ESR - EXECUTIVE SUMMARY REPORT

HIPNOS

Prepared by: David Gonzalez-Arjona (GMV)

Approved by: David Gonzalez-Arjona

Authorized by: David Gonzalez-Arjona

Digitally signed by David González Arjona
DN: dc=es, dc=gmv, ou=Usuarios, ou=GMV-ASP, cn=David González Arjona, email=dgarjona@gmv.com

Code: GMV-HIPNOS-D10

Version: 1.1

Date: 06/11/2017

Internal code: GMV 23875/17 V2/17

DOCUMENT STATUS SHEET

Version	Date	Pages	Changes
1.0	10/10/2017	14	Creation of document
1.1	06/11/2017	14	Removed GMV copyright and added reference to ESA HIPNOS contract

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

1. INTRODUCTION

1.1. PURPOSE

The objective of this document is to provide a summary report of the work developed in the ESA project for the study and definition for the **high-performance avionics architecture** to deal with **current and future complex GNC** systems targeting in this case the **Active Debris Removal** scenario.

1.2. SCOPE

The objective of this document is the presentation of the final synthesis of the activity **"HIPNOS", Development of a representative HW/SW solution for a high-performance processing platform for Active Debris Removal missions. Implement COTS-based solution as Demonstration of the activity**.

HIPNOS includes trade-off study of processing technologies, processing space-grade and COTS devices, processing architectures, trade-off for computer vision algorithms, and demonstrator of the selected solution.

The scope and objectives of the activity are:

- study/define the high-level architecture of a high-performance computing system for space avionics for GNC in ADR missions
- design one high-accuracy & high-complexity chain of Computer Vision algorithms to support the ADR scenario of e.Deorbit
- select the most appropriate acceleration platform in terms of speed, power, rad-hardness, mass/size, flexibility, future trends
- accelerate the CV algorithms on FPGA or GPU, or DSP, or multi-core CPU to achieve the high-speed processing required in ADR
- develop and demonstrate a preliminary, proof-of-concept system (by using COTS components and high-definition videos) with a representative ADR use-case.
- present the feasibility of implementing demanding algorithms with real-time performance on future space-grade platforms

This work has been performed under the **ESA contract No. 4000117700/16/NL/LF**. The consortium participating in this project was formed by GMV Aerospace and Defense as prime, and the Greek universities of NTUA and FORTH as subcontractors.

1.3. DEFINITIONS AND ACRONYMS

1.3.1. DEFINITIONS

Concepts and terms used in this document and needing a definition are included in the following table:

Table 1-1 Definitions

Concept / Term	Definition

1.3.2. ACRONYMS

Acronyms used in this document and needing a definition are included in the following table:

Table 1-2 Acronyms

Acronym	Definition
ADCSS	Avionics, Data, Control, and Software Systems
ADR	Active Debris Removal
AFE	Autonomous Features Extraction
AIT	Action Item List

Acronym	Definition
AOCS	Attitude and Orbit Control System
CAM	Collision Avoidance Manoeuver
CFI	Customer Furnished Item
CIL	Critical Item List
CDR	Critical Design Review
COTS	Commercial-off-the-shelf
CV	Computer Vision
DoF	Degrees of Freedom
ECSS	European Cooperation for Space Standardization
ESA	European Space Agency
ESTEC	European Space Research and Technology Centre
FDIR	Failure Detection, Isolation, and Recovery
FP	Final Presentation
FPGA	Field Programmable Gate Array
GNC	Guidance, Navigation, and Control
HIPNOS	High Performance Avionics Solution For Advanced and Complex GNC Systems
IMA	Integrated Modular Avionics
IRP	Inventory Routing Problem
LEO	Low-Earth Orbit
MCU	Micro Controller Unit
OBC	On-Board Computer
OBSW	On-Board Software
OBDH	On-Board Data Handling
RCS	Reaction Control System
VBN	Vision Based Navigation
VHDL	VHSIC Hardware Description Language
VHSIC	Very High Speed Integrated Circuit
VLIW	Very Long Instruction Word

2. REFERENCES

2.1. APPLICABLE DOCUMENTS

The following documents, of the exact issue shown, form part of this document to the extent specified herein. Applicable documents are those referenced in the Contract or approved by the Approval Authority. They are referenced in this document in the form [AD.X]:

Table 2-1 Applicable Documents

Ref.	Title	Code	Version	Date
[AD.1]	e.Deorbit CDF study report	CDF-135	C	Sep 2012
[AD.2]	SAVOIR-GS-001: SAVOIR generic OBC specification	TECSW/12-536/JLT	1.0	15/01/2015
[AD.3]	e.deorbit Final Report	EDEORBITASD-RP-0009		11 Sep 2014
[AD.4]	e.deorbit Final Report	e.deorbit-TAS-TNFR-0005		15 Sep 2014
[AD.5]	e.deorbit Final Report	EDE-KTRP-0006		15 Oct 2014
[AD.6]	e.deorbit Mission and System Requirements Document (MSRD)	GSTP MSRD e.deorbit		09 Mar 2015
[AD.7]	HIPNOS - High Performance Avionics Solution for Advanced and Complex GNC Systems	ESA AO 1-8410/15/NL/LF		2015

2.2. REFERENCE DOCUMENTS

The following documents, although not part of this document, amplify or clarify its contents. Reference documents are those not applicable and referenced within this document. They are referenced in this document in the form [RD.X]:

Table 2-2 Reference Documents

Ref.	Title	Code	Version	Date
[RD.1]	TN1 – Applications Survey And Analysis Of User Requirements Initial Technology Survey And Trade-Off Product Requirement Specification	GMV-HIPNOS-D1	1.1	29/08/2016
[RD.2]	TN2 – Preliminary Breadboard Design	GMV-HIPNOS-D2	1.0	10/08/2016
[RD.3]	TN3 – Full Breadboard Design	GMV-HIPNOS-D3	1.0	08/02/2017
[RD.4]	TN4 – Breadboard Test Procedure and Test Report	GMV-HIPNOS-D4	1.0	01/07/2017
[RD.5]	TN5 – Target Test Case Definition	GMV-HIPNOS-D4	1.0	01/07/2017
[RD.6]	TN3 update – Full Breadboard Design	GMV-HIPNOS-D6	1.0	10/10/2017

3. HIPNOS PROJECT SCENARIO OVERVIEW

The Clean Space initiative of the European Space Agency (ESA) seeks to decrease the environmental impact of space programmes by focusing, among others, on Active Debris Removal (ADR) and the e.Deorbit mission. In this direction, one of the main challenges is to autonomously track and approach a big non-cooperative satellite such as ENVISAT. To achieve the high level of autonomy and reliability required in this phase of the ADR mission, vision based navigation will guide a chaser spacecraft in real-time based on high-definition images acquired and processed on-board at high frame-rates. The increased complexity of these computer vision algorithms mandates the development and use of high-performance avionics to provide one order of magnitude faster execution than today's conventional space-grade processors.

In the context of ESA's project **HIPNOS** *(HIgh Performance avionics solutioN for advanced and complex GNC Systems*), we study computer vision algorithms and design avionics architectures suitable for high-performance embedded computing in space, such as in the ADR mission scenario.

The rendezvous phase of the e.Deorbit mission will rely on sophisticated imaging sensors and advanced autonomous control to continuously estimate the distance/pose of the target and synchronize to its spin. The chaser will be equipped with sensors such as LIDAR, far range and close range cameras, multispectral camera, and infrared camera. The acquisition rate and the high-definition resolution of these cameras will result in a huge amount of data to be processed on-board, at real-time, which together with the complexity of the related computer vision algorithms, will increase the demand for processing power to unprecedented levels for space applications. Conventional space-grade processors, such as LEON3 and RAD750, can provide only a very limited amount of performance, e.g., in the range of 50-400 DMIPS (with 50-200MHz operating frequency). Put into perspective, for image processing, such performance is 10-100x lower than that of contemporary desktop CPUs, which are used today by the algorithm designers during the early development stages of the Vision-based Navigation blocks (VBN). Even with the latest devices, such as LEON4 and RAD5545 that achieve 10x more speed than their predecessors, the space-grade CPU performance seems to be one order of magnitude less than what will be needed to run highly-accurate vision- algorithms on-board the future spacecrafts. Therefore, building blocks for Advanced Image Processing Systems and complex Guidance Navigation Control (GNC) mandate the design of new generation space avionics to include hardware accelerators, e.g., FPGAs, GPUs, or multi-core VLIW DSP processors.

3.1. HIPNOS TRADE-OFF

In the first part of the activity a great effort was made to analyze and perform a **wide trade-off** of processing technologies towards designing such a high-performance architecture for embedded computing in space. We perform a survey and trade-off analysis involving a wide number of diverse processing units. The set includes older and novel devices, with both radiation-hardened and COTS technology. We note here that COTS devices can be considered for certain missions, e.g., short duration flights mission in closer distances to Earth in non-harsh environment, as in the case of ADR in LEO orbit. Gradually, the analysis focuses on high-performance embedded accelerators, which are also compared to more conventional devices for the sake of perspective. At the same time, we study the application scenario of visual based navigation (VBN) and perform an extended exploration to collect generic specifications and commonly used algorithms, i.e., to define a set of representative benchmarks. However, the profiling of each device depends both on the complexity of the benchmark and the underlying architecture, and hence, the relative performance varies greatly among groups of devices (or even singletons) and creates clouds of overlapping results (of range up to 10x within the same group). This challenge is tackled in our comparative study by combining numerous literature results with consortium in-house development/testing, which ultimately leads to a big consistent picture. Table 3-1 summarizes the devices and algorithms used in our study: we examine 30 representatives from all four processor categories (CPU, FPGA, GPU, and DSP) by utilizing more than 10 well-known benchmarks (with images of 512x384 and 1024x1024 pixels).

Table 3-1 Summary of platforms and benchmarks considered in the comparative study during HIPNOS

		In-house development & testing	Literature survey
platforms	FPGA	Xillinx Virtex6 VLX240T-2, Zynq7000 (Z7020, Z7045)	Altera Stratix III E260, Xilinx Virtex 5QV Virtex4 VLX100, Virtex6 VSX475T, Zynq7000
	CPU	desktop Intel i5-4590, laptop i3-4010U embedded: ARM CortexA9, Intel X1000 space: LEON3, OCE E698PM (LEON4)	desktop: Intel i7-3820, AMD FX-8120 embedded ARM CortexA15 space-grade BAE RAD5545
	GPU	Nvidia GeForce GTX 670, GTX 680, GTX 960	Nvidia GTX 980Ti / 650Ti / 295, Tesla C2050, mobile: Nvidia Tegra K1/X1, ARM Mali-T760
	DSP	space-grade Xentium MPPB embedded multi-core Myriad2	TI multi-core TMS320C6678 and 66AK2H14, 1-core C674x, 64-core RH RC64 (MACSPACE)
benchmarks		2D convolutions (5x5 to 11x11), Harris-corner & Canny-edge detectors, Stereo Matching, Hyperspectral search, Pose Estimation (incl. feature detection, description, matching), Super-Resolution	2D convolutions and SAD (up to 25x25), Harris-corner & Canny-edge detectors, Stereo Matching, Image Denoising and Block Matching, Hyperspectral imaging, *etc.* (plus nominal DMIPS and MFLOPS figures)

Overall, the results show that new generation space-grade CPUs are 10x faster than their predecessors, however, they are still one order of magnitude slower than what will be needed for reliable autonomous VBN. Separated by orders of magnitude, the FPGA accelerators provide the highest performance per Watt than all platforms, whereas the CPUs provide the lowest (any CPU type). In terms of speed, alone, the desktop GPUs and the FPGAs are difficult to distinguish (as groups, they provide similar clouds of results). Likewise, high-end mobile-GPUs and many-core DSPs are difficult to distinguish, although the latter have the potential for slightly better performance and power. In terms of speed, alone, desktop GPUs and FPGAs are clearly better than mobile-GPUs and many-core DSPs. In terms of power, a 10 Watt budget is enough to allow many-core DSPs or mobile-GPUs or FPGAs to accelerate a conventional space-grade CPU by 1-3 orders of magnitude. In a high-performance embedded computing scenario of limited power and mass availability, it would be preferable to utilize COTS 28nm SoC FPGAs, which provide 2-29x faster execution than 28nm DSPs. More specifically, Table 3-2 reports the performance (as normalized throughput, in a linear scale such that the ratio of any two values equals the performance ratio of the compared platforms), power, and performance per Watt, for the most promising devices: one representative FPGA (Xilinx Zynq7045), two multi-core DSP processors (12-core Movidius Myriad2 & 8-core TI 66AK2H14), and for reference, one of latest rad-hard CPUs (OCE E698PM, also SoC) and desktop GPUs (Nvidia GTX 670/680/960, also at 28nm). The table results base on multiple benchmarks with a fair level of optimization on each platform. We note that the throughput of E698PM is hypothetical, i.e., it is extrapolated from our measurements (by assuming very efficient 4-core parallelization and 600 MHz clock) as a best case scenario for the near-future space-grade CPU performance (in practice, it could prove even 10x lower). We also note that the DSP values of Table 3-2 are derived with maximum clock frequency, i.e., 1.2 GHz for 66AK2H14; when/if this rate is decreased during space flight, then the speed gap between DSPs and FPGA will increase proportionally (the relatively low clock rate is an advantage of FPGAs).

Table 3-2 Comparison of Performance (normalized throughput) & Power for selected SoC at 28nm technology (excl. rad-hard LEON4)

	4-core LEON4 (E698PM) 600 MHz	12-core VLIW (Myriad2) 600 MHz	8-core DSP (66AK2H14) 1200 MHz	FPGA (Zynq7045) 200–300 MHz	desktop GPU (GTX 670/680/960) $\geq$ 1 GHz
Performance	< 20	50	70–240	430–1460	600–1800
Power (W)	$\simeq 2.8$	1	$\simeq 10$	4–6	> 70
Perf/Watt	< 7	50	$7 - 24$	110–240	< 25

Beyond performance and power, we consider parameters such as the connectivity of the COTS board, radiation tolerance, size/mass, re-programmability, development effort, and cost. The majority of available COTS boards provide multiple interfaces (Eth, USB, PCIe, etc.), however, the FPGA boards offer the advantage of extra GPIO and FMC pins allowing us to communicate with custom connectors and daughter-boards. A second advantage of FPGAs lies in the various mitigation techniques that can be applied for error correction due to radiation. Considering programmability, the dynamic

reconfiguration of SRAM FPGAs renders them almost as useful as the remaining many-core platforms, even for remote updating. Considering size/mass, the FPGAs are among the most competitive COTS boards (e.g., 5:72 _ 10:16 cm2 for the Zynq MMP Zedboard). On the downside, as expected, the development effort is increased for FPGAs (e.g., 4x compared to SW platforms, even though efficient programming of many-core chips, with multiple levels of parallelization, also requires an increased amount of effort compared to the conventional SW coding).

Following the aforementioned analysis, in HIPNOS, we selected a SoC-FPGA (Zynq XC7Z100, on MMP board) to co-design and accelerate a proof-of-concept algorithm for the pose estimation of ENVISAT. The computer vision algorithm bases on edge detection, rendering of a-priori known 3D models of the satellite, and linear algebra techniques to extract a 6D pose vector. The goal is to achieve processing rates of 5-10 FPS for 1024x1024 input images (monocular camera, 8-bit values). For testing, we use synthetic datasets depicting the ENVISAT tumbling and approaching the camera with various poses, at a distance of 70-30m away from the chaser. The SW algorithm is currently under development in C/C++. A preliminary HW architecture is defined, together with HW/SW partitioning and scheduling, so that the most computationally demanding kernels of the algorithm will be developed in VHDL. The HW acceleration of the image processing kernels provides in general a 20-50x acceleration vs the embedded ARM Cortex-A9 of Zynq. Therefore, the proposed HW/SW co-processing solution will provide a game-changer faster execution than conventional approaches and demonstrates the benefits of employing such an avionics solution in future space missions.

HIPNOS Vision-Based Navigation is based on pose estimation relying on a passive monocular camera using intensity edges as the primary features. With these considerations in mind, a RAPiD-like model-based approach is selected as the basis of HIPNOS. The key idea behind RAPiD is to consider a set of 3D object points, called control points, which are likely to project on high-contrast image edges. By measuring the 2D displacement of these control points on the image plane, the 3D motion of the object between two consecutive frames can be recovered. In the original formulation of RAPiD, the control points were manually sampled offline along the 3D model edges and in areas of rapid albedo change. In our case, they will be generated dynamically by combining information from the detected image edges and a rendered depth map. RAPiD required tedious pre-processing of wireframe object models to determine the visibility of control points from various camera positions (this was achieved by partitioning a view-sphere around the model). In the proposed approach, control points are determined dynamically after rendering, which takes care of visibility constraints automatically. Such an approach provides increased flexibility and better constrained pose refinements at the cost of additional computations. As an indication, we note that around 1000 to 1500 control points are generated with this approach, a number which constitutes an increase by two orders of magnitude compared to the original RAPiD. First, edges are detected with Canny's algorithm. Canny is primarily chosen for performing edge thinning, i.e. removing the unwanted spurious edge points via non-maximum suppression. The 3D model of the object is rendered with the current pose to produce a depth map, which is then combined with the detected edges to determine the control points. Following this, edges are matched between successive images along their perpendicular direction and the matches determined drive the estimation process for the pose refinement.

The following images depict the experimental setup in and the output screen during execution. The MMP is connected via Ethernet to a local PC (with a local-only connection, point-to-point), which is used to remotely execute HIPNOS and collect the data (via ssh). The left monitor in the picture shows the gif file of the recorded demo (independently of the current MMP status), whereas the right monitor shows the remote console (with printed pose and execution time, per frame) and some debug information (edgels superimposed on the ENVISAT image). The MMP board is actually in operation (green LEDs are for power, red LEDs encode partially the state of the algorithm) and the right monitor displays information at run-time.

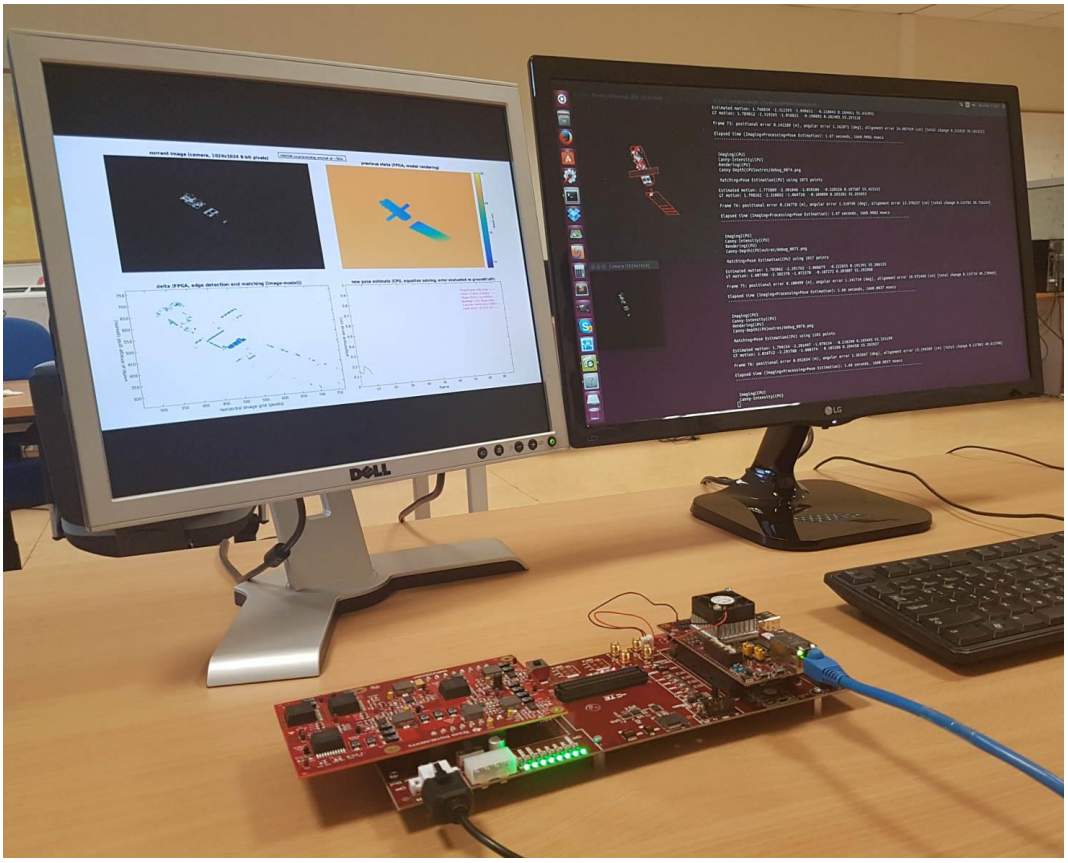

Figure 3-1: Experimental setup (MMP connected to local PC, showing run-time execution)

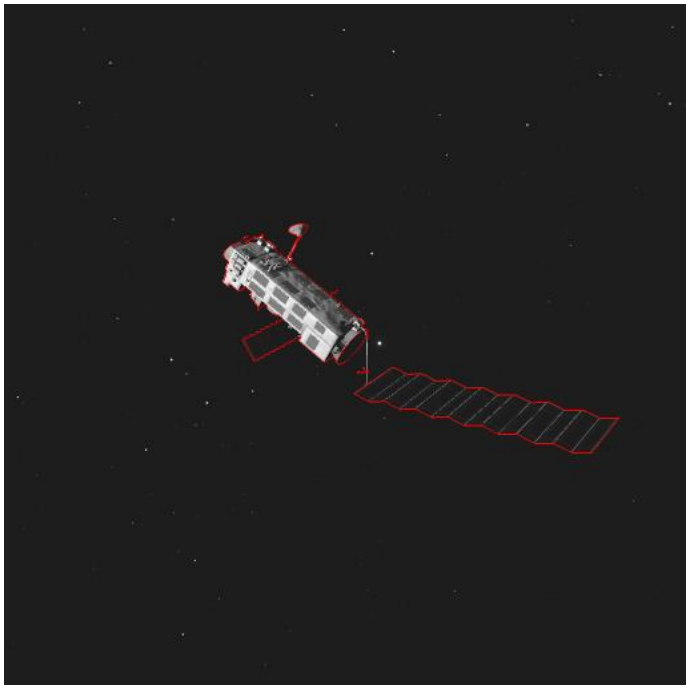

Figure 3-2. Estimated pose of ENVISAT. Depth edges are shown in red, superimposed on the intensity image. Their alignment with the image indicates the correctness of the estimated pose.

4. CONCLUSIONS AND LESSONS LEARNT

The HIPNOS project completed successfully two distinct phases of study regarding Computer Vision for Active Debris Removal: the platform evaluation and the proof-of-concept development of a complete algorithm on Zynq.

First stage of the project consisted on an extensive trade-off of avionics processing elements, devices and architectures, based on previous components, new components and future generation both for space in the form of space-grade rad-hard or rad-tolerant devices and COTS components that could be used in short missions in LEO orbits as an example of a space environment no so harsh. The trade-off also involved the analysis for the pose estimation pipeline algorithm to be implemented in HIPNOS as the Vision-Based Navigation solution to be demonstrated. RAPiD pipeline was selected as the basis algorithms that with some improvements has been developed in HIPNOS following a HW/SW co-design methodology. The pipeline is based on image preprocessing function followed by Canny edge detection, depth rendering, edge matching, pose refinement and Kalman filtering.

The platform evaluation study resulted in a big consistent picture regarding the capabilities of multiple diverse platforms. It showed that new generation space-grade CPUs are 10x faster than their predecessors, however, they are still one order of magnitude slower than what will be needed for reliable autonomous VBN. Therefore, we must design high-performance avionics architectures utilizing HW accelerators. In particular, instead of utilizing multiple chips, it is preferable to utilize SoC devices, which integrate general purpose processors and HW accelerators towards size/mass minimization, power reduction, fast intra-communication, re-programmability, and increased connectivity. Separated by orders of magnitude, the FPGA accelerators provide the highest performance per Watt than all platforms, whereas the CPUs provide the lowest (any CPU type). In terms of speed alone, high-end desktop GPUs and FPGAs are difficult to distinguish (as groups, they provide similar clouds of results). Likewise, high-end mobile-GPUs and many-core DSPs are difficult to distinguish, although the latter have the potential for slightly better performance and power. In terms of speed alone, desktop GPUs and FPGAs are clearly better than mobile-GPUs and many-core DSPs. In terms of power, desktop CPUs and GPUs seem prohibitive for space avionics, however, a 10 Watt budget is enough to allow many-core DSPs or mobile-GPUs or FPGAs to accelerate a conventional space-grade CPU by 1– 3 orders of magnitude. In such a high-performance embedded computing scenario, with relaxed constraints on radiation tolerance due to mission specifications, it would be preferable to utilize COTS 28nm SoC FPGAs, which provide 2– 29x faster execution than 28nm DSP processors.

The algorithmic study showed that

- 3D tracking of satellites can employ image edges as features.
- The use of rendering allows any 3D object model to be employed, avoiding any pre-processing or assumptions about the model's nature.
- Outliers can be handled without the definition of arbitrary outlier thresholds.
- The accuracy of the developed tracker has been shown to exceed that of established tracking techniques.

The development phase of HIPNOS showed in practice that

- It is possible to gain one order of magnitude in performance and performance/watt by VHDL acceleration on Zynq (with kernels providing speedup 60x over ARM). Specifically for CV in VBN, we achieved
  - 10+ FPS for 1024x1024 images (5+ FPS for 2048x2048 images)
  - pose estimation error around 1% (most often less than that)
  - power consumption around 5 Watt (peak <10W)

- Rendering is a very expensive process with respect to logic resources (e.g., 10x more than feature detection, to achieve a similar speedup). However, it is possible to use even a rough model of the tracked object and decrease the complexity.
- Xillybus (or any similar PS-PL communication) is more efficiently used in burst mode and on-the-fly processing (instead of transferring data in blocks and applying local buffering). Pipelining and good scheduling significantly masks the communication overhead
- Customizing the algorithm and the architecture to the given problem proved, once again, very effective (e.g., double speed at component level, and/or decreased resources). Moreover, opting for irregular processing of features might be preferable than regular processing of big datasets (e.g., images) even on the most regular-oriented computational platforms/architectures (like FPGAs)
- Arguably, through significant algorithmic and SW optimization, the latest space-grade CPUs (like LEON4 in OCE E698PM) could achieve a processing rate in the area of 1 FPS. However, if

more than 1Mpx or 5-10 FPS is required, then acceleration on FPGA is mandatory.

Compared to the system requirements identified at the early stages of HIPNOS (in RR reported in DN1), today, we can show the fulfillment of the main contributors to the high-performance avionics that the project was looking for. Derived requirements for the processing board of HIPNOS based on e.Deorbit MSRD and relevant projects was identified and are compared now with the obtained results:

- Electrical Power
  - requirement = 10 Watt (TBC)
  - results = average 4.3 Watt, peak 9 Watt
- Mass/size
  - requirement = 0.5 Kg (TBC) and 20x20x10 cm^3 (TBC)
  - results = 0.065 Kg and 5.7x10.2cm2 (excluding base-board and power supply)
- Processing Power
  - requirement = 10 fps relative navigation over 1024x1024 pixel images
  - results = 12±2 fps relative navigation over 1024x1024x pixel images
- Processing Accuracy
  - requirements = errors <3%
  - results = errors around 1% most often less

4.1. ROADMAP

Final trade-off study will never end because the technology is alive, continuously being improved and new providing new developments. The trade-off should be update with the latest up to date devices and extrapolation of future implementations.

In order to consolidate the proposed avionics solution road to a more mature system road to be flown, several aspects should be taken into account to increase the TRL of the system:

- Real-Time Operating System (already GMV provides RTEMS on the Zynq, but not in the final demonstration)
- FDIR, EDAC, Scrubbing memory (Zynq ARM processors)
- Redundancy Mitigation: TMR, Dual-coreLock-Step (margins)
- ¿HW Shielding, SOI process, lead, current limiters, power cycle?
- Fault-tolerant additional SW design
- In-flight Reconfiguration and Supervisor
- Validation and Verification:
  - (MIL → SIL → PIL) → HIL Tests in Representative Environments
- Fault-Injection Tests
- Radiation Characterization
- Radiation Tests Campaigns
- Road to Flight Model, HW including camera detector Embedded/integrated Solution

END OF DOCUMENT