The FitOptiVis ECSEL Project:
Highly Efficient Distributed Embedded Image/Video Processing in Cyber-Physical Systems

Invited Paper

Zaid Al-Ars[1], Twan Basten[2], Ad de Beer[3], Marc Geilen[2], Dip Goswami[2], Pekka Jääskeläinen[4], Jiří Kadlec[5], Marcos Martinez de Alejandro[6], Francesca Palumbo[7], Geran Peeren[3], Luigi Pomante[8], Frank van der Linden[3], Juka Saarinen[9], Tero Säntti[10], Carlo Sau[12], Maria Katiuscia Zedda[11]

[1]Delft University of Technology, [2]Eindhoven University of Technology, [3]Philips, [4]Tampere University, [5]Institute of Information Theory and Automation, [6]Thales Alenia Space, [7]Università degli Studi di Sassari, [8]Università degli Studi dell'Aquila, [9]Nokia, [10]University of Turku, [11]Abinsula, [12]Università degli Studi di Cagliari

ABSTRACT

Cyber-Physical Systems (CPS) are systems that are in feedback with their environment, possibly with humans in the loop. They are often distributed with sensors and actuators, smart, adaptive and predictive and react in real-time. Image- and video-processing pipelines are a prime source for environmental information improving the possibilities of active, relevant feedback. In such a context, FitOptiVis aims to provide end-to-end multi-objective optimization for imaging and video pipelines of CPS, with emphasis on energy and performance, leveraging on a reference architecture, supported by low-power, high-performance, smart devices, and by methods and tools for combined design-time and run-time multi-objective optimization within system and environment constraints.

CCS CONCEPTS

• Computer systems organization~Distributed architectures • Computer systems organization~Embedded and cyber-physical systems • General and reference~Surveys and overviews

KEYWORDS

Image-video processing, distributed systems, heterogeneous system, energy and performance optimization

1 Introduction

FitOptiVis addresses smart integration of image- and video-processing pipelines for Cyber-Physical Systems (CPS). Smart systems integration is one of the essential capabilities required to maintain and to improve the competitiveness of European industry in the application domains of ECSEL. This is especially relevant for CPS that increasingly are autonomous distributed integrations of electronic systems, subcomponents and software that are tightly interacting with physical systems and their environment. Several of these components are smart systems themselves. Novel design approaches and methodologies are needed to address this new class of computing systems.
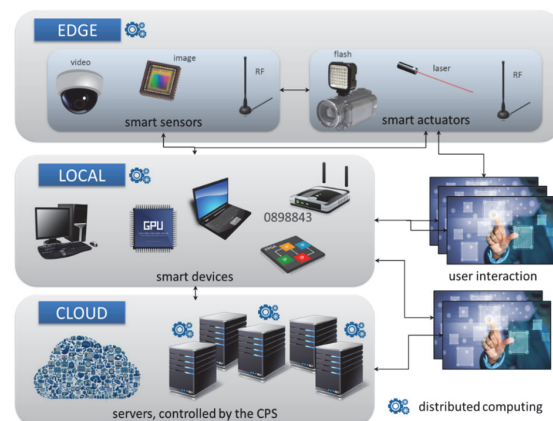

Figure 1: Generic configuration of imaging and video pipelines in CPS

In general, images certainly play a central role in human perception and understanding of our environment. In the same way, CPS need visual context and awareness to make the correct decisions and take the appropriate actions. However, advanced image and video processing is computing intensive, whereas for adequate behaviour the results need to be available with low latencies and high throughput. In many cases, several devices of a distributed CPS need to operate with low energy, as power sources might be scarce, or heat dissipation must be limited to protect the system, its environment or humans interacting with it. Optimization for other quality aspects, such as security, safety, or adaptability, is also important. A generic distributed image- and video-processing pipeline for CPS is shown in Figure 1.

FitOptiVis, which stands for *"From the cloud to the edge - smart IntegraTion and OPtimisation Technologies for highly efficient Image and VIdeo processing Systems"*, has started on June 1st, 2018 and its duration is three years. In the following, we highlight project objectives, explain the selected approach, describe target use cases, and discuss implementation issues.

2 Project Objectives

The main objective of FitOptiVis is to develop an integral approach for smart integration of image- and video-processing pipelines for CPS covering a reference architecture, supported by low-power, high-performance, smart devices, and by methods and tools for combined design-time and run-time multi-objective optimization within system and environment constraints.

CPS are systems that are in feedback with their environment, possibly with humans in the loop. They are often distributed with sensors and actuators, smart, adaptive and predictive and react in real-time [1]. Image- and video-processing pipelines are a prime source for environmental information improving the possibilities of active, relevant feedback. In such a context, advanced imaging and video applications are very complex. Complexity further increases when multiple heterogeneous sensor inputs are combined for analysis and through the distributed nature and the integration of both generic and specialized devices.

On top of that, CPS need to satisfy stringent constraints on real-time behaviour, safety, security, form factor, performance, energy consumption, reliability and quality.

From the implementation point of view, distributed pipelines (as the one depicted in Figure 1) consist of a heterogeneous configuration of legacy devices, state-of-the-art multi-vendor devices and components, and newly developed application-specific ones. Effective smart system integration for imaging and video applications must be built evolutionarily upon earlier developments, and the CPS must be able to cope with individual component (HW and SW) upgrades during its lifetime.

In line with the quote of Thomas Sowell: *"There are no solutions; there are only trade-offs"* [2], achieving the above mentioned main FitOptiVis objective will lead to a generic and unified platform and methodologies that enable the quick creation of image and video processing infrastructure skeletons. These can be used in many domains for different distributed CPS implementations and will allow customized management and optimization techniques for trade-offs, both at run-time and at design time. Skeletons will facilitate evolution. CPS are highly evolvable, they need to adapt to changing contexts and today's needs may rapidly change into something else, where solutions can be borrowed from another application field.

FitOptiVis integral approach on multi-objective optimization is built on a reference architecture. This reference architecture consists of a suitable set of *component abstractions*. It further defines *composable, scalable resource virtualization and distributed run-time support for multi-objective optimization targeting quality and resource management*. The component abstractions capture essential metrics of interest (functionality, performance, energy, resource usage, and so on) of video and imaging tasks and devices and components such as sensors, actuators, processing platforms, storage devices and network components. The architecture is supported by design-time methods and tools targeting performance and energy optimization. In addition, FitOptiVis design-time support involves methods and tools for seamless, compositional integration of the image and video pipelines. Design- and run-time models provide abstractions for performance and energy related to the distributed system configurations and the use of processing, communication and storage resources. The reference architecture will be provided in the form of template solutions for a flexible virtual platform built from the component abstractions and offering multi-objective run-time optimization support for quality and resource management. Imaging and video pipelines can be designed targeting virtual platforms in various configurations corresponding to different points in the multi-objective performance, energy usage, and resource cost space. These virtual platform configurations can be mapped at run-time onto physical resources, depending on their availability and the needs of other applications. Energy-efficient, high performance, smart devices and components will be developed to support, demonstrate, and exploit the reference architecture. The developed tooling ensures effective resource usage predictions and simulations for design-space exploration at design-time, and multi-objective optimizations during run-time.

The increase in power demand by additional functionality is tackled by novel ultra-low power technology. Since the added functionality typically includes both an increase in the number of image generation components and the need for additional processing, the complete imaging or video pipeline must be adapted before a new optimum in power and performance is met. FitOptiVis will integrate ultra-low power and high-performance devices and components into the reference architecture for imaging and video pipelines for CPS to enable performance to be optimized against power by exploiting the advantages of distributed resources in such an imaging pipeline, supported by efficient computing near the sensors or actuators (edge computing) through, for example, configurable hardware accelerators.

FitOptiVis will **"master the complexity, … while reducing the cost of utilizing powerful software intensive products/systems, … enabling the development of dependable and robust, cognitive and collaborative autonomous systems"** [3]. This will lead to shorter development times and improved products with richer functionality.

The FitOptiVis objective will be reached by pushing the state of the art in distributed imaging and video pipelines, development support and run-time support, targeting primarily real-time performance and energy usage. FitOptiVis will provide a reference architecture supporting composability built on suitable component abstractions and embedded sensing, actuation and processing devices adhering to those abstractions. The reference architecture will support design portability, on-line multi-objective quality and resource management and run-time adaptation guaranteeing system constraints and requirements based on platform virtualization. A cloud in the FitOptiVis context is a set of connected servers under the control of the CPS. Public clouds are not considered because they cannot provide the required service guarantees. Non-functional aspects other than performance and energy, such as reliability and security, are taken in to account to meet use case specific objectives but are not an explicit target of research and development. FitOptiVis enables integration of state-of-the-art technology and new developments in the mentioned domains through its reference architecture. The following list summarizes the operational objectives (OBJ) of FitOptiVis:

1. *Reference architecture and virtual platform* - Template solutions
    a. Component abstractions (covering video and imaging tasks and heterogeneous processing, storage and network devices and components);
    b. Virtualization supporting scalability, portability and composability principles;
    c. Multi-objective quality and resource management (support for run-time decision making, adaptation, re-distribution and upgrades).
2. *Design-time support* - Model driven development
    a. Apply performance and energy models, image/video processing models, model combinations;
    b. Multi-objective optimization for performance and energy;
    c. Co-simulation, Design-space exploration, HW/SW co-design targeting FPGAs and custom accelerators;
    d. Integration methods, Heterogeneous devices, evolutionary development.
3. *Run-time support* - On-line multi-objective optimization
    a. Real-time multi-objective combinatorial optimization;
    b. Data and processes distribution;
    c. Run-time adaptation, energy driven adaptations, workload (re-)distribution, support for run-time upgrades.
4. *Energy-efficient, high-performance, smart devices and components*
    a. Sensors, actuators, communication and processing components;
    b. Programmable and Non-Programmable accelerators;
    c. Ultra-low-power devices.
5. *Effective impact*
    a. Ensure exploitation of the project results in distributed image and video processing CPS and their development.

3 Concept and Approach

FitOptiVis developments are driven by industrial use cases that serve as a basis for requirements, demonstration, and validation. Models and abstractions play a crucial role as a basis for an integrated reference architecture, design methods and run-time operation of CPS imaging and video pipelines. Devices and components expose their functional and performance set points, each with the corresponding resource requirements, through minimal interfaces. Model-driven design methods aim to design set points that trade-off functional and performance capabilities against resource usage, for maximal design-time and run-time flexibility [7][8]. Resource virtualization and predictable, composable reconfiguration enable modular, scalable run-time multi-objective adaptation and optimization ensuring quality and real-time performance, even across reconfigurations. On-line monitoring and learning techniques enable the evolution of set points and their resource requirements in changing contexts, both operational environments and use patterns. A common reference architecture will capture the essential aspects of the envisioned approach in appropriate component abstractions, virtualization techniques and quality- and resource management protocols. The approach is grounded in the development of smart, high-performance, energy-efficient devices and components to validate the concepts and as a basis to exploit the FitOptiVis results. Figure *2* illustrates the organizational approach of FitOptiVis while Figure 3 sketches the technical approach. The reference architecture provides the framework for the other activities.

The FitOptiVis operational objectives are motivated by industrial *use cases* from diverse application domains. The use cases share demanding image and video processing, while they differ in other aspects like available processing power, energy budgets, safety requirements, configuration distribution, configuration dynamics, etc. The use cases dictate the *requirements* for the development of an integral approach for smart integration of image- and video-processing pipelines for CPS. The use cases moreover form the basis for *demonstrators*. These demonstrators serve for validation and analysis of the project results, and as a basis for *exploitation*. The use cases will be detailed in the next subsection. Following an industry driven approach guarantees that all the technical specifications for the FitOptiVis platform, components and tools are going to be derived from the use case needs and requirements, leveraging on appropriate requirements gathering and management techniques. Both cross-domain and domain-specific requirements will be considered and prioritized to cover all the necessary design set points. The use cases and demonstrators are chosen to have diversity in image and video processing tasks. The different application domains will benefit from a common

reference architecture, shared design methods, and generally usable devices and components. As already said, CPS are highly evolvable systems. Therefore, the solutions for the problems of tomorrow can then be borrowed from other application fields. The generic and unified platform and design methods delivered by FitOptiVis enable the quick creation of image and video processing infrastructure in different domains, and simplified connections between systems across domains. This will become increasingly important in our rapidly evolving society where systems are without exception connected in the Internet of Things. FitOptiVis will deliver generic results that address this development and that are applicable in many domains (even outside of FitOptiVis) and across domains.
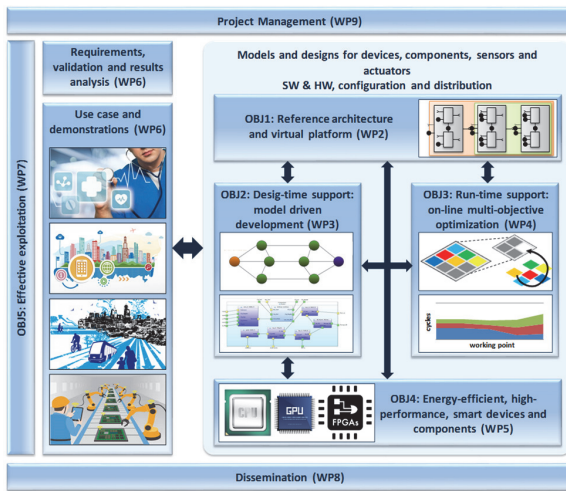

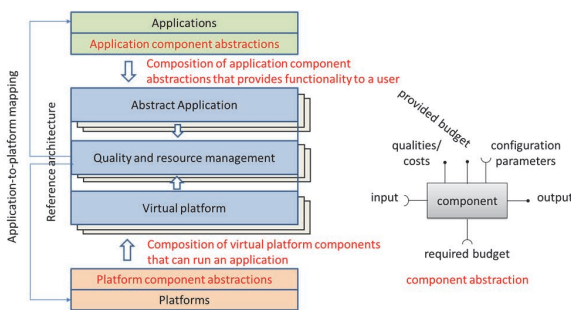**Figure 2: Organizational approach and dependencies**

Figure 3: Technical approach and interfaces – the blue parts form the reference architecture

4 Target Use Cases

The proposed integrated FitOptiVis approach will target end-to-end energy and performance optimization of distributed and adaptive image/video processing pipelines, from the edge to the cloud. FitOptiVis will take its main requirements, architecture,

designs, implementations, demonstration and validation from the use cases (UC) listed below:

1. *Water Supply*
   a. The water supply system constitutes a critical distributed infrastructure that requires:
      i. continuous monitoring, to detect the facility's integrity (i.e. damages and burst pipes) and unauthorized accesses;
      ii. fast surveillance and maintenance intervention, to restore missing service states.
2. *Virtual Reality*
   a. Virtual reality systems are everyday more important in medical, gaming and military applications and cinema industry. Recent advances in digital photography and video led to the development of advanced 3D vision and display systems. Emerging technology of virtual reality applications should be supported by:
      i. high-quality video capture;
      ii. efficient coding and processing technologies;
      iii. an accurate, fast positioning system.
3. *Habit Tracking*
   a. Habit Tracking at home is particularly relevant for the elderly population, to:
      i. assess their physical habits and identify situations where methodological and behavioural concepts can promote healthy lifestyle;
      ii. detect deviations from a standard behavioural pattern, early discovering potential deterioration of users' health or wellness.
4. *Industrial Inspection*
   a. Industrial inspection applies machine vision to quality control in production processes.
5. *Road Traffic Surveillance*
   a. Road and railway traffic surveillance for vehicle detection and recognition.
6. *Multi Source Streaming Composition*
   a. This use case is meant to create an embedded high-performance video compositor capable of rendering 8 simultaneous streams on a screen. It requires:
      i. tight coordination among the video sources, and with the compositor;
      ii. being adaptable to variations in setting. If the screen layout changes, the compositor needs to dynamically adapt.
7. *Sustainable Safe MRI*
   a. To facilitate MRI use in long interventional procedures, in paediatric imaging, and in the emergency setting, we intend to develop:
      i. a low energy deposition method for 4D imaging;
      ii. energy consumption and thermal losses models to predict control and minimize thermal load, while maintaining superior image quality.
8. *Robots Calibration*
   a. Robots, robotic arms or manipulators require calibration to be done periodically or after each geometry change. This use case is meant to provide fast and automated calibration, getting rid of any mechanically connected measurement device.
9. *Surveillance of Smart-Grid Critical Infrastructure*

a. The aim of this use case is active surveillance for the prevention of potential harm to EU Smart-Grid Critical Infrastructures (CI). An adequate level of protection must be ensured, and the detrimental effects caused by disruptions on the society and citizens must be minimized. Given the potential harms to the citizens' well-being, the economy, or the environment, their protection requires:

    i. a hybrid network that combines the information from video surveillance (best effort traffic) with critical control information (which requires deterministic time and low latency) on a single network without impact on each other's service;

    ii. the proposed strategy is to integrate both service types using TSN. TSN can combine on a single network both information flows.

10. *Autonomous Exploration*

a. Definition of the next generation video processors that will constitute the core of the next earth observation and robotic planetary exploration missions. The challenge is to make the processing adaptable to:

    i. different non-functional critical parameters (i.e. available power or connection bandwidth);

    ii. unexpected functionalities under degraded conditions;

    iii. failures of some of system elements due to the challenging environment.
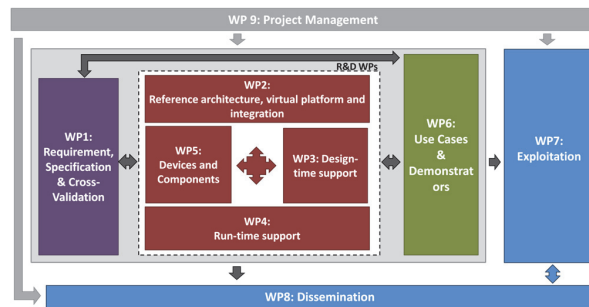
5 Implementation

Figure 4: FitOptiVis Work Plan

5.1 Work Plan

The mission of the FitOptiVis implementation work plan (Figure 4) is: (i) to achieve the technical objectives of the proposal (OBJ 1-4), (ii) to maximize the collaboration among partners and (iii) to guarantee a solid impact of the project (OBJ5). We aim to have impact in terms of advancements with respect to the state of the art, as well as improvements with respect to current industrial best practices and, last but certainly not least, a concrete market uptake of the project outcomes.

FitOptiVis activities are organized as follow:

1. *Requirements, specification and cross-validation of the results* - We apply two sources for requirements: the use cases and the industrial/academic knowledge on the state-of-the-art and current issues. Validation will take results from demonstrators on the use cases and the evaluation against the requirements, its outcome is a report with the KPI metrics.

2. *Reference architecture, virtual platform and integration* – This work package addresses objective 1. The reference architecture will provide component abstractions for the image and video pipelines and their implementation platforms, as well as platform virtualization techniques. Together they form the basis for quality- and resource management for image and video pipelines in terms of qualities like frame rates and power dissipation and required processing, communication and storage resource budgets. Emphasis will be on technology supporting multi-objective optimization for, at least, performance and energy. System-level concerns like distribution and (re-)configuration will be addressed. The architecture needs to make sure that the image and video pipelines and the run-time support can work on a heterogeneous network of hardware devices. Each use case will be built on top of the reference architecture.

3. *Design-time support* – This work package addresses objective 2. It defines a model-based working methodology involving methods and tools for predicting, simulating and estimating at design-time resource usage. In addition, methods, software libraries, reference designs, HW/SW co-design [5] and compilation techniques will improve the resource behaviour of the final system, considering the heterogeneous and changing structure and resource needs of the final system. The development concentrates on the development of all video/image processing building blocks and the run-time support, including hardware IP/accelerators, software applications and sensors.

4. *Run-time support* – This work package addresses objective 3. It gathers all the technology that will implement real time resource management within the system. It will deliver components in a similar way as WP3 both the actual implementation that will run on the final product and the models to integrate in the system model. It involves monitoring, measuring components [6], and control components implementing the algorithms.

5. *Devices and components* – This work packages addresses objective 4. It develops and selects hardware and software devices, components and configurations that are best suitable for optimal energy and performance use. Each component will have a FitOptiVis compliant model view, which may present different levels of abstraction, depending on its usage and the configurations it may embed, such that the run-time support can manage the resources.

6. *Use cases and demonstrators* – This work package, together with WP7 and WP8, addresses objective 5. The use-case teams will develop demonstrators, adopting the FitOpviVis reference architecture, along with all the design- and run-time strategies coupled to it. The demonstrators are validated according to the originally defined use-case requirements, considering both design- and/or run-time resource behaviour.

7. *Exploitation* – This work package, together with WP6 and WP8, addresses objective 5. Nokia ensures that project developments and results will be proactively exploited within the consortium in several domains of the development of intelligent CPS. Since exploitation activities are of paramount importance to reach the overall project impact the

Exploitation manager will actively participate to this WP to manage the implementation of the exploitation tasks in agreement with the overall FitOptiVis exploitation plan.

8. *Dissemination* – This work package contributes to address objective 5 ensuring that project developments and results will be proactively communicated and disseminated within the consortium itself and to relevant stakeholders, the media and the public. These activities are meant to foster the concrete and wide impact of the project findings.

9. *Project management* – Coordination and management of the project, work packages and the consortium, including quality and risk management.
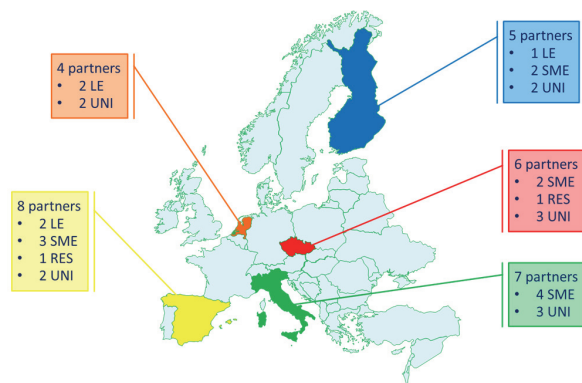

Figure 5: FitOptiVis Consortium

5.2 Consortium as a Whole

Philips is the project coordinator and has the overall responsibility for the FitOptiVis consortium. The FitOptiVis consortium is composed of 30 partners from 5 countries (Figure 5): 5 Large Enterprises, 11 Small and medium enterprises, 2 Research institutes, 12 Universities. All the FitOptiVis partners are highly active at both European and international levels and exhibit a core set of expertise in contributing towards medium- to large-scale European and national projects.

Several large, small and medium size enterprises producing CPS where image processing is a crucial element of correct execution. In all cases the CPS has a distributed set of image/video sensors whose inputs need to be processed fast, and in many cases with low energy use. Some partners have other dependability requirements that need to be regarded as well. Other SME partners, the research institutes and the universities provide appropriate innovative technologies and methodologies to tackle the multi-objective optimization for image and video processing. The different companies will provide demonstrations involving their own CPS, instantiating the FitOptiVis architecture, run-time support and/or components and applying the design time technologies. The solid expertise and the cooperation of all the partners together will allow to successfully carry out the project, challenging our ambitious objectives and demonstrating the project results by means of 10 different demonstrators.

5 Where We Are

In the following table a summary of technical activities carried on in FitOptiVis in the first 9 months.

WP2: The reference architecture will provide template solutions for component abstractions, virtualization techniques and quality- and resource management. We are developing a domain-specific language (DSL) to specify component abstractions (qualities, costs, inputs, outputs, configuration parameters, provided and/or required budgets of components), their compositions and quality- and resource management requirements. The DSL can be specialized for specific use cases.

WP3: Partners are working on different tools for simulation, exploration, optimization and synthesis of image processing pipelines [9]. As an initial release, a design flow based on the Xilinx SDSoC 2018.2 for four development systems (i.e., ZynqBerry, Zynq4x5, UltraScale4x5, UltraScale8x5), capable to accelerate video processing SW functions in the programmable HW logic on the same device, has been released [4].

WP4: In WP4, the focus has been on identifying connections to WP2 concepts and designing a distributed heterogeneous SW runtime stack. A first prototype of such a stack was released internally with a proof-of-concept application distributing execution seamlessly to cloud-edge from a low-end device.

WP5: According to the scenario needs, we have currently defined 7 soft IPs to accomplish specific algorithms/functions, 8 HW components (including processors templates and application specific accelerators), and 3 communication blocks. The release of the first library of components is expected by May 2019.

ACKNOWLEDGMENTS

REFERENCES

[1] http://cyberphysicalsystems.org/ (2019)
[2] Thomas Sowell. 1986. A Conflict of Visions: Ideological Origins of Political Struggles. Basic Books 2002, ISBN 0465081428.
[3] ECSEL GB 2017.78 - MASP 2017, Section 8.1, page 83.Conference Name:ACM Woodstock conferenceConference Short Name:WOODSTOCK'18Conference Location:El Paso, Texas USA
[4] https://www.trenz-electronic.de/ (2019)
[5] Muttillo, V., Valente, G., Ciambrone, D., Stoico, V., & Pomante, L. (2018). Hepsycode-RT: A real-time extension for an ESL HW/SW co-design methodology. ACM International Conference Proceeding Series.
[6] Moro, A., Federici, F., Valente, G., Pomante, L., Faccio, M., & Muttillo, V. (2015). Hardware performance sniffers for embedded systems profiling. 12th International Workshop on Intelligent Solutions in Embedded Systems. ISBN:978-1-4503-0000-0/18/06.
[7] H. Alizadeh Ara, M. Geilen, A. Behrouzian, T. Basten and D. Goswami, "Compositional Dataflow Modelling for Cyclo-Static Applications," 2018 21st Euromicro Conference on Digital System Design (DSD), Prague, 2018, pp. 121-129.
[8] Carlo Sau, Paolo Meloni, Luigi Raffo, Francesca Palumbo, Endri Bezati, Simone Casale-Brunet, and Marco Mattavelli. 2016. Automated Design Flow for Multi-Functional Dataflow-Based Platforms. J. Signal Process. Syst. 85, 1 (October 2016), 143-165
[9] C. Sau, L. Fanni, P. Meloni, L. Raffo and F. Palumbo, "Reconfigurable coprocessors synthesis in the MPEG-RVC domain," 2015 International Conference on ReConFigurable Computing and FPGAs (ReConFig), Mexico City, 2015, pp. 1-8.