DSP and FPGA – Competition, Synergy, and Future Integration

R. Trautner[a] , J. Both[a], D. Merodio[a], R. Jansen[a], R. Weigand[a]

[a]ESA/ESTEC, 2200 AG Noordwijk, The Netherlands

Roland.Trautner@esa.int

Abstract

Digital Signal Processors (DSPs) have been popular devices for computation-intensive data processing for many decades. More recently, programmable logic devices (PLDs) have seen a dramatic evolution in terms of performances, popularity and capabilities of devices and programming tools.
The application spectrum for programmable devices with respect to General Purpose Processors (GPPs) and DSPs has therefore evolved from a complementary, supportive role to a more competitive/ synergetic one.
The evolution of chip technology follows a long-term trend towards increasingly complex Systems on Chip (SoC), integrating fixed design elements with reconfigurable blocks and in many cases also mixed signal elements into Application Specific Integrated Circuits (ASICs). For commercial digital data processing chips, a trend towards increased mix of reconfigurable logic with fixed digital functions can be observed. This is where we see a major opportunity for future Digital Signal Processing in space applications.

In this paper, we first recall the basic technology trends for the implementation of data processing chains. We then summarize and compare the specific advantages and drawbacks of processor ASICs and FPGAs in the area of space based data processing. The advantages expected for systems on chip that integrate processors and FPGA fabric are explained, and typical application cases are addressed. The SoC design trade spaces for the mix of processor and FPGA IP are discussed, and relevant technology developments in Europe are summarized. The Scalable Sensor Data Processor (SSDP) as a promising technology basis for future, large flexible DSP SoCs is presented, and an architecture concept for a new high performance flexible European FPGA-equipped DSP is introduced.

I. Introduction

Digital Signal Processors have been popular devices for computation-intensive data processing for many decades. In comparison to General Purpose Processors (GPPs), their specific architectural designs support efficient processing of digital data via separate data and instruction memories, combined operations such as multiply-accumulate (MAC), hardware support for efficient loop execution, execution of multiple parallel operations (SIMD / VLIW), Direct Memory Access (DMA) mechanisms and other specific features. Ever increasing clock speeds and, more recently, many-core designs have led to significant performance increases, a trend that is still continuing. Recent chip developments for space applications such as the Scalable Sensor Data Processor (SSDP) [1] include the combination of GPP and DSP cores, combing their respective strengths in execution of control code and efficient processing of data samples.

On the other hand, programmable logic devices (PLDs) have been developed towards impressive levels of performance. Originally starting from relatively modest complexity level that allowed the implementation of glue logic and other specific circuitry, the recent generation of programmable devices, in the form of memory based Field Programmable Gate Arrays (FPGAs), allows not only to complement dedicated ASICs including GPPs and DSPs, but can replace them entirely in many application cases.

However, both FPGAs and ASICs have specific advantages but also drawbacks, which cannot be overcome by choosing one of these technologies while discarding the other. In the commercial world, an increasing number of products provide evidence for a trend to integrate reconfigurable logic with hard-wired functionality within complex SoCs. The combination of these technologies is expected to provide a maximum of application versatility and performance for future data processing systems and applications, including those for on-board payload data processing in space applications.

II. ASICs and FPGAs – A brief comparison

When implementing a digital design in ASIC or FPGA, a number of important considerations and tradeoffs apply. For ASICs, new developments or use of commercially available ASICs are possible. For FPGA, one is restricted to commercially available products. For the purposes of this comparison, we assume an identical or similar technology node (like 65nm, 180nm) for both technologies.

Application performance / speed is often a key requirement. Dedicated ASIC developments are providing the highest performances, often by a factor of 4 up to 10 higher than FPGA implementations. However, for projects that are restricted to commercially available parts and involve solutions that do not map well onto an existing ASIC products, FPGAs typically provide superior performances.

Power consumption is typically among the key drawbacks of FPGA solutions. Dedicated ASICs or commercial ASIC products that meet the performance specifications provide more power efficient solutions.

Radiation hardness is another factor favouring ASICs over re-programmable FPGAs. While the TID is typically adequate also for FPGAs, in the past ASICs have generally been superior in SEE tolerance. However, the application of TMR techniques and the use of antifuse FPGAs as well as the

upcoming use of flash-based FPGAs does enable the design of very robust FPGA based systems.

Development time is a key criterion in favour of FPGA based solutions. The manufacturing, testing and validation, and space qualification of a dedicated ASIC typically consumes between 1 and 2 years (even more in case of complications) of additional development time.

Flexibility is another key FPGA advantage. Late design changes, design modifications towards derived products, or even in-orbit re-programming are all possible with FPGAs.

Cost is a factor that depends on some key parameters. The NRE cost of an ASIC development is typically high (several M€), but with moderate or even high numbers of use cases the ASIC's unit cost may drop significantly below the cost of a comparable FPGA solution. For functions that are highly recurrent (GPP, GNSS ASICs, etc.) ASICs beat FPGAs easily. However, for one-time product development, or small series up to few 10 products, FPGA solutions are typically superior or competitive.

Today, it is in most cases the available time for product development, and the envisaged market size / sales volume for a product, that drives the decisions for development of a dedicated ASIC or an FPGA based solution. Commercially available ASICs (standard products) are typically used wherever they can meet the application's performance needs.

III. FPGA-EQUIPPED PROCESSORS

For space applications, the combination of processors with FPGAs has so far mostly been done on Printed Circuit Board (PCB) level. . More recently, FPGA dies have been combined with processor silicon in Multi-Chip Modules (MCMs) such as Intel's combination of commercial Xeon® processors with Altera® FPGAs [2], and the Atmel ATF697FF [3] which combines a AT697 SPARC processor with an ATF280 SRAM based FPGA for space applications.

In the commercial world, a next step – the integration of FPGA and GPP / DSP on the same die – is already taking place [4, 5], with a trend towards increasingly complex SoCs. It is common knowledge that commercial processor technology trends are arriving in the space processor technology area with a typical delay of 10-15 years. Therefore it is reasonable to assume that the integration of processors / DSPs and FPGAs on the same chip is about to arrive in qualified space ASICs around 2020, probably first in the form of Multi-Chip Modules (MCMs) or in FPGAs with integrated processor macros, and – possibly some years later - followed by DSPs and processors with integrated FPGA fabric. In the following paragraphs, we summarize the added value of such designs for space applications, and discuss the related tradeoffs and design spaces available to SoC developers.

A. Added value of FPGA integration

So far, the development of increasingly powerful processors and more and more capable FPGAs has been done independently, with HDL and software development done separately and often for similar applications. An efficient combination of FPGA and processor(s) in the same SoC would not only allow to combine the specific advantages of these technologies, but could also allow to re-use software and HDL from the traditional separate development lines for the specific elements where they perform best: software for control code, real-time needs, FDIR, fast application switching and FPGA reconfiguration, and HDL for high speed co-processing, glue logic, and interfacing.

It is therefore expected that the flexibility and application performance of a new DSP chip can be maximized when traditional processor architectures are combined with on-chip FPGA fabric. This is illustrated in Fig. 1.

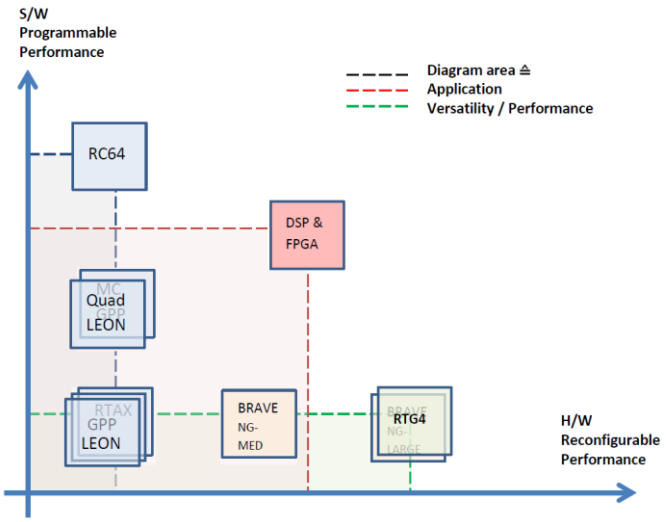

Figure 1: Postulated advantages of an FPGA equipped Many-core DSP chip

For space applications, many advantages of integrated FPGA fabric on space qualified DSP chips are obvious.

The hard-wired functionality (such as GPP, DSP cores, NoC infrastructure, interface IPs) could provide

- Reliable and efficient control and FDIR functions for the application
- Software programmable processing power for algorithms that map well on programmable IPs
- High energy efficiency and small footprint for these hard-wired elements
- GPP-like real time performances / properties (short reaction time in software via branch / interrupt)
- Management of the FPGA configuration as a soft-ware function via the GPP / control processor

On the other side, the on-chip FPGA fabric could provide

- Functionality in line with application needs (logic dominated or DSP slice dominated FPGA tiles)
- Reconfigurable and quickly re-programmable acceleration of specific functions
- Flexible glue logic for external companion chips such as ADC, DAC, accelerator chips
- Flexible logic for non-standard I/Os, tailored protocols, additional I/O functions

The integration of FPGA fabric on-chip would also reduce the pin-count and increase associated reliability (in comparison to solutions using separate processor and FPGA), reduce the power consumption, and lower the system cost, as DSP/GPP plus separate FPGA could be replaced by a single chip. The TID of the system would be uniform, and the SEE sensitivity / radiation related FDIR could be managed on chip level. The PCB footprint of the system would be reduced as well.

B. Typical Space Application Cases

Once a SoC that combines the advantages of hard-wired, software-based processors and FPGAs is available, it is expected that these features are exploited in a range of relevant application cases. These would include

- Data Processing Units (DPUs): The very high processing performance of both processor cores and FPGA fabric would be utilized; in addition, different and non-standard interfaces for DPUs on different spacecraft could easily be implemented without board / hardware modifications. In comparison to separate (processor + FPGA) solutions, the lower footprint, smaller pin count, lower power consumption, and uniform radiation hardness level would be advantageous.
- Future platform OBCs: Following a recent trend / desire to perform more subsystem related data processing tasks on the platform processor (example: star tracker / navigation software running on platform OBC), a platform processor will need high – and often specific - processing performance which could be provided by the on-chip FPGA fabric.
- Payloads and instruments: In many payload applications, an instrument controller (typically a GPP or DSP) is needed in combination with detectors / image sensors / ADCs / DACs. In such applications that are typically connected via FPGA. The envisaged SoC would allow to replace two chips by one, and provide the associated mass / power / footprint savings at lower cost and higher reliability.
- Telecom processors: in telecom applications, high bandwidth and processing power can be provided by the SoC's NoC, HSSL, and processor cores. Application specific processing such as codecs can be implemented in the FPGA fabric, and provide flexibility via re-programming and re-configuration.

In addition to these generic space-related application cases, other uses may exist for niche applications and in terrestrial areas like nuclear industry and specific R&D fields.

C. Tradeoffs and Design Spaces

For the design of a SoC that integrates hard-wired processors and FPGA fabric, a wide design space exists between the extremes of processor-only and FPGA-only designs. Some of the intermediate options have been adopted by commercial products already; these include the Xilinx Zynq FPGAs [6] which combine quad-core and / or dual-core processors with FPGA fabric, and some variations of the Xilinx Virtex5 family which includes PowerPC processor cores. For processors with small on-chip FPGA, so far no commercial products have surfaced. For space, as of early 2016 no qualified products featuring on-chip FPGA and hard-wired processors exist. Future versions of the European FPGAs that are now under development may include a processor core, but with a die footprint that is dominated by the FPGA fabric. For the FPGA, further trade-offs between logic-dominated or DSP slice dominated fabric are needed.

Therefore, the overall design space for FPGA-equipped SoCs for space applications remains unexplored to a significant degree, and invites associated R&D towards an optimized mix of hard-wired and reconfigurable SoC elements.

IV. Recent Technology Developments

A number of relevant technology developments have been performed or started in ESA contracts in the past few years. Here, we provide only a brief summary; more information is available in the provided references, including [7].

A. DSP and NoC IP and related Chips

Recent ESA DSP chip and IP core developments can be traced back to a study performed from 2008 onwards, called "Massively Parallel Processor Breadboarding Study", ESA contract 21986 [8]. In this study, a heterogeneous multi-core architecture (2 DSPs, 1 GPP) have been prototyped successfully together with a proprietary NoC, space standard interfaces, and other features. This was followed by an ASIC prototyping activity that proved the DSP cores, NoC, and other key features in rad-hard silicon ("DARE plus – ASICs for Extreme Radiation Hardness and Harsh Environments", ESA contract Nr. 4000104087) [9]. The most recent development in this line is the Scalable Sensor Data Processor (SSDP) ASIC (ESA contract Nr. 4000109670). This chip provides a LEON3 SPARC compatible GPP and two NoC-connected VLIW Xentium® DSPs. At a system clock of 100 MHz, the chip provides up to 1600 MOps (800 MMACs) of DSP processing power, space standard interfaces, and a wide range of digital and mixed signal features that are expected to be attractive for applications in space instrumentation, data processing units, robotics, and various types of spacecraft subsystems. SSDP will be available world-wide as a commercial product via Cobham Gaisler [10]. More information on SSDP is provided in a companion paper [1].

Based on the DSP and NoC IP used in MPPB and SSDP, additional developments are ongoing for the development of a floating point version of the Xentium ® DSP (CCN to ESA contract 21986) as well as an advanced version of the NoC (ESA contract 4000115252/15/NL/LF) which provides enhanced features for reliability, radiation hardness, and FDIR. Both developments are expected to be completed in 2017. A comprehensive overview on DSP and NoC IP, related software, and ongoing ESA contracts is available in a companion paper [11] and from [7].

B. FPGA IP and chips

ESA and CNES, together with European space industry, have been working for several years on the European high-performance, radiation hardened reprogrammable FPGA family know under the acronym BRAVE for "Big Reprogrammable Array for Versatile Environments".

The first FPGA of the family, under development since 2015, is the NG-MEDIUM (also known as NXP-32000) under ESA contract 4000113670/15/NL/LvH [12]. NG-MEDIUM is based on the STMicroelectronics 65nm ASIC technology and the associated radiation-hardened library. In order to achieve high-density (i.e. high capacity) and high performance, most of the NG-MEDIUM design has been created using a full-custom flow (i.e. not using the radiation-hardened library). Radiation hardening has been the main focus, from the full-custom cell-level to the system-level. For instance, at system level it includes EDAC for the internal memories, as well as a transparent configuration memory integrity check. The NG-

MEDIUM has a logic capacity similar to the Microsemi® RTAX4000 FPGA with the advantage of including 112 hard-macro DSP blocks that bring high-performance for any algorithm requiring arithmetic operations.

The configuration of the BRAVE FPGA family is based on radiation-hardened SRAM-memory cells, which provide unlimited re-programmability of the FPGA. The re-programmability allows BRAVE based systems to perform different processing functions, and enables the implementation of adaptive systems for space.

The first samples of NG-MEDIUM will be available in Q4 2016, with an expected qualification in Q4 2017. The next FPGA of the family, called NG-LARGE, will use the same 65nm ASIC technology and will have a complexity similar to the RTG4 FPGA: the first samples are expected to be available in 2017. There are plans to have the third FPGA of the family, NG-ULTRA, in 2018.

For integration into larger SoCs, it is possible to use an embedded FPGA IP (eFPGA) based on the NanoXplore BRAVE family. There are other European eFPGA solutions (Menta and Adicsys) based on the use of standard-cell digital flow: they are more portable solutions across ASIC technologies, but they provide less density and performance. Future SoC designers therefore have several options available for their developments.

V. An Architecture Concept for an FPGA equipped High Performance DSP

There are multiple options for the integration of FPGA fabric with contemporary processor and DSP IP. For space applications and related product acceptance, an evolutionary approach that supports both software and HDL re-use is considered advantageous. For ESA, such an approach would suggest the evaluation of integration possibilities with recently developed DSP core and NoC IP technology performed by European companies. In this paper, we therefore consider an example design (tentatively called "FlexDSP") that could be derived from the latest European Space DSP, the Scalable Sensor Data Processor [1], in a rather straightforward and evolutionary way.

A. SSDP

The SSDP has already been mentioned in chapter IV A, and details can be found in the corresponding companion paper [1]. Here, we limit the description to those elements and concepts that are relevant for the derived "FlexDSP" concept.

SSDP is based on an architecture that uses a GPP (LEON3) to control a NoC based subsystem (called "Omnific", illustrated in Fig. 2) via an AHB bus bridge and interrupt signals. All high bandwidth data traffic takes place within the NoC subsystem; the GPP controls data traffic and related DMA transfers, assignment of jobs (data and binary program code) to DSP cores, use of buffers, and overall system con-figuration. Once DSP jobs or data transfers are completed, the GPP is notified via interrupt and can schedule the next activity within a data processing sequence. Some of these tasks (like control of DMA transfers, or synchronization of DSP cores) may be delegated to the DSPs in case a lower

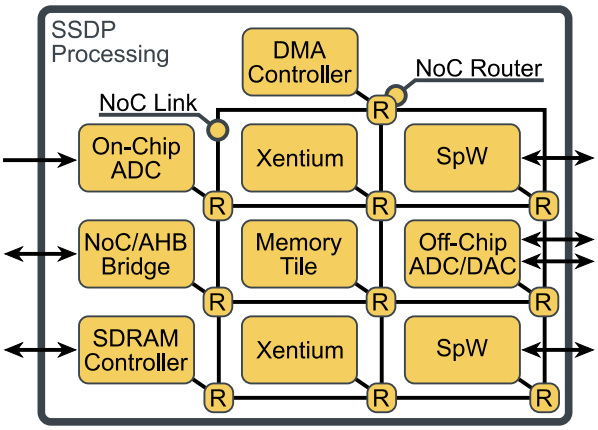

Figure 2: SSDP's "Omnific" NoC Subsystem

GPP interrupt rate is desirable; this may be especially useful for future SoCs with significantly larger numbers of DSP cores and NoC nodes, including the "FlexDSP" example.

DSP cores are controlled via their network interfaces (NIs). These interfaces provide a number of registers ("mailboxes") which are used to pass information to the DSP, and for assigning the DSP's program code to be executed. A typical DSP activity sequence may look as follows:

- Input data is transferred to the DSP's local data memory via DMA controlled by the GPP.
- The location (address) of the DSP's assigned program code is written to the NI's mailbox.
- The NI then fetches the code and transfers it to the DSP's instruction cache; once this is completed, the DSP is started and code execution commences.
- Once the DSP code execution is complete, the GPP is notified via interrupt and the DSP core changes to idle mode.
- Output data in the DSP's local memory is transferred to a suitable location via DMA under GPP control.
- A new sequence as above can be initiated.

This concept is scalable to larger NoCs within certain boundaries, and is re-used in the "FlexDSP" concept introduced in the next section.

B. FlexDSP

Taking an evolutionary path towards future SoCs, and specifically towards an FPGA-equipped many-core DSP as discussed here, is desirable for a number of reasons, including

- Re-use of available GPP code and tools (LEONx family, SSDP) to reduce cost / increase maturity
- Re-use of available DSP code and tools (SSDP) to reduce cost / increase maturity
- Re-use of existing HDL for the embedded FPGA fabric (glue logic to ADC/DAC, codecs, etc)
- Reducing development risk by re-using accessible, well known and validated IP cores

- Exploiting the accumulated expertise and familiarity with IP and underlying concepts available in the user community

A concept for a new, FPGA-equipped many-core DSP could therefore be based on the following features:

- GPP controller with local memory and associated subsystem
- NoC with routers, bridges, DMA, network interfaces
- DSP cores with their local memories / caches, connected via network interfaces
- One or more FPGA tiles, connected via network interfaces
- NoC connected memory tiles and external memory controllers

In order to re-use the software concepts developed for SSDP, the FPGA tiles could have an interface that supports a utilization in a way that is similar or identical to that of the DSP cores. The network interface could provide access to the FPGA's integrated memories (corresponding to the DSP core's local data memory) and support loading the FPGA's configuration bitstream (corresponding to loading code into the DSP core's instruction cache). The NI would also facilitate access to NoC connected memories for the FPGA, and provide configuration registers and interrupt outputs to the GPP that controls the system.

The type of FPGA configuration memory could be chosen between SRAM and FLASH; from a first assessment it seems that SRAM based FPGA might be more advantageous, as SRAM cells are available in the relevant DSM ASIC processes (65nm, possibly 28nm) and would support fast re-configuration as well as unlimited numbers of re-programming cycles.

Final memory choice and aspects such as radiation effects mitigation will need further study.as radiation effects mitigation will need further study. From the foregoing considerations, a concept arises that is depicted in Fig. 3. It extends the basic architecture known from SSDP towards a larger number of cores, and integrates FPGA fabric that is connected to the NoC and to external I/Os. In all application cases, the SoC would start up by booting the GPP which then configures the on-chip resources. Then, depending on the application, two distinct types of FPGA fabric utilization can be considered:

Data processing with frequent re-configuration

Here, the FPGA fabric is used for data processing under the control of the GPP. A configuration bitstream would be assigned and loaded via the NI. Data can either be provided via the NI-connected local FPGA memories, or stored in memory tiles or external memories and accessed via the NI. The FPGA fabric would then process the data in line with its configuration, and notify the GPP via interrupt once the output data is ready. The GPP can then assign new data and either re-start the processing, or assign a different configuration bitstream for a different type of job. The FPGA bitstream may be loaded via the NoC from either on-chip memory tiles, or from external (presumably DDR2/3 SDRAM) memory. It is expected that in this scenario, which is depicted in Fig. 3, FPGA reconfiguration can occur in a very short time span, much faster than in contemporary SRAM based FPGAs for which the bitstream is provided via a comparatively slow interface from external memory.

Static interfacing, pre-processing and glue logic

For applications where the on-chip FPGA is used for static functions like interfacing to external chips (ADC, DAC, co-processor chips), provision of custom / additional interfaces (non-standard interfaces of specialized ASICs, additional SpW / SPI / CAN / other interfaces) or simple pre-processing (averaging / filtering / other pre-processing of input data) the FPGA bitstream is loaded only once at the time of system configuration, and kept static during the application execution. For a product, interface configuration bitstreams for popular companion devices might be provided with the DSP's SDE. The mitigation of radiation effects (configuration memory scrubbing) would be performed frequently under GPP or NI control. This scenario is depicted in Fig. 4. It must be noted that the number of DSP cores, on-chip memory tiles, IOs, and FPGA tiles depicted in the example (Fig. 3 and Fig. 4) are chosen here for the purposes

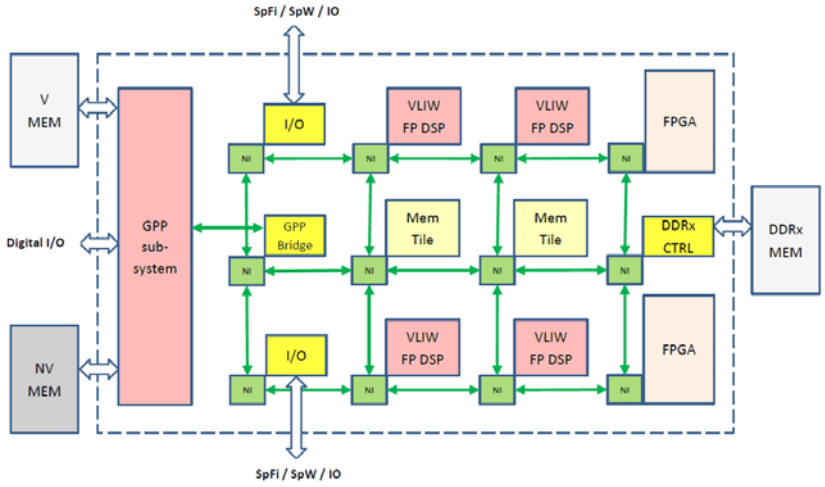

Figure 3: FPGA-equipped many-core DSP example (FlexDSP") configured for reconfigurable processing

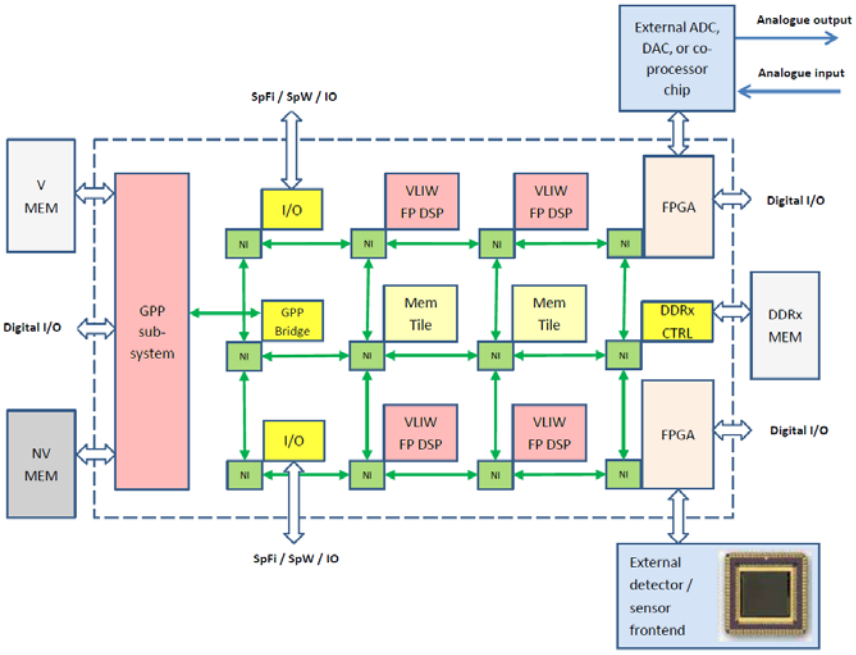

Figure 4: FPGA-equipped many-core DSP example ("FlexDSP") configured for static interfacing

of concept illustration, and are not the outcome of a requirements analysis and trade-off that would be necessary for a product development.

The performances that can be achieved for such a SoC increase with the chosen number of DSP cores and FPGA tiles. The following performances can be expected for the individual SoC elements when implemented in a 65nm ASIC process:

- System clock (GPP clock, DSP core and NoC) of ca. 250 MHz
- DSP cores providing ca. 1 GFLOPS per core (2 MACs per clock cycle, 32 bit floating point)
- Ca. 8 Gbps NoC speed (32 bit parallel lanes, bi-directional links)
- FPGA performance varying with fabric size and type; several to several 10 GOPS can be expected
- Process typical TID, radiation hardened design, 6.25 Gbps HSSL, DDR2/3 expected

Upcoming 28 nm processes would further boost the achievable performance and allowable SoC size. Related considerations are however out of scope for this paper.

VI. Summary

DSP ASICs and FPGAs enable different solutions for space based data processing applications. Both technologies have their advantages and drawbacks. Following a trend that started in the commercial world, it is expected that an optimized combination of hardwired GPP / DSP functionality with FPGA fabric on the same chip will enable more power-, size-, mass-, and cost-efficient solutions for future space applications. The design space for such future SoCs is still largely unexplored and invites further study, followed by prototyping activities that enable future product developments.

The underlying technologies (GPP, NoC, DSP cores, FPGA fabric, DSM ASIC processes) are available or under development in Europe, and would support the envisaged SoC developments in the near future. A first NoC-based DSP chip is under development in Europe, and its scalable architecture provides a possible basis for future SoCs that integrate FPGA fabric. A conceptual example for a large, European FPGA-equipped NoC based DSP SoC has been introduced. Its key features and possible operating modes have been explained, and areas inviting further study have been identified.

At ESA, further studies on FPGA integration into GPP and DSP SoCs are planned, and are expected to inform future design decision for the next generation of high performance components for space based data processing applications.

VII. References

[1] Scalable Sensor Data Processor: Architecture and Development Status, R. Pinto, L. Berrojo, E. Garcia, R. Trautner, G. Rauwerda, K. Sunesen, S. Redant, S. Habinc, J. Andersson, J. López, ESA DSP Day, Gothenborg, 2016.

[2] http://www.eweek.com/servers/intel-begins-shipping-xeon-chips-with-fpga-accelerators.html

[3] ATF697FF Rad-hard reconfigurable processor with embedded FPGA, http://www.atmel.com/devices/ATF697FF.aspx

[4] https://www.altera.com/products/soc/overview.html

[5] http://www.microsemi.com/products/fpga-soc/soc-fpga/smartfusion2#overview

[6] http://www.xilinx.com/products/silicon-devices/soc.html

[7] R. Trautner, Development of New European VLIW Space DSP ASICs, IP cores and related software via ESA contracts in 2015 and beyond, Proceedings of DASIA, Barcelona, 2015

[8] Walters, K.H.G. and Gerez, S.H. and Smit, G.J.M. and Rauwerda, G.K. and Baillou, S. and Trautner, R., Multicore soc for on-board payload signal processing., AHS, 2011.

[9] G. Thys et al., Radiation Hardened Mixed-Signal IP with Dare Technology, AMICSA, 2012

[10] Cobham Gaisler, www.gaisler.com

[11] Multi-Core DSP sub-system OIP, G. Rauwerda, K. Sunesen, T. H. Thanh, J. Potman, ESA DSP Day, Gothenborg, 2016.

[12] O. Lepape, NanoXplore NXT-32000 FPGA, SpacE FPGA Users Workshop, 3rd edition, ESA/ESTEC, Noordwijk, 2016.

VIII. Acknowledgements

The authors would like to thank K. Sunesen (Recore Systems, NL), A. Notebaert (Airbus D&S, F), and P. Lombardi (Syderal, CH) for their inputs and contributions to this paper.