCENARIO-AWARE approach improves the BASE-  
 650 LINE OOCF by a factor of 14.0% and 15.2%—at overlay- and cluster-level—  
 651 while the SCENARIO-AGNOSTIC by factors of 2.7% and 7.3%.

652 *5.3.4. Scalability analysis*

653 This experiment assesses the scalability of the proposed methodology for  
 654 larger UAV companion computers. We assume that a total of  $N_{acc} = 16$  accel-

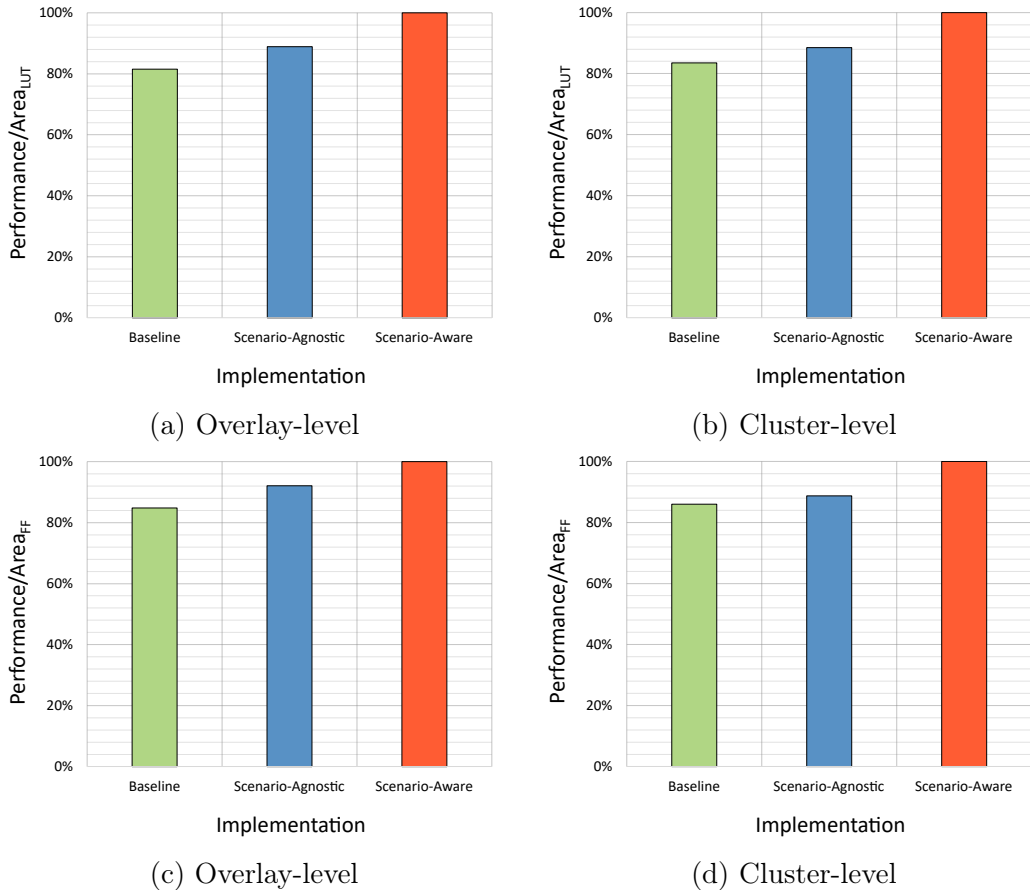


Figure 10: Trade-off between performance and FPGA resource utilization (LUT, FF).

655 erators is sufficient to gain meaningful insights, as commercial HeSoCs tend  
 656 to integrate low-tens of accelerators into shared-memory clusters [89, 90, 91].  
 657 Further scaling introduces challenges due to single-cluster limitations, such  
 658 as contention and interference among shared components. A widely adopted  
 659 solution is to aggregate accelerators into multi-cluster, accelerator-rich archi-  
 660 tectures. This approach reduces pressure on shared cluster components but  
 661 incurs additional FPGA resource overhead, thereby necessitating an investi-  
 662 gation into the scalability of our overlay.

663 To this end, we assemble five multi-cluster OOCF variants using our  
 664 SLD methodology. Each variant includes  $N_{\text{acc}}$  accelerators that are evenly  
 665 distributed across up to 16 clusters. Additionally, we scale the L2 memory

Table 8: Scalability analysis of the proposed FPGA overlay.

Overlay <sup>1</sup>	N <sub>CL</sub> <sup>2</sup>	N <sub>L2P</sub> <sup>3</sup>	LUT		FF		BRAM		DSP	
			Abs <sup>4</sup>	Rel <sup>5</sup>	Abs	Rel	Abs	Rel	Abs	Rel
1CL	1	1	89495	1.0	56395	1.0	184	1.0	9	1.0
2CL	2	2	129342	1.45	98755	1.75	368	2.0	18	2.0
4CL	4	4	222986	2.49	173135	3.07	736	4.0	36	4.0
8CL	8	8	407514	4.55	314644	5.58	1472	8.0	72	8.0
16CL	16	16	899865	10.05	590244	10.47	2944	16.0	144	16.0

<sup>1</sup> **Overlay**: Multi-cluster variants of the FPGA overlay.

<sup>2</sup> **N<sub>CL</sub>**: Number of clusters.

<sup>3</sup> **N<sub>L2P</sub>**: Number of L2 memory ports.

<sup>4</sup> **Abs**: Absolute FPGA utilization.

<sup>5</sup> **Rel**: Relative FPGA utilization normalized to the absolute value of the 1CL variant.

666 subsystem according to the number of clusters ( $N_{CL}$ ), i.e., we expose memory  
 667 ports and banks ( $N_{L2P}$ ) enabling each cluster to access data and instructions  
 668 independently. We scale  $N_{CL}$  and  $N_{L2P}$  simultaneously. The largest variant  
 669 is 16CL, with each accelerator integrated into a dedicated cluster.

670 Table 8 reports absolute and relative values from the experiment. Relative  
 671 values are normalized to the value of the 1CL variant to gain insights into  
 672 multi-cluster scalability relative to a baseline, i.e., the single-cluster instance  
 673 of our FPGA overlay. These values represent the overlay implementation  
 674 costs, excluding any application-specific accelerator datapath or wrapper. In  
 675 principle, BRAMs and DSPs scale poorly, exposing a linear trend correlated  
 676 with  $N_{CL}$  because of the additional L2 memory ports and proxy cores, i.e., 1  
 677 per cluster. In contrast, LUTs and FFs scale a bit better — around  $10\times$  when  
 678 16CL are instantiated. Re-aggregating accelerators into multiple clusters  
 679 reduces the cost per cluster, as fewer LIC accelerator ports are exposed, and  
 680 some components, such as those dedicated to host communication, are not  
 681 replicated when scaling, as previously shown in Table 6. However, a critical  
 682 contributor to LUTs and FFs is the fully-connected AXI4 XBAR at HeSoC-  
 683 level. To this end, it is worth noting that we are working on extending the IP  
 684 portfolio of our FPGA overlay with a Network-on-Chip (NoC) architecture  
 685 to enhance the scaling capability of our methodology [92].

686 The former discussion provides insights into the implementation cost of an  
 687 empty overlay, i.e., the “skeleton” of a multi-cluster, accelerator-rich HeSoC.  
 688 We conclude by presenting an additional analysis that integrates the previ-  
 689 ous findings. Figure 11 shows five example accelerators from the literature,  
 690 including: (i) DPU-B512, AMD Deep Learning Processor Unit (DPU) for  
 691 Convolutional Neural Network (CNN) acceleration (B512 configuration) [93];




Figure 11: Examples of five accelerators from the literature, including: (i) DPU-B512, AMD DPU for CNN acceleration (B512 configuration) [93]; (ii) AD-Lane, CV accelerator-rich pipeline for lane detection [94]; (iii) AD-PF, PF to accelerate the localization component in a perception-plan-act AD stack [80]; (iv) DPU-B4096, AMD DPU for CNN acceleration (B4096 configuration) [93]; and (v) C4D-UAV, our accelerator-rich use-case. We analyze the FPGA resource usage of each accelerator, normalized over a single cluster overlay resource consumption. Results are presented using two vertical axes: the *left axis* for LUTs, FFs, BRAMs, and the *right axis* for DSPs.

692 (ii) AD-Lane, Computer Vision (CV) accelerator-rich pipeline for lane detec-  
 693 tion [94]; (iii) AD-PF, Particle Filter (PF) to accelerate the localization com-  
 694 ponent in a perception-plan-act Autonomous Driving (AD) stack [80]; (iv)  
 695 DPU-B4096, AMD DPU for CNN acceleration (B4096 configuration) [93]; and  
 696 (v) C4D-UAV, our UAV companion computer for SPA applications imple-  
 697 menting the SCENARIO-AWARE strategy, described in Section 5.3. As a side  
 698 note, more details concerning AD-Lane and AD-PF are also found in Section 6.

699 We evaluate the FPGA resource usage of each accelerator, normalized  
 700 to the resource consumption of a single cluster overlay.<sup>4</sup> This normaliza-  
 701 tion enables a meaningful comparison of the FPGA resources required to  
 702 implement real-world accelerators, including those from our accelerator-rich  
 703 use-case scenario. The results intuitively reveal how efficiently resources are  
 704 utilized: higher normalized values suggest more effective use of FPGA re-  
 705 sources by the accelerator itself, while lower values imply greater overhead  
 706 introduced by the overlay. However, it is important to emphasize that the

<sup>4</sup>A normalized value of 1.0 indicates that the accelerator uses the same amount of resources as the overlay; a value of 2.0 indicates twice the resource usage.

707 overlay cluster should not be regarded as overhead per se. As discussed  
708 in Section 4.1, the cluster and its components are fundamental to building  
709 high-performance, accelerator-rich systems [95, 36, 37].

710 The results are presented using two vertical axes: the left axis shows LUT,  
711 FF, and BRAM utilization, while the right axis displays DSP usage. This  
712 dual-axis format improves readability, as the normalized values for these two  
713 groups of FPGA resources differ significantly in scale.

714 It can be observed that, for simple accelerators and pipelines such as  
715 AccDpuSmall and AD-Lane, the resource overhead introduced by the overlay  
716 is comparable, and in some cases equivalent ( $\approx 1\times$ ), to that of individual  
717 application-specific datapaths. This indicates that, even in fine-grained sce-  
718 narios, the cost associated with the overlay is not prohibitive and does not  
719 hinder its practical adoption.

720 More notably, in more complex configurations—such as DPU-B4096, AD-PF,  
721 and our **C4D-UAV**—the resource impact of the overlay becomes virtually  
722 negligible, with overheads ranging from slightly above  $2\times$  to up to  $6\times$  across  
723 all resource classes.

## 724 6. Discussion

725 In this section, we discuss further application scenarios enabled with our  
726 SLD methodology. These demonstrate its broader applicability, not limited  
727 to the SPA context. In particular, we focus on two AD use-cases concerning  
728 F1TENTH autonomous racing competitions [96].

729 F1TENTH is an open-source platform for autonomous racing with 1:10  
730 scaled cars for safe and rapid experimentation of AD systems and algorithms.  
731 Given their small size, these cars mount very constrained computing plat-  
732 forms, thus mandating deep AD stack optimization to attain target Key Per-  
733 formance Indicators (KPIs), e.g., performance, area, and energy consump-  
734 tion.

735 These examples demonstrate that our FPGA overlay, with its automation  
736 flow and proxy core availability, streamlines the integration of application-  
737 specific accelerators, as well as the HW/SW partitioning and optimization  
738 of complex applications deployed to FPGA-based HeSoCs.

### 739 6.1. AD-Lane: A computer vision pipeline for lane detection

740 The first use-case concerns AD-Lane, a CV pipeline for *lane detection*  
741 that autonomously identifies road lane markings and guarantees that the

742 AD system follows the correct path [94].

743 Using the same methodology applied to the UAV companion computer,  
744 we examine the pipeline stages and accelerate the most critical stages with  
745 the AMD Vitis Vision HLS Library [97]. Non-accelerated tasks are executed  
746 in SW, given the limited capability of the AMD Kria KV260 board. We  
747 integrate **AD-Lane** using our FPGA overlay and explore the benefits of the  
748 following HW/SW optimizations: (i) the *proxy core* improves the perfor-  
749 mance of accelerators with many input parameters with respect to looser  
750 and costly interactions with the host processor; (ii) the *L1 memory* benefits  
751 memory-bound stages, given that a tightly-coupled accelerator SPM pro-  
752 vides lower access latency and better predictability compared to the off-chip  
753 main memory; and (iii) *pipelining* the memory and compute phases accord-  
754 ing to a double buffering memory scheme, orchestrated by the proxy core,  
755 improves overall performance compared to sequential execution. Ultimately,  
756 we achieve a  $22\times$  speedup, compared to a SW-only implementation.

## 757 6.2. AD-PF: Real-time LiDAR-based autonomous localization

758 The second use-case concerns AD-PF, a typical perception-plan-act AD  
759 stack, where we accelerate the localization component, a Monte-Carlo method  
760 named *PF* [80]. We leverage the PF to localize the car in a pre-built map  
761 using perception data from a LiDAR sensor, while the optimal car trajec-  
762 tory is computed offline and the control loop is closed with a pure pursuit  
763 algorithm. As the application concerns racing vehicles, HW/SW optimiza-  
764 tions aim to maximize the average speed on single laps and peak speed in  
765 head-to-head situations. Experiments are conducted on real-life race tracks  
766 with two boards (AMD ZCU102, Avnet Ultra96). We examine the impact  
767 of the single pipeline stages and accelerate the most critical one, i.e., the  
768 *Ray Marching (RM)* stage that takes up to 90% of the overall PF execution.  
769 Hence, we design a *Ray Marching Engine (RME)* with HLS, that achieves a  
770  $2.5\times$  speedup compared to a pure SW implementation. Performance is fur-  
771 ther improved by a factor of  $2 - 3\times$  integrating the RME within our FPGA  
772 overlay and exploiting the *proxy core* to optimize the memory transfers be-  
773 tween the HeSoC subsystems—host and device.

## 774 7. Conclusions

775 In this work, we have contributed an innovative SLD methodology for  
776 the design of overlay-based UAV companion computers. In particular, the

777 methodology tackles the challenge of facilitating the assembling of full-fledged  
778 UAV companion computers, targeting *SPA* applications, as is one of the use-  
779 case scenarios of the EU-funded C4D research project.

780 The workflow consists of a modular and scalable accelerator-rich RISC-V  
781 HeSoC, a heterogeneous SW stack to streamline the development of the het-  
782 erogeneous application and a toolflow to automatically generate and integrate  
783 application-specific CGR accelerators inside the FPGA overlay.

784 We show the results of extensive benchmarking, based on three alternative  
785 variants of the OOCF architecture. Experimental results are retrieved by  
786 executing the C4D use-case application on a COTS FPGA-based HeSoC,  
787 belonging to the AMD Zynq US+ family. These demonstrate improvements  
788 in performance and area usage, up to 18.5%.

789 The proposed methodology not only demonstrates the integration capa-  
790 bilities of accelerator-rich systems, but also paves the way for new research  
791 directions. First, the architectural template can be further extended with  
792 additional IPs to achieve better scalability in increasingly accelerator-dense  
793 environments. In this regard, we are already working on expanding the IP  
794 portfolio of our overlay with FlocNoC<sup>5</sup> [92]. A second promising direction is  
795 to investigate and integrate automated design space exploration methods for  
796 optimal selection of integration and merging strategies for accelerators.

## 797 References

- 798 [1] R. Gebbers, V. I. Adamchuk, Precision agriculture and food security,  
799 *Science* 327 (5967) (2010) 828–831.
- 800 [2] R. Bongiovanni, J. Lowenberg-DeBoer, Precision agriculture and sus-  
801 tainability, *Precision agriculture* 5 (2004) 359–387.
- 802 [3] A. Monteiro, S. Santos, P. Gonçalves, Precision agriculture for crop and  
803 livestock farming—brief review, *Animals* 11 (8) (2021) 2345.
- 804 [4] H. Azadi, S. M. Moghaddam, S. Burkart, H. Mahmoudi, S. Van Pas-  
805 sel, A. Kurban, D. Lopez-Carr, Rethinking resilient agriculture: From  
806 climate-smart agriculture to vulnerable-smart agriculture, *Journal of*  
807 *Cleaner Production* 319 (2021) 128602.

---

<sup>5</sup><https://github.com/pulp-platform/FlocNoC>

- 808 [5] N. Ahmed, D. De, I. Hussain, Internet of things (iot) for smart precision  
809 agriculture and farming in rural areas, *IEEE internet of things journal*  
810 5 (6) (2018) 4890–4899.
- 811 [6] T. A. Shaikh, T. Rasool, F. R. Lone, Towards leveraging the role of  
812 machine learning and artificial intelligence in precision agriculture and  
813 smart farming, *Computers and Electronics in Agriculture* 198 (2022)  
814 107119.
- 815 [7] H. Yin, Y. Cao, B. Marelli, X. Zeng, A. J. Mason, C. Cao, Soil sen-  
816 sors and plant wearables for smart and precision agriculture, *Advanced*  
817 *Materials* 33 (20) (2021) 2007764.
- 818 [8] D. Xie, L. Chen, L. Liu, L. Chen, H. Wang, Actuators and sensors for  
819 application in agricultural robots: A review, *Machines* 10 (10) (2022)  
820 913.
- 821 [9] P. K. R. Maddikunta, S. Hakak, M. Alazab, S. Bhattacharya, T. R.  
822 Gadekallu, W. Z. Khan, Q.-V. Pham, Unmanned aerial vehicles in smart  
823 agriculture: Applications, requirements, and challenges, *IEEE Sensors*  
824 *Journal* 21 (16) (2021) 17608–17619.
- 825 [10] E. Karunathilake, A. T. Le, S. Heo, Y. S. Chung, S. Mansoor, The path  
826 to smart farming: Innovations and opportunities in precision agriculture,  
827 *Agriculture* 13 (8) (2023) 1593.
- 828 [11] J. Liu, J. Xiang, Y. Jin, R. Liu, J. Yan, L. Wang, Boost precision agricul-  
829 ture with unmanned aerial vehicle remote sensing and edge intelligence:  
830 A survey, *Remote Sensing* 13 (21) (2021) 4387.
- 831 [12] Y. Kalyani, R. Collier, A systematic survey on the role of cloud, fog,  
832 and edge computing combination in smart agriculture, *Sensors* 21 (17)  
833 (2021) 5922.
- 834 [13] M. P. Christiansen, M. S. Laursen, R. N. Jørgensen, S. Skovsen, R. Gis-  
835 lum, Designing and testing a uav mapping system for agricultural field  
836 surveying, *Sensors* 17 (12) (2017) 2703.
- 837 [14] P. Katsigiannis, L. Misopolinos, V. Liakopoulos, T. K. Alexandridis,  
838 G. Zalidis, An autonomous multi-sensor uav system for reduced-input

- 839 precision agriculture applications, in: 2016 24th Mediterranean Confer-  
840 ence on Control and Automation (MED), IEEE, 2016, pp. 60–64.
- 841 [15] J. Primicerio, S. F. Di Gennaro, E. Fiorillo, L. Genesisio, E. Lugato,  
842 A. Matese, F. P. Vaccari, A flexible unmanned aerial vehicle for precision  
843 agriculture, *Precision Agriculture* 13 (4) (2012) 517–523.
- 844 [16] D. Madroñal, F. Palumbo, A. Capotondi, A. Marongiu, Unmanned ve-  
845 hicles in smart farming: A survey and a glance at future horizons, in:  
846 Proceedings of the 2021 Drone Systems Engineering and Rapid Simu-  
847 lation and Performance Evaluation: Methods and Tools Proceedings,  
848 2021, pp. 1–8.
- 849 [17] Y. Liu, X. Ma, L. Shu, G. P. Hancke, A. M. Abu-Mahfouz, From in-  
850 dustry 4.0 to agriculture 4.0: Current status, enabling technologies, and  
851 research challenges, *IEEE transactions on industrial informatics* 17 (6)  
852 (2020) 4322–4334.
- 853 [18] J. Chen, X. Ran, Deep learning with edge computing: A review, *Pro-  
854 ceedings of the IEEE* 107 (8) (2019) 1655–1674.
- 855 [19] H. A. Alharbi, M. Aldossary, Energy-efficient edge-fog-cloud architec-  
856 ture for iot-based smart agriculture environment, *Ieee Access* 9 (2021)  
857 110480–110492.
- 858 [20] J. Su, X. Zhu, S. Li, W.-H. Chen, Ai meets uavs: A survey on ai empow-  
859 ered uav perception systems for precision agriculture, *Neurocomputing*  
860 518 (2023) 242–270.
- 861 [21] P. P. Ray, A review on tinyml: State-of-the-art and prospects, *Journal of  
862 King Saud University-Computer and Information Sciences* 34 (4) (2022)  
863 1595–1623.
- 864 [22] P. Radoglou-Grammatikis, P. Sarigiannidis, T. Lagkas, I. Moscholios,  
865 A compilation of uav applications for precision agriculture, *Computer  
866 Networks* 172 (2020) 107148.
- 867 [23] L. Valente, A. Nadalini, A. Veeran, M. Sinigaglia, B. Sa, N. Wistoff,  
868 Y. Tortorella, S. Benatti, R. Psiakis, A. Kulmala, et al., A heteroge-  
869 neous risc-v based soc for secure nano-uav navigation, arXiv preprint  
870 arXiv:2401.03531.

- 871 [24] A. Di Mauro, M. Scherer, D. Rossi, L. Benini, Kraken: A direct  
872 event/frame-based multi-sensor fusion soc for ultra-efficient visual pro-  
873 cessing in nano-uavs, arXiv preprint arXiv:2209.01065.
- 874 [25] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld,  
875 T. Laengle, G. Cioffi, Y. Song, A. Loquercio, et al., Agilicious: Open-  
876 source and open-hardware agile quadrotor for vision-based flight, *Science*  
877 *robotics* 7 (67) (2022) eabl6259.
- 878 [26] A. Douklias, L. Karagiannidis, F. Misichroni, A. Amditis, Design and  
879 implementation of a uav-based airborne computing platform for com-  
880 puter vision and machine learning applications, *Sensors* 22 (5) (2022)  
881 2049.
- 882 [27] D. C. Tsouros, S. Bibi, P. G. Sarigiannidis, A review on uav-based ap-  
883 plications for precision agriculture, *Information* 10 (11) (2019) 349.
- 884 [28] I. Munasinghe, A. Perera, R. C. Deo, A comprehensive review of uav-  
885 ugv collaboration: Advancements and challenges, *Journal of Sensor and*  
886 *Actuator Networks* 13 (6) (2024) 81.
- 887 [29] M. Mammarella, L. Comba, A. Biglia, F. Dabbene, P. Gay, Cooperation  
888 of unmanned systems for agricultural applications: A case study in a  
889 vineyard, *biosystems engineering* 223 (2022) 81–102.
- 890 [30] J. Chen, X. Zhang, B. Xin, H. Fang, Coordination between unmanned  
891 aerial and ground vehicles: A taxonomy and optimization perspective,  
892 *IEEE transactions on cybernetics* 46 (4) (2015) 959–972.
- 893 [31] S. J. Undertaking, et al., European drones outlook study: unlocking the  
894 value for europe.
- 895 [32] R. Nouacer, M. Hussein, H. Espinoza, Y. Ouhammou, M. Ladeira,  
896 R. Castiñeira, Towards a framework of key technologies for drones, *Mi-*  
897 *croprocessors and Microsystems* 77 (2020) 103142.
- 898 [33] N. Aliane, A survey of open-source uav autopilots, *Electronics* 13 (23)  
899 (2024) 4785.
- 900 [34] E. Ebeid, M. Skriver, J. Jin, A survey on open-source flight control  
901 platforms of unmanned aerial vehicle, in: *2017 euromicro conference on*  
902 *digital system design (dsd)*, IEEE, 2017, pp. 396–402.

- 903 [35] A. Wilson, A. Kumar, A. Jha, L. R. Cenkeramaddi, Embedded sensors,  
904 communication technologies, computing platforms and machine learning  
905 for uavs: A review, *IEEE Sensors Journal* 22 (3) (2021) 1807–1826.
- 906 [36] J. Cong, M. A. Ghodrati, M. Gill, B. Grigorian, K. Gururaj, G. Reinman,  
907 Accelerator-rich architectures: Opportunities and progresses, in: 2014  
908 51st ACM/EDAC/IEEE Design Automation Conference (DAC), IEEE,  
909 2014, pp. 1–6.
- 910 [37] M. J. Lyons, M. Hempstead, G.-Y. Wei, D. Brooks, The accelerator  
911 store: A shared memory framework for accelerator-based systems, *ACM*  
912 *Transactions on Architecture and Code Optimization (TACO)* 8 (4)  
913 (2012) 1–22.
- 914 [38] Z. Wan, A. Lele, B. Yu, S. Liu, Y. Wang, V. J. Reddi, C. Hao, A. Ray-  
915 chowdhury, Robotic computing on fpgas: Current progress, research  
916 challenges, and opportunities, arXiv preprint arXiv:2205.07149.
- 917 [39] Y. Chi, W. Qiao, A. Sohrabizadeh, J. Wang, J. Cong, Democratizing  
918 domain-specific computing, *Communications of the ACM* 66 (1) (2022)  
919 74–85.
- 920 [40] L. P. Carloni, The case for embedded scalable platforms, in: *Proceedings*  
921 *of the 53rd Annual Design Automation Conference*, 2016, pp. 1–6.
- 922 [41] A. Amid, D. Biancolin, A. Gonzalez, D. Grubb, S. Karandikar, H. Liew,  
923 A. Magyar, H. Mao, A. Ou, N. Pemberton, et al., Chipyard: Inte-  
924 grated design, simulation, and implementation framework for custom  
925 socs, *IEEE Micro* 40 (4) (2020) 10–21.
- 926 [42] M. H. Quraishi, E. B. Tavakoli, F. Ren, A survey of system architectures  
927 and techniques for fpga virtualization, *IEEE Transactions on Parallel*  
928 *and Distributed Systems* 32 (9) (2021) 2216–2230.
- 929 [43] G. Bellocchi, A. Capotondi, F. Conti, A. Marongiu, A risc-v-based fpga  
930 overlay to simplify embedded accelerator deployment, in: 2021 24th  
931 *Euromicro Conference on Digital System Design (DSD)*, IEEE, 2021,  
932 pp. 9–17.
- 933 [44] A. Lambertini, E. Mandanici, M. A. Tini, L. Vittuari, Technical chal-  
934 lenges for multi-temporal and multi-sensor image processing surveyed

- 935 by uav for mapping and monitoring in precision agriculture, *Remote*  
936 *Sensing* 14 (19) (2022) 4954.
- 937 [45] C. Chen, Z. Zheng, T. Xu, S. Guo, S. Feng, W. Yao, Y. Lan, Yolo-based  
938 uav technology: A review of the research and its applications, *Drones*  
939 7 (3) (2023) 190.
- 940 [46] A. Sassu, J. Motta, A. Deidda, L. Ghiani, A. Carlevaro, G. Garibotto,  
941 F. Gambella, Artichoke deep learning detection network for site-specific  
942 agrochemicals uas spraying, *Computers and Electronics in Agriculture*  
943 213 (2023) 108185.
- 944 [47] F. A. Almalki, B. O. Soufiene, S. H. Alsamhi, H. Sakli, A low-cost  
945 platform for environmental smart farming monitoring system based on  
946 iot and uavs, *Sustainability* 13 (11) (2021) 5908.
- 947 [48] C. Qu, J. Boubin, D. Gafurov, J. Zhou, N. Aloysius, H. Nguyen,  
948 P. Calyam, Uav swarms in smart agriculture: Experiences and oppor-  
949 tunities, in: *2022 IEEE 18th International Conference on e-Science (e-*  
950 *Science)*, IEEE, 2022, pp. 148–158.
- 951 [49] A. I. de Castro, Y. Shi, J. M. Maja, J. M. Peña, Uavs for vegetation  
952 monitoring: Overview and recent scientific contributions, *Remote Sens-*  
953 *ing* 13 (11) (2021) 2139.
- 954 [50] A. Abbas, Z. Zhang, H. Zheng, M. M. Alami, A. F. Alrefaei, Q. Abbas,  
955 S. A. H. Naqvi, M. J. Rao, W. F. Mosa, Q. Abbas, et al., Drones in plant  
956 disease assessment, efficient monitoring, and detection: a way forward  
957 to smart agriculture, *Agronomy* 13 (6) (2023) 1524.
- 958 [51] N. Kitpo, M. Inoue, Early rice disease detection and position mapping  
959 system using drone and iot architecture, in: *2018 12th South East Asian*  
960 *Technical University Consortium (SEATUC)*, Vol. 1, IEEE, 2018, pp. 1–  
961 5.
- 962 [52] J. Su, C. Liu, M. Coombes, X. Hu, C. Wang, X. Xu, Q. Li, L. Guo, W.-  
963 H. Chen, Wheat yellow rust monitoring by learning from multispectral  
964 uav aerial imagery, *Computers and electronics in agriculture* 155 (2018)  
965 157–166.

- 966 [53] T. Talaviya, D. Shah, N. Patel, H. Yagnik, M. Shah, Implementation  
967 of artificial intelligence in agriculture for optimisation of irrigation and  
968 application of pesticides and herbicides, *Artificial Intelligence in Agri-*  
969 *culture* 4 (2020) 58–73.
- 970 [54] B. S. Faiçal, H. Freitas, P. H. Gomes, L. Y. Mano, G. Pessin, A. C.  
971 de Carvalho, B. Krishnamachari, J. Ueyama, An adaptive approach for  
972 uav-based pesticide spraying in dynamic environments, *Computers and*  
973 *Electronics in Agriculture* 138 (2017) 210–223.
- 974 [55] D. Debnath, F. Vanegas, J. Sandino, A. F. Hawary, F. Gonzalez, A  
975 review of uav path-planning algorithms and obstacle avoidance methods  
976 for remote sensing applications, *Remote Sensing* 16 (21) (2024) 4019.
- 977 [56] S. Ahmed, B. Qiu, F. Ahmad, C.-W. Kong, H. Xin, A state-of-the-  
978 art analysis of obstacle avoidance methods from the perspective of an  
979 agricultural sprayer uav’s operation scenario, *Agronomy* 11 (6) (2021)  
980 1069.
- 981 [57] L. Wang, Y. Lan, Y. Zhang, H. Zhang, M. N. Tahir, S. Ou, X. Liu,  
982 P. Chen, Applications and prospects of agricultural unmanned aerial  
983 vehicle obstacle avoidance technology in china, *Sensors* 19 (3) (2019)  
984 642.
- 985 [58] C.-L. Sotiropoulou, C. Gentsos, S. Nikolaidis, A. Rjoub, Fpga-based  
986 canny edge detection for real-time applications, in: published at the  
987 26th Conference on Design of Circuits and Integrated Systems (DCIS),  
988 Albufeira, Portugal, 2011.
- 989 [59] M. Bouhali, F. Shamani, Z. E. Dahmane, A. Belaidi, J. Nurmi, Fpga  
990 applications in unmanned aerial vehicles-a review, in: *Applied Recon-*  
991 *figurible Computing: 13th International Symposium, ARC 2017, Delft,*  
992 *The Netherlands, April 3-7, 2017, Proceedings 13, Springer, 2017, pp.*  
993 *217–228.*
- 994 [60] B. H. Y. Alsalam, K. Morton, D. Campbell, F. Gonzalez, Autonomous  
995 uav with vision based on-board decision making for remote sensing and  
996 precision agriculture, in: *2017 IEEE Aerospace Conference, IEEE, 2017,*  
997 *pp. 1–12.*

- 998 [61] B. Dai, Y. He, F. Gu, L. Yang, J. Han, W. Xu, A vision-based au-  
999 tonomous aerial spray system for precision agriculture, in: 2017 IEEE  
1000 International Conference on Robotics and Biomimetics (ROBIO), IEEE,  
1001 2017, pp. 507–513.
- 1002 [62] P. Horstrand, R. Guerra, A. Rodríguez, M. Díaz, S. López, J. F. López,  
1003 A UAV platform based on a hyperspectral sensor for image capturing and  
1004 on-board processing, *IEEE Access* 7 (2019) 66919–66938.
- 1005 [63] A. Saddik, R. Latif, A. El Ouardi, Low-power FPGA architecture based  
1006 monitoring applications in precision agriculture, *Journal of Low Power  
1007 Electronics and Applications* 11 (4) (2021) 39.
- 1008 [64] N. H. Malle, F. F. Nyboe, E. S. M. Ebeid, Onboard powerline perception  
1009 system for UAVs using mmwave radar and FPGA-accelerated vision, *IEEE  
1010 Access* 10 (2022) 113543–113559.
- 1011 [65] K. Guo, S. Zeng, J. Yu, Y. Wang, H. Yang, A survey of FPGA-based  
1012 neural network accelerator, arXiv preprint arXiv:1712.08934.
- 1013 [66] A. K. Jain, X. Li, P. Singhai, D. L. Maskell, S. A. Fahmy, Deco: A  
1014 DSP block based FPGA accelerator overlay with low overhead intercon-  
1015 nect, in: 2016 IEEE 24th Annual International Symposium on Field-  
1016 Programmable Custom Computing Machines (FCCM), IEEE, 2016, pp.  
1017 1–8.
- 1018 [67] PYNQ - Python productivity for Zynq.  
1019 URL <http://www.pynq.io/>
- 1020 [68] V. Mayoral-Vilches, S. M. Neuman, B. Plancher, V. J. Reddi, Robot-  
1021 core: An open architecture for hardware acceleration in ROS 2, in: 2022  
1022 IEEE/RSJ International Conference on Intelligent Robots and Systems  
1023 (IROS), IEEE, 2022, pp. 9692–9699.
- 1024 [69] D. Giri, K.-L. Chiu, G. Eichler, P. Mantovani, L. P. Carloni, Accelerator  
1025 integration for open-source SoC design, *IEEE Micro* 41 (4) (2021) 8–14.
- 1026 [70] C. Heinz, J. Hofmann, J. Korinth, L. Sommer, L. Weber, A. Koch,  
1027 The Tapasco open-source toolflow: For the automated composition of  
1028 task-based parallel reconfigurable computing systems, *Journal of Signal  
1029 Processing Systems* 93 (2021) 545–563.

- 1030 [71] B. Hutchings, M. Wirthlin, Rapid implementation of a partially reconfig-  
1031 urable video system with pynq, in: 2017 27th International Conference  
1032 on Field Programmable Logic and Applications (FPL), IEEE, 2017, pp.  
1033 1–8.
- 1034 [72] The Composable Video Pipeline — PYNQ Composable Overlays 1.0.2  
1035 documentation.  
1036 URL [https://pynq-composable.readthedocs.io/en/latest/  
1037 video\\_pipeline.html](https://pynq-composable.readthedocs.io/en/latest/video_pipeline.html)
- 1038 [73] J. Goeders, T. Gaskin, B. Hutchings, Demand driven assembly of  
1039 fpga configurations using partial reconfiguration, ubuntu linux, and  
1040 pynq, in: 2018 IEEE 26th Annual International Symposium on Field-  
1041 Programmable Custom Computing Machines (FCCM), IEEE, 2018, pp.  
1042 149–156.
- 1043 [74] G. Tagliavini, D. Cesarini, A. Marongiu, Unleashing fine-grained par-  
1044 allelism on embedded many-core accelerators with lightweight openmp  
1045 tasking, IEEE Transactions on Parallel and Distributed Systems 29 (9)  
1046 (2018) 2150–2163.
- 1047 [75] A. Kurth, A. Capotondi, P. Vogel, L. Benini, A. Marongiu, Hero: An  
1048 open-source research platform for hw/sw exploration of heterogeneous  
1049 manycore systems, in: Proceedings of the 2nd Workshop on AutotuniNg  
1050 and aDaptivity AppRoaches for Energy efficient HPC Systems, 2018, pp.  
1051 1–6.
- 1052 [76] D. Rossi, F. Conti, A. Marongiu, A. Pullini, I. Loi, M. Gautschi,  
1053 G. Tagliavini, A. Capotondi, P. Flatresse, L. Benini, Pulp: A paral-  
1054 lel ultra low power platform for next generation iot applications, in:  
1055 2015 IEEE Hot Chips 27 Symposium (HCS), IEEE, 2015, pp. 1–39.
- 1056 [77] Zynq UltraScale+ MPSoC.  
1057 URL [https://www.xilinx.com/products/silicon-devices/soc/  
1058 zynq-ultrascale-mpsoc.html](https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html)
- 1059 [78] V. Boppana, S. Ahmad, I. Ganusov, V. Kathail, V. Rajagopalan, R. Wit-  
1060 tigt, Ultrascale+ mpsoC and fpga families, in: 2015 IEEE Hot Chips 27  
1061 Symposium (HCS), IEEE, 2015, pp. 1–37.

- 1062 [79] F. Zaruba, L. Benini, The cost of application-class processing: Energy  
1063 and performance analysis of a linux-ready 1.7-GHz 64-bit RISC-v core  
1064 in 22-nm FDSOI technology 27 (11) 2629–2640. doi:10.1109/TVLSI.  
1065 2019.2926114.
- 1066 [80] A. Bernardi, G. Brilli, A. Capotondi, A. Marongiu, P. Burgio, An fpga  
1067 overlay for efficient real-time localization in 1/10th scale autonomous  
1068 vehicles, in: 2022 Design, Automation & Test in Europe Conference &  
1069 Exhibition (DATE), IEEE, 2022, pp. 915–920.
- 1070 [81] G. Valente, G. Brilli, T. Di Mascio, A. Capotondi, P. Burgio, P. Valente,  
1071 A. Marongiu, Fine-grained qos control via tightly-coupled bandwidth  
1072 monitoring and regulation for fpga-based heterogeneous socs, IEEE  
1073 Transactions on Parallel and Distributed Systems.
- 1074 [82] P. D. Schiavone, F. Conti, D. Rossi, M. Gautschi, A. Pullini, E. Fla-  
1075 mand, L. Benini, Slow and steady wins the race? a comparison of ultra-  
1076 low-power risc-v cores for internet-of-things applications, in: 2017 27th  
1077 International Symposium on Power and Timing Modeling, Optimization  
1078 and Simulation (PATMOS), IEEE, 2017, pp. 1–8.
- 1079 [83] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi,  
1080 E. Flamand, F. K. Gürkaynak, L. Benini, Near-threshold RISC-v core  
1081 with DSP extensions for scalable IoT endpoint devices 25 (10) 2700–  
1082 2713. doi:10.1109/TVLSI.2017.2654506.
- 1083 [84] F. Conti, L. Benini, A ultra-low-energy convolution engine for fast brain-  
1084 inspired vision in multicore clusters, in: 2015 Design, Automation Test  
1085 in Europe Conference Exhibition (DATE), pp. 683–688, ISSN: 1558-  
1086 1101. doi:10.7873/DATE.2015.0404.
- 1087 [85] F. Conti, P. D. Schiavone, L. Benini, XNOR neural engine: A hardware  
1088 accelerator IP for 21.6-fJ/op binary neural network inference 37 (11)  
1089 2940–2951. doi:10.1109/TCAD.2018.2857019.
- 1090 [86] C. Sau, T. Fanni, C. Rubattu, L. Raffo, F. Palumbo, The multi-dataflow  
1091 composer tool: An open-source tool suite for optimized coarse-grain re-  
1092 configurable hardware accelerators and platform design, Microprocessors  
1093 and Microsystems 80 (2021) 103326.

- 1094 [87] T. Fanni, A. Rodríguez, C. Sau, L. Suriano, F. Palumbo, L. Raffo,  
1095 E. de la Torre, Multi-grain reconfiguration for advanced adaptivity in  
1096 cyber-physical systems, in: 2018 International Conference on ReCon-  
1097 Figurable Computing and FPGAs (ReConFig), IEEE, 2018, pp. 1–8.
- 1098 [88] J. Sérot, F. Berry, S. Ahmed, Caph: a language for implementing  
1099 stream-processing applications on fpgas, in: Embedded Systems Design  
1100 with FPGAs, Springer, 2013, pp. 201–224.
- 1101 [89] E. Flamand, D. Rossi, F. Conti, I. Loi, A. Pullini, F. Rotenberg,  
1102 L. Benini, Gap-8: A risc-v soc for ai at the edge of the iot, in: 2018  
1103 IEEE 29th International Conference on Application-specific Systems,  
1104 Architectures and Processors (ASAP), IEEE, 2018, pp. 1–4.
- 1105 [90] B. D. De Dinechin, R. Ayrignac, P.-E. Beaucamps, P. Couvert,  
1106 B. Ganne, P. G. de Massas, F. Jacquet, S. Jones, N. M. Chaisemartin,  
1107 F. Riss, et al., A clustered manycore processor architecture for embedded  
1108 and accelerated applications, in: 2013 IEEE High Performance Extreme  
1109 Computing Conference (HPEC), IEEE, 2013, pp. 1–6.
- 1110 [91] E. Lindholm, J. Nickolls, S. Oberman, J. Montrym, Nvidia tesla: A  
1111 unified graphics and computing architecture, IEEE micro 28 (2) (2008)  
1112 39–55.
- 1113 [92] T. Fischer, M. Rogenmoser, T. Benz, F. K. Gürkaynak, L. Benini,  
1114 Floonoc: A 645-gb/s/link 0.15-pj/b/hop open-source noc with wide  
1115 physical links and end-to-end axi4 parallel multistream support, IEEE  
1116 Transactions on Very Large Scale Integration (VLSI) Systems.
- 1117 [93] Y. Lei, Q. Deng, S. Long, S. Liu, S. Oh, An effective design to improve  
1118 the efficiency of dpus on fpga, in: 2020 IEEE 26th International Confer-  
1119 ence on Parallel and Distributed Systems (ICPADS), IEEE, 2020, pp.  
1120 206–213.
- 1121 [94] A. Magnani, G. Brilli, A. Marongiu, Architectural design exploration  
1122 of a lane detection vision pipeline for fpga-based fltenth autonomous  
1123 vehicles, in: *To appear* CF '25 Companion: Proceedings of the 22st  
1124 ACM International Conference on Computing Frontiers: Workshops and  
1125 Special Sessions, ACM, 2025.

- 1126 [95] G. Bellocchi, A. Capotondi, L. Benini, A. Marongiu, Enabling fast  
1127 system-level integration and prototyping of accelerator-rich platforms,  
1128 in: Proceedings of the 22nd ACM International Conference on Comput-  
1129 ing Frontiers: Workshops and Special Sessions, 2025, pp. 54–57.
- 1130 [96] M. O’Kelly, H. Zheng, D. Karthik, R. Mangharam, Fltenth: An open-  
1131 source evaluation environment for continuous control and reinforcement  
1132 learning, Proceedings of Machine Learning Research 123.
- 1133 [97] Vitis Libraries.  
1134 URL [https://www.xilinx.com/products/design-tools/vitis/  
1135 vitis-libraries.html](https://www.xilinx.com/products/design-tools/vitis/vitis-libraries.html)