Project HIPNOS: Case Study of High Performance Avionics for Active Debris Removal in Space

George Lentaris, Ioannis Stratakos, Ioannis Stamoulias, Konstantinos Maragos, Dimitrios Soudris
School of Electrical and Computer Engineering
National Technical University of Athens
Athens, Greece
Email: {glentaris,dsoudris}@microlab.ntua.gr

Manolis Lourakis, Xenophon Zabulis
Institute of Computer Science
Foundation for Research and Technology - Hellas
Heraklion, Greece
Email: lourakis@ics.forth.gr

David Gonzalez-Arjona
GMV Aerospace & Defence SAU
Tres Cantos, Madrid, Spain
Email: dgarjona@gmv.com

***Abstract*—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 eDeorbit. 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 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 algorithms and avionics architectures suitable for ADR. The examined algorithms base on image feature extraction and the architectures base on COTS SoC-FPGA devices. Preliminary analysis highlights the benefits of employing this avionics solution in future space missions.**

I. Introduction

The Clean Space initiative of ESA is a pioneering eco-friendly approach to human space activity [1]. Such an approach is essential for preserving key orbits for future use. The safety of Earth's orbital environment relies on the design of non-debris creating missions, as well as on mitigating the amount of debris already proliferating in space. Overall, Clean Space has three branches: *EcoDesign*, to address environmental impacts and foster green technologies on ground, *CleanSat*, to prevent the production of debris in space in the future, and *eDeorbit*, to remove a large piece of space debris already orbiting the Earth.

Numerous analyses worldwide have confirmed the need for novel technologies and approaches to remove space debris [6]. Decades of launches have left Earth surrounded by a halo of space junk including more than 17,000 trackable objects larger than a coffee cup [2]. At orbital velocity, even a 1 cm nut could hit with the force of a hand grenade. Therefore, this debris threatens on-going and future missions with catastrophic collision. To date, official reports include 5 cases of unintentional high-speed collisions between active satellites and orbital debris, as well as multiple collision avoidance manoeuvres with major cost in satellite fuel. The studies confirm that, even if the industry complies with the existing mitigation measures and stops creating any new space debris, due to the Kessler effect [5], it will not be sufficient to prevent the continuous exponential growth of the debris population already in space.

The problem becomes more important in certain Low Earth Orbits (LEO), which are highly populated by satellites and could be rendered unstable. The most effective way to stabilize the LEO environment is to remove from orbit the large non-functional satellites and launch vehicles. Such Active Debris Removal (ADR) mission scenarios consider that a specialized spacecraft, the *chaser*, will approach the non-cooperative debris, the *target*, to capture and de-orbit it, e.g., perform a controlled burn-up in the atmosphere [6]. More specifically [2], in the first ESA ADR mission "e.Deorbit" of 2020+, the target will be ENVISAT, i.e., a massive drifting satellite left in an uncertain state of possibly rapid tumbling rate (excluding solar panels, approximately, it weights 8 tons and has a size of $2.5 \times 2.5 \times 10\ \mathrm{m}^3$, while orbiting 772Km above Earth, fig. 1). Multiple capture mechanisms with various trade-offs have been examined by e.Deorbit, including throw-nets, robotic arm, and harpoons. In all of these cases, one of the most critical and challenging phases of the mission will be to achieve the rendezvous and steady stationkeeping with the target, autonomously.

The rendezvous phase will rely on sophisticated imaging sensors and advanced autonomous control to continuously estimate the distance/pose of the target and synchronize to its spin [3]. Generally, the chaser's sensors will include LIDAR, far range and close range cameras, multispectral camera, and infrared camera [4]. 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. Clearly,

2159-3477/17 $31.00 © 2017 IEEE
DOI 10.1109/ISVLSI.2017.68

Fig. 1. Visualization of ENVISAT ©ESA, the target debris of e.Deorbit.

such requirements mandate the use of new generation space avionics providing much higher performance than the hitherto utilized space-grade processors.

Advanced Image Processing Systems and complex Guidance Navigation Control (GNC) will be the enabling building blocks in ADR, as well as in similar space proximity operations in the future (inspection, in-orbit servicing, cargo and satellite delivery, space tug, in-orbit assembly, etc.) [4], [7]. However, most often, terms such as “advanced”, “complex”, and “autonomous”, conceal increased processing workload and decreased available time. In contrast, 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 $50-200$MHz operating frequency). Put into perspective, for image processing, such performance is $10-100$x 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). The latest space-grade devices, such as LEON4 and RAD5545, achieve 10x more speed than their predecessors; however, even these devices seem to be one order of magnitude slower than what will be needed to run highly-accurate vision algorithms on-board the future spacecrafts. Therefore, to support autonomous VBN, the community must advance the space avionics architectures to include hardware accelerators, e.g., FPGAs, GPUs, or multi-core VLIW DSP processors.

Designing new avionics architectures involves in-depth trade-off analysis with multiple metrics: performance, power, energy, cost, size, reconfigurability, application-design complexity, fault tolerance, etc. The analysis must consider factors such as limited power availability in space (e.g., due to solar powering and power dissipation capabilities) and limited performance availability (e.g., due to the characteristics of the space qualified electronics). The analysis becomes even more complicated when involving novel accelerators for the reasons explained above.

In general, the selection of avionics must adhere to space-oriented constraints, such as minimal resources (power consumption, mass and volume) and radiation tolerance. However, due to the nature of the ADR mission, the chaser spacecraft will remain in flight for only 3-4 months and at Low Earth Orbit, i.e., it will avoid the extremely harsh radiation environment and the total ionizing doses will be rather limited. Therefore, it becomes feasible to consider non-rad-hard Commercial-of-The-Shelf (COTS) technology for its avionics. Due to higher performance at lower cost, COTS parts compete with radiation-hardened components within the space market and, lately, their use in hi-rel space applications is on the rise. Their use ranges from memory chips in flagship planetary explorers to various parts in pico-nano sats (currently, there is a growing number of operational constellations providing useful services which are based on COTS, selected and tested to meet certain mission requirements).

In the context of ESA project HIPNOS (*HIgh Performance avionics solutioN for advanced and complex GNC Systems*), we study and develop new avionics architectures able to provide considerably increased performance, possibly with the use of COTS accelerators. We take into consideration the requirements of computer vision algorithms for VBN and the space-oriented constraints to propose and demonstrate a HW/SW solution suitable for the future ADR missions. More specifically, regarding the algorithms, we develop image processing pipelines and test them on synthetic datasets depicting various ENVISAT proximity operations. We select the most accurate of the SW pipelines to study their complexity and extract realistic requirements for the underlying HW to meet. Regarding the HW architecture, we perform a trade-off analysis, as mentioned above, to explore the possible solutions with and without rad-hard components. We focus on System-on-Chip devices integrating embedded processors and FPGA resources to accelerate the selected algorithms and evaluate all of their performance metrics.

The remainder of the paper reports the status of our on-going project HIPNOS providing a survey of relevant algorithms (Section 2), overview of relevant avionics (Section 3), and our preliminary analysis/estimations for ADR (Section 4).

II. Previous Work on Computer Vision Algorithms for Autonomous Rendezvous in Space

Autonomous rendezvous will be the most computationally intensive phase of the ADR mission for the on-board computers, either at "inspection" distance around 50m away from ENVISAT, or at 2m "capture" distance (e.g., robotic arm scenario) [3],[6]. Therefore, the rendezvous is considered as the most appropriate application case for studying the performance requirements to design the underlying HW (worst case analysis). The main role of the VBN subsystem will be to continuously estimate the 6DoF relative pose of ENVISAT, i.e. its position and attitude (orientation) in space, relative to the chaser [8]. Pose can be described mathematically by means of a rotation and translation transformation, which brings an object from a reference pose to the observed one. Accurate and high-rate pose estimations will be used to control the motion of the chaser spacecraft during the rendezvous with the uncommunicative/uncooperative orbiting ENVISAT.

Estimating the pose of such a target is quite challenging, because no strong prior information regarding its attitude and position will be available. The situation is aggravated due to the varying illumination conditions, largely varying chaser-target distance that results in dramatic aspect changes, self-occlusions, background clutter, and measurement noise. Such challenges can be tackled by elaborate and computationally demanding processing of images.

Optical approaches to real-time pose estimation for autonomous space rendezvous and docking can be categorized as active or passive, according to the sensing modality they employ. Another taxonomy of existing approaches can be defined by classifying them as either target-based or shape-based [9]. Target-based systems employ a set of easily detectable artificial markers (either retro-reflectors or visual fiducials) placed in a known arrangement on the tracked spacecraft. These markers are tracked using a camera or laser and relative pose is computed from the transformation mapping the known marker pattern onto the currently acquired images. Examples of such systems include the Space Vision System (SVS), the Trajectory Control System (TCS) or the TriDAR (Triangulation and LIDAR Automated Rendezvous and Docking). Such systems have been successfully deployed on the International Space Station or the Space Shuttle [10].

The aforementioned target-based systems achieve highly accurate results, however, they are not applicable in the current mission scenario (ENVISAT has no such customized markers). Instead, we must employ more complex algorithms, such as shape-based approaches that exploit the spacecraft's known geometry. These shape-based approaches utilize 3D models of the target spacecraft and seek the pose that best aligns the spacecraft's model with the captured data/images. In most cases, the required model is known beforehand, however a few methods use Structure from Motion (SfM) techniques to construct the model incrementally as they approach the target. At the core of the shape-based systems, the algorithms process the acquired image to compute a variety of visual features, either keypoints [11],[12],[13],[14],[15],[16], edges [17],[18],[19] or blobs [20], and depth maps [9],[21],[22]. The shape-based algorithms provide increased flexibility to the VBN system, however at the cost of greater sensitivity to error (mismatches between sensed features and the model, which can have a devastating effect on the quality of the estimated pose). Such errors can be reduced by employing more sophisticated algorithms, which in turn demand an increase in the processing power of the underlying HW.

III. Previous Work on High-Performance Avionics Architectures for Space

To date, the majority of space missions utilize rad-tolerant and rad-hard IBM PowerPC and Gaisler LEON processors. Especially for European missions, LEON seems to be preferred in critical functions, e.g., GNC, whereas PowerPC is preferred as payload in less critical situations, e.g., on rovers. In terms of power consumption, the LEON-based solutions are less demanding than PowerPC, e.g., $1-2$ Watts versus $10-20$ Watts, respectively. In terms of radiation tolerance, one can find reliable implementations in both families, e.g., withstanding $100-1000$ Krad(Si) TID and suffering only 10^{-6} SEU per bit-day. The latest devices 64-bit quad-core RAD5545 and SPARC V8 E698PM can deliver up to $2-5$K DMIPS performance, i.e., they provide a tenfold increase compared to their predecessors. However, even so, for advanced GNC with VBN, on-board accelerators are essential.

Excluding ASIC, the most common general-purpose accelerators are Field Programmable Gate Arrays (FPGA), Very Long Instruction Word (VLIW) DSP multi-core processors, and Graphics Processing Units (GPU). The latter are not yet considered for space applications, mainly due to their increased power dissipation, in the order of 100 Watts (embedded GPUs decrease power to 10W but at the cost of 10x lower performance). The VLIW DSP solution is being examined through projects such as MACSPACE (with 64 X1643 DSP cores from CEVA) and SSDP (with Xentium VLIW DSP cores from RECORE), or devices such as Texas Instruments 66AK2H12 (8-core C66x DSP cores) or even Movidius Myriad2 VPU (12-core VLIW). However, the most promising solution, especially in terms of performance per Watt, seems to be the FPGA (in the industry, so far, the only possible high-performance avionics solution has been to use FPGAs, since no space grade DSPs are available).

The FPGA market offers a multitude of devices, either space-grade or COTS. However, given the increased requirements of image processing in terms of logic and memory, the most suitable space-grade FPGAs for VBN seem to be the Xilinx Virtex-5QV and the Microsemi RTG4. The former provides double on-chip memory, which is a critical resource for image processing. The space-grade European market offers today the ATMEL devices (which are considered quite small for high-performance tasks) and, recently, the NanoXplore NX-eFPGA (also dubbed BRAVE, or NG-MEDIUM) which will be able to support high-performance tasks in the near future. Besides space-grade, cutting-edge technology includes System-on-Chip COTS devices, like Xilinx Zynq and Microsemi Smartfusion and Altera Arria, which could be also considered in the context of HIPNOS.

Multiple research groups study plausible/efficient avionics architectures utilizing FPGAs. Most efforts are oriented to the use of System-on-Chip (SoC) computers and, moreover, to exploring the possibility of relying on COTS devices. As a result of combining the aforementioned trends (high-performance, SoC, COTS), a reasonable solution currently being studied for future use in space is the Xilinx Zynq FPGA [23],[24]. The JPL proposes for NASA APEX [23], an FPGA-based SoC, which is currently prototyped on Zynq-7Z100. Next to the Zynq chip, APEX integrates unvolatile SSD and DDR memories (all on ZedBoard mini-ITX, plus an FMC connected card providing ADC functionality). Also, [23] proposes to use certain functionality of the COTS Zynq (e.g., internal watchdogs) and mitigation techniques on the FPGA side to increase the resilience of the system to space radiation. In total, the volume of this avionics module is $17 \times 17 \times 5$ cm^3, whereas its power dissipation is around 5 Watts. The Florida University proposes CHREC [24], which is also based on Xilinx Zynq7000. The developed board includes rad-hard flash memory and has a form factor of 1U, i.e., it fits in a $10 \times 10 \times 11.35$ cm^3 cubesat; it weighs less than 0.1Kg and consumes around 2 Watts of electrical power. The Goddard centre of NASA proposes using SpaceCube [25][26], which is an architecture based on space-grade FPGAs (e.g., Virtex-5QV). They develop a hybrid computing platform (CPU, plus DSP processors, plus FPGA logic) with various versions already tested in space environment. “SpaceCube mini” is the smallest module (form-factor 1U consisting of three U-folded boards, one of which includes the Xilinx Virtex-5QV as the main processing unit), it weighs less than 1.4 Kg (with a housing 2.5mm thick), and consumes around 5 Watts, whereas its estimated TDI is $50{-}700$ Krad(Si). “Spacecube 2.0” is larger (flight unit with 4 boards inserted in a rack of 3U form factor), i.e., $127 \times 178 \times 229$ mm^3 with 5.8Kg, and consumes 20 Watts.

Before concluding this section, we note that in the process of designing new and efficient avionics architectures in Europe, ESA supports the SAVOIR initiative (Space Avionics Open Interface aRchitecture) [27]. SAVOIR seeks to federate the space avionics community to work together in order to improve the avionics subsystems, i.e., to influence standardisation processes by defining the governance model to be used for the products, generic specifications, interface definition of the elements being produced. The primary output of SAVOIR is a reference avionics architecture for spacecraft platform hardware and software.

IV. Preliminary Analysis in HIPNOS

The current section presents some key decisions taken early in the development phase of HIPNOS in order to setup preliminary platforms and experiments, which will be gradually refined towards a high-performance avionics solution. Our preliminary analysis bases on combined literature survey and in-house testing.

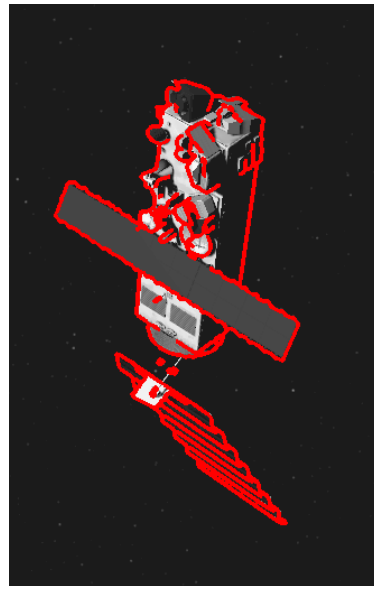

Fig. 2. Example of ENVISAT model edges superimposed on grayscale image (resolution 1024x1024, camera is at a distance of 50m, edge pixels are in red)

Starting with the sensors to be used in our study, we note that cameras have the advantages of compact size, low cost, robust hardware and low power consumption. On the downside, cameras are sensitive to ambient illumination and are associated with more error than LIDARs when estimating the target’s pose at certain distances (e.g., at very close range). In HIPNOS, we will study the scenario of utilizing a high-definition camera located $60{-}30$m away from ENVISAT, with a rate of $5 - 10$ frames per second (monocular, 8- or 10-bit greyscale pixels, resolution 1024x1024). For testing, we will use synthetic datasets depicting the ENVISAT tumbling and approaching the camera with various poses (fig. 2).

Concerning the building blocks of the Computer Vision (CV) algorithms to be developed, our preliminary analysis favors shape-based approaches that utilize edge features and a-priori known models of ENVISAT. Edges are preferred to interest points because the former call for highly textured objects and are more sensitive to ambient illumination conditions. Our preliminary experiments indicated that relatively few point features could be extracted from the test images. In general, such points/corners could be successfully matched between consecutive frames, however, tracking them along longer trajectories or matching them against offline databases had limited success. Instead, geometrical edge-based features (defined by pixels where image intensities undergo a sharp variation) seem to be more robust against illumination conditions and image noise. Also, they are more suitable for matching when the target object is weakly textured, i.e., as is the case with ENVISAT. Fig. 2 shows an example of edges detected using the Canny algorithm: the shape of ENVISAT becomes clear allowing for successful matching to offline models. That is, we can adopt a 3D mesh model of the target spacecraft, which will be graphically rendered at arbitrary, hypothesized poses, to provide a strong prior on the expected appearance of ENVISAT in the captured images (currently, we prefer the a-priori known model approach over Structure-from-Motion techniques, owing to the practical difficulties involved in visually reconstructing a 3D model of adequate quality on-the-fly).

Concerning the high-level architecture design, we have considered the following generic scenarios: 1) many-core CPU, 2) CPU plus multi-core DSP, 3) CPU plus FPGA, 4) multi-FPGA including soft-core CPU, and 5) System-on-Chip. The first four assume distinct devices interconnected with, e.g., a high-speed network like SpaceWire at 100 Mbps. The drawback of these approaches will be the increased size/mass (multiple chips) and power dissipation. Additionally, specifically for multicores with SW, developing a hard real-time system would require functional isolation together with time analysis against single-core chips (which is difficult due to state interruption between processes/applications, constraints of high-criticality over low-criticality applications, safety and worst-case execution time requirements of SW, as is the common practice in the space industry). FPGAs can avoid such complicated SW analysis, although they require increased development time. Considering softcore CPUs (e.g., LEON synthesized in an FPGA), the benchmarks show significantly low processing power (e.g., $20-70$ MIPS), and hence, they would hinder the speed of HW/SW co-processing. Instead, ARM-based ASIC CPUs integrated in a SoC FPGA can provide more than 1K MIPS (even with single-core execution). Another bottleneck for HW/SW co-processing could arise in cases 2-4 due to slow interconnection between devices (limited communication bandwidth between functions of the algorithm). We note that, in all five cases above, one can today choose among space-grade and COTS parts.

Considering all the aforementioned factors, we identify the SoC approach as the most appropriate for HIPNOS. In particular, we select SoC FPGAs, which offer a wide range of advantages owing to the coupling of FPGA logic with processors and peripherals. They can lead to architectures with single-device, low-power, small size/weight, and fast intra-chip communication able to support efficient HW/SW co-processing. The SW part allows for increased flexibility with respect to the algorithm design, whereas the HW part allows for significant acceleration. Furthermore, SoC FPGAs allow for developing/prototyping complicated systems in a reasonable amount of time and cost by utilizing COTS devices. Comparing the available high-performance COTS devices in the market leads to the Microsemi SmartFusion2 and Xilinx Zynq7000 candidates. The latter is preferred in HIPNOS due to the nature of the examined application, i.e., high-definition image processing in not very harsh radiation environments: roughly, Zynq provides more than 2x logic and 6x more on-chip memory, with a more powerful processor (Cortex-A9 versus Cortex-M3), under an affordable increase of power consumption (e.g., double than SmartFusion2, but still around 6 Watts in in-house tests). We also note that, in the near future, a member of the Zynq Ultrascale+ family will become rad-tolerant (with even more resources, e.g. 10x more logic and 7x more memory than space-grade Virtex-5QV, which is already considered in [28] for implementing computer vision for planetary rover applications). As a representative platform for development in HIPNOS, we use the Zynq MMP board (fig. 3), with size 5.72×10.16 cm^2 and mass 65gr. (MMP alone, not together with its baseboard), integrating the XC7Z100 FPGA (277K LUTs, 555K DFFs, 2K DSPs, 755 RAMB36 or 26Mbit, and dual-core ARM Cortex-A9 CPU at 667MHz).

TABLE I
EXAMPLE TIME FOR BASIC CV FUNCTIONS PER 1-MPIX IMAGE (ON ARM CORTEX-A9 AT 667MHZ), COMPARED TO EXAMPLE BUDGETS FOR VBN.

basic function	Time (ms) on Cortex-A9	HIPNOS budget (ms)	required acceleration
Harris detect	2801	100	28x
Canny detect	505	50	10x
contrast enhance	27	50	1x

Using the Zynq board to measure the SW execution time of few representative computer vision functions leads to the preliminary results reported in Table I. Feature extraction algorithms like Harris and Canny (for detecting interest point/corners and edges, respectively) require $0.5-2.8$ sec to process one 1-Mpix image (1024x1024 8-bit values) on ARM Cortex-A9 (1-thread, gcc -O3, 667 MHz). In contrast, given the HIPNOS specification of $5-10$ FPS for the entire pose estimation pipeline, we expect that such functions need to run, at most, within $50-100$ msec. Hence, the FPGA implementation must provide $10-28$x acceleration factors when porting the most demanding parts of the CV pipeline. This one order of magnitude faster execution is within the reach of FPGA accelerators, as recently demonstrated in [28], and hence, we expect that Zynq will be able to support the future VBN of ADR.

V. CONCLUSION AND FUTURE WORK

Advanced and complex GNC systems with accurate and autonomous Vision Based Navigation will require new avionics architectures providing one order of magnitude higher performance than the latest space-grade CPUs. Currently in the HIPNOS project, we design such architectures based on COTS SoC FPGAs and we develop computer vision algorithms suitable for the ADR scenarios. In the future, we will implement the most representative pose estimation algorithms in HW/SW to quantify the benefits of our proposed avionics solution.

ACKNOWLEDGMENT

The authors would like to thank Gianluca Furano from ESTEC, ESA, The Netherlands. The work is supported by the European Space Agency via the HIPNOS project (programme GSP, ref. 4000117700/16/NL/LF, duration 12 months).

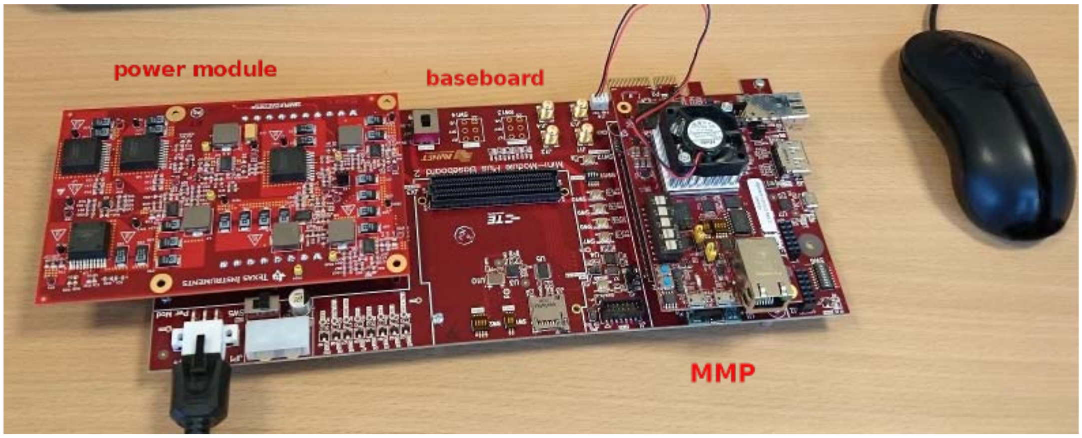

Fig. 3. The MMP development kit with Zynq XC7Z100 SoC FPGA (size 5.72×10.16 cm^2, mass 65gr.), plus baseboard (below) and power module (left).

References

[1] ESA Clean Space initiative, http://www.esa.int/Our_Activities/Space_Engineering_Technology/Clean_Space

[2] ESA document ESA-TEC-SC-TN-2015-007, "e.Deorbit Implementation Plan", Clean Space, European Space Research and Technology Centre, 18 December 2015.

[3] ESA document GSTP MSRD e.deorbit, "e.Deorbit Phase B1 -Mission & System Requirements Document (MSRD)", European Space Research and Technology Centre, 9 March 2015.

[4] ESA document ESA-GSTP-TECT-PL-003997, "GSTP Element 1 Develop - Compendium of Potential Activities: Clean Space", TEC-TI, European Space Research and Technology Centre, 20 January 2017.

[5] D.J. Kessler and B.G. Cour-Palais, "Collision Frequency of Artificial Satellites: The Creation of a Debris Belt", *Journal of Geophysical Research*, vol. 83, no. A6, pp. 2637–2646, 1 June 1978.

[6] C. Bonnal, J.-M. Ruault, and M.-C. Desjean, "Active debris removal: Recent progress and current trends", *Acta Astronautica*, vol. 85, pp. 51–60, 2013.

[7] A. Flores-Abad, O. Ma, K. Pham, and S. Ulrich, "A review of space robotics technologies for on-orbit servicing", *Progress in Aerospace Sciences*, vol. 68, pp. 1–26, Jul. 2014.

[8] E. Kervendal, T. Chabot and K. Kanani, "GNC Challenges and Navigation Solutions for Active Debris Removal Mission", in Advances in Aerospace Guidance, Navigation and Control, pp. 761–779, Springer 2013.

[9] C. English, G. Okouneva, P. Saint-Cyr, A. Choudhuri and T. Luu, "Real-time Dynamic Pose Estimation Systems in Space: Lessons Learned for System Design and Performance Evaluation", in Special issue on Quantifying the Performance of Intelligent Systems, International Journal of Intelligent Control and Systems, 16(2), 79-96, 2011

[10] J.A. Christian, and S. Cryan, "A Survey of LIDAR Technology and its Use in Spacecraft Relative Navigation", AIAA Guidance, Navigation, and Control Conference, Boston, MA, 19–22 August 2013.

[11] L. Vacchetti, V. Lepetit, and P. Fua, "Stable Real-Time 3D Tracking Using Online and Offline Information", IEEE Trans. on Pattern Analysis and Machine Intelligence. 26(10), pp. 13851391, 2004.

[12] F. Qureshi and D. Terzopoulos: Intelligent perception and control for space robotics. Machine Vision and Applications 19(3): pp. 141-161, 2008.

[13] S. Augenstein, "Monocular Pose and Shape Estimation of Moving Targets for Autonomous Rendezvous and Docking", Ph.D. thesis, Dept. of Aeronautics and Astronautics, Stanford University, Palo Alto, CA, June 2011

[14] A. Ansar and Y. Cheng "Vision Technologies for Small Body Proximity Operations", Proceedings of the International Symposium on Artificial Intelligence and Robotics & Automation in Space (i-SAIRAS 10), Sapporo, Japan, 2010

[15] M. Lourakis, X. Zabulis, "Model-Based Pose Estimation for Rigid Objects", International Conference on Computer Vision Systems, vol. 7963 of Lecture Notes on Computer Science, Springer Berlin Heidelberg, pp. 83-92, 2013

[16] M.A. Post, X.T. Yan, J. Li and C. Craig, "Visual Pose Estimation System for Autonomous Rendezvous of Spacecraft", 13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA15), 11-13 May 2015, Noordwijk, 2015

[17] J.M. Kelsey, J. Byrne, M. Cosgrove, S. Seereeram and R.K. Mehra, "Vision-based relative pose estimation for autonomous rendezvous and docking", 2006 IEEE aerospace conference, March 411, Big Sky, MT, USA; 2006. p. 120, 2006

[18] T. Drummond and R. Cipolla, "Real-time Visual Tracking of Complex Structures", IEEE Trans. on Pattern Analysis and Machine Intelligence, vol.24, no.7, pp.932-946, 2002.

[19] A. I. Comport, É. Marchand, M. Pressigout, and F. Chaumette, "Real-time markerless tracking for augmented reality: The virtual visual servoing framework," IEEE Trans. Vis. Comput. Graph., vol. 12, no. 4, pp.615–628, 2006.

[20] D. Cho and P. Tsiotras, G. Zhang and M. Holzinger, "Robust Feature Detection, Acquisition and Tracking for Relative Navigation in Space with a Known Target", Proc. AIAA Guidance, Navigation, and Control Conference, 2013

[21] P. Jasiobedzki, M. Greenspan and G. Roth, "Pose Determination and Tracking for Autonomous Satellite Capture", Proceedings of the sixth International Symposium on Artificial Intelligence and Robotics & Automation in Space (i-SAIRAS 01), Montreal, Canada, 2001

[22] X. Zabulis, M. Lourakis, P. Koutlemanis, "3D Object Pose Refinement in Range Images", International Conference on Computer Vision Systems, vol. 9163 of Lecture Notes on Computer Science, Springer International Publishing, pp. 263-274, 2015

[23] Iturbe, Xabier, Didier Keymeulen, Emre Ozer, Patrick Yiu, Daniel Berisford, Kevin Hand, and Robert Carlson, "An integrated SoC for science data processing in next-generation space flight instruments avionics", in 2015 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC), pp. 134-141. IEEE, 2015.

[24] Rudolph, Dylan, Christopher Wilson, Jacob Stewart, Patrick Gauvin, Alan George, Herman Lam, Gary Alex Crum, Mike Wirthlin, Alex Wilson, and Aaron Stoddard, "CSP: A Multifaceted Hybrid Architecture for Space Computing", 28th annual AIAA/USU conf. on small satellites, 2014.

[25] Flatley, Thomas P. "SpaceCube: A Family of Reconfigurable Hybrid On-Board Science Data Processors", Keynote 2, 2014 International Conference on ReConFigurable Computing and FPGAs (ReConFig14), 2014

[26] Petrick, David, Nat Gill, Munther Hassouneh, Robert Stone, Luke Winternitz, Luke Thomas, Milton Davis, Pietro Sparacino, and Thomas Flatley, "Adapting the SpaceCube v2. 0 data processing system for mission-unique application requirements", In Adaptive Hardware and Systems (AHS), 2015 NASA/ESA Conference on, pp. 1-8. IEEE, 2015

[27] Space AVionics Open Interface aRchitecture, http://savoir.estec.esa.int/

[28] G. Lentaris, I. Stamoulias, D. Soudris, and M. Lourakis, "HW/SW co-design and FPGA acceleration of visual odometry algorithms for rover navigation on Mars", IEEE Transactions on Circuits and Systems for Video Technology, vol. 26, no 8, pp.1563-1577, Aug 2016.