

Dataflow-Functional High-Level Synthesis for Reconfigurable Accelerators

Dr. Loknath Sarangi,
College of Engineering Bhubaneswar

Abstract— Nowadays, domain-specific acceleration is essential across the whole computing spectrum, from embedded systems to high performance computers. Unfortunately, from the perspective of design productivity, system specialization is by definition a nightmare. However, the combination of coarse-grained reconfiguration (CGR) architectures and dataflow (DF) models of computing might be especially useful in situations when the kernels to be accelerated are inherently streaming oriented. We provide in this letter an innovative approach, based on DF and functional programming principles, to assemble and characterize virtually reconfigurable accelerators. This methodology can address design productivity difficulties for CGR accelerators. When compared to state-of-the-art high-level synthesis, the primary benefit of the suggested methodology is accurate IP-level latency predictability that improves design space exploration.

I. CONTEXT, BACKGROUND, AND MOTIVATION

FLEXIBILITY and performance are two highly valued properties of processing systems in many application domains. Modern systems include ubiquitous and upgradable devices for IoT or cyber-physical applications that also communicate with cloud infrastructures. These systems must adapt to mutable conditions, while avoiding unpredictable performance degradation.

To boost performance, high performance computing (HPC) and embedded systems designers are pushed to opt for

domain-specific accelerators (DSAs), built from ad-hoc hardware-coded kernels exploiting parallelism for optimizing performance. Two representative examples of this trend are Amazon Web services (AWS) and Movidius. On the HPC side, AWS offers field programmable gate arrays (FPGAs) in a cloud environment [2]. On the embedded side, the Intel Movidius vision processing unit is a low-power DSA used in smartphones and drones for computer vision and artificial intelligence applications.

In terms of system flexibility, hardware DSAs based on coarse-grained reconfigurable (CGR) architectures offer flexibility by adapting to variable application parameters. CGR architectures are traditionally composed of a mesh of processing elements (PEs) whose interconnections are reconfigured over time [16] to offer flexibility. CGR architectures inherit from their hardware nature the capacity of executing compute-intensive kernels in an energy-efficient way. However, efficiency and flexibility offered do not come for free. Hardware design is a complex and error-prone task, leading to productivity losses, especially for heterogeneous and irregular targets. High-level synthesis (HLS) approaches have been proposed to cope with these losses, obtaining design productivity gains by separating functional system verification, performed from a time-agnostic high-level language, from timed system verification, performed after automatically inferring hardware-specific code [7].

Many HLS tools are now available, such as Xilinx Vivado HLS [17] and Intel FPGA SDK for OpenCL [4]. These tools are based upon imperative, C-like, languages. The algorithm to be implemented is formulated as a *sequence* of instructions operating on mutable data. This choice is motivated by the very large number of developers trained to manipulate imperative languages that have been dominating computer sciences for decades. But, from a hardware perspective, this choice presents two major drawbacks.

- 1) Imperative formulations generally do not distinguish iterations over time from iterations over space, which do not translate uniformly in hardware (the latter do not imply causality and can therefore be parallelized using replication).
- 2) They implicitly rely on the concept of global memory at the implementation level, which immediately leads to a “bottleneck” on memory accesses [1].

These drawbacks can be circumvented by relying upon so-called applicative or functional languages in which algorithms are described as a (mathematical) composition of side-effect free functions. This approach naturally fits dataflow (DF) models of computation (MoCs). An application is decomposed into independent processing actors, communicating with first-in first-out data queues (FIFOs), with no global storage or synchronization. This is particularly true for *stream processing* applications, processing data “on the fly” and which benefit

significantly from CGR architectures-based acceleration, as found in signal/image processing, media coding/compression, cryptography, video analytics, etc.

This letter, in Section II, introduces a novel methodology for the optimal characterization of virtually reconfigurable DSAs exploiting a DF functional (DF) HLS approach as an alternative to traditional HLS tools based on imperative languages. The proposed methodology is targeted for heterogeneous and irregular CGR architectures, leveraging on application specific PEs and tailoring the interconnect to minimize FIFOs. For prototyping purposes, and to compare with commercial flows, experiments target FPGA technologies and demonstrate that, thanks to its modularity and abstraction capabilities, the proposed approach guarantees latency predictability (see Section III). Beside we also show (Section IV) that our DFF approach leads to improved performance for DSAs with CGR architectures.

II. PROPOSED DATAFLOW-FUNCTIONAL HLS

The proposed framework is built from two existing tools developed by the authors. The Caph just Aint Plain HdL (CAPH) compiler and the multidataflow composer (MDC) toolset. CAPH [12] is an open-source, domain-specific framework for the specification, simulation, and implementation of stream-processing applications based on a dynamic DF MoC. Applications are specified as DF networks (DFNs) using a higher-order, polymorphic functional language. The behavior of each DF actor is defined as a set of transition rules using pattern matching on structured data, resulting in improved abstraction capabilities. The CAPH toolchain provides graphical visualization of DFNs; code simulation with trace facilities; and HLS producing SystemC code for simulation and resource-monitoring purposes, as well as platform-agnostic ready-for-synthesis VHDL code. CAPH provides 1:1 DFN PEs synthesis, but no support for reconfiguration. MDC is a suite for the design and development of heterogeneous and irregular DSAs based on a set of DFNs to specify the desired behavior of the system [6]. It builds reconfigurable datapaths, leveraging on a two-step approach.

- 1) A *model-to-model compiler*, the multidataflow generator (MDG), derives a multifunctional DFN [10] starting from a set of disjointed DFN specifications.
 - 2) A *DF-to-hardware mapper*, the Platform Composer (PC), deploys an HDL description of the CGR datapath.
- Up to this letter, MDC was specific to RVC-CAL language and the actor communication protocol was hardwired on the resulting CGR architecture. Now users define custom communication protocols, and MDC accordingly implements the handshake among PEs. MDC provides N:1 DFNs mapping on a CGR DSA, but it does not support PE synthesis.

Based on the described complementarities, this letter proposes a fully automated toolchain, integrating MDC and CAPH, for specifying and deploying CGR DSAs. The toolchain architecture is depicted in Fig. 1. The proposed DFF HLS flow consists of three main phases.

- 1) *Composition*: Model-to-model compilation performed by MDC MDG, taking as inputs generic DFNs. Output: high-level multidataflow DFN of the DSA.
- 2) *Optimization*: Optimal sizing of the actor-connecting FIFOs. Optimization starts from an estimation produced by the CAPH Compiler SystemC backend, which is adapted to the multifunctional DFN case by worst case analysis.
- 3) *Generation*: Deployment of the CGR DSA. The MDC PC component outputs a top-level HDL module,

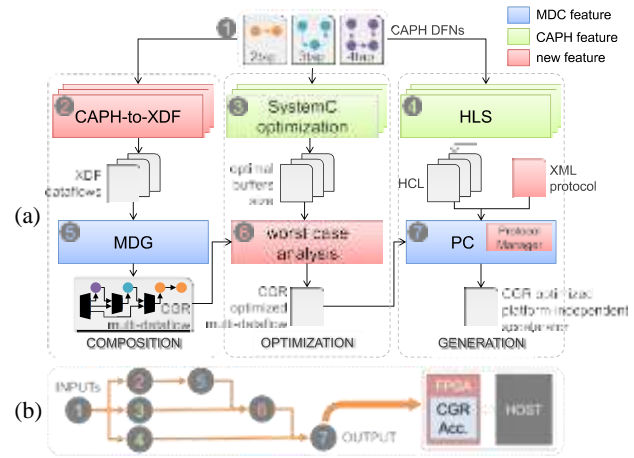


Fig. 1. Proposed flow: (a) components; and (b) execution sequence.

TABLE I
 PROPOSED FLOW VERSUS STATE-OF-THE-ART

| | [8] | [15] | [9] | this work |
|--------------------------|---------------|---------------|-------------|------------|
| Framework: Input Specs | C | C | dataflow | dataflow |
| Framework: (Re-)mapping | Y | Y | Y | Y |
| Framework: PE generation | N | Y | Y | Y |
| DSA: technology | custom device | custom device | Xilinx FPGA | ASIC, FPGA |
| DSA: heterogeneity | N | Y | Y | Y |
| DSA: topology | fixed | fixed | custom | custom |

corresponding to the optimized multifunctional DFN, using the CAPH generated hardware component library (HCL) actors.

Fig. 1(b) summarizes the flow steps. (1) Users are required to provide the input set of specifications (N different DFNs) to be accelerated using the CAPH language. Starting from these inputs, three parallel steps take place to make CAPH DFNs compliant with MDC through the CAPH-to-XDF parser (2), optimize sizing of the buffer connecting DF actors (3), and generate the target-independent HCL (4). Three more steps come in succession. The N XDF networks produced by (2) are merged in a multidataflow network (5), which is optimized (6) using the worst case analysis results of (3). Finally, the CGR platform-independent DSA is deployed in (7). During the merging process (5), common actors are identified and connections minimization is guaranteed. In the CGR DSA, combinatorial switching elements are used to access shared modules, this may affect frequency but not latency.

III. PROPOSED APPROACH: FEATURES AND NOVELTY

CGR architectures provide the flexibility needed by modern signal processing applications, while achieving high efficiency through specialization. Their features and granularity are highly varied [3], and several works in literature address their design issues. Raffin *et al.* [8] proposed the ROMA framework for the scheduling, binding and routing of reconfigurable accelerators for multimedia applications. The ROMA architecture is composed of custom reconfigurable PEs and the framework is specific to them, following an approach similar to general purpose devices, thus losing specialization and, in turn, efficiency. Voros *et al.* [15] leveraged on a fixed CGR structure that comes along with its synthesis flow, and where three different types of PEs can be used for accelerating purposes. Our effort focuses on a more specific problem: we discuss in this letter a flow to characterize and synthesize circuits to be embedded in a DSA infrastructure. MDC has been combined

in previous work with the Xronos HLS to offer DFF HLS to build a Xilinx-specific DF-to-hardware design environment [9]. Table I reports a qualitative comparison among this letter and other frameworks available in literature for CGR DSAs.

CGR DSAs can also be obtained with sole HLS, while modeling reconfiguration by hand. State-of-the-art HLS frameworks are typically imperative-based, and most of them are also target-dependent [5]. As we are going to demonstrate in Section IV, imperative-based HLS is not the best choice for developing CGR DSAs, differently from DFF HLS. Synflow Studio [14] provides a proprietary HLS framework for multi-vendor FPGAs using the Cx DF language, which does not support ASICs. Platform-agnostic HDL specification is challenging, as the resulting hardware has difficulties to compete with platform-specific code performance. To the best of our knowledge, only two academic flows support platform-agnostic DF-oriented HLS: the Orcc HDL backend [13] leveraging on the RVC-CAL language, and CAPH [12]. CAPH, adopted in this letter, has the advantage of using functional programming semantics that, contrary to the RVC-CAL language, removes imperative semantics from actors. The Orcc HDL backend is no longer available, thus it is not possible to compare performance with [13].

Performance predictability is important in a design flow. Indeed, knowing the properties, and in particular the latency, of the implemented system at early design stages improves productivity and reduces design effort. This is a key feature of our approach. HLS tools based on imperative languages can only provide IP latency estimates *after synthesis* when a reconfigurable system is designed. In Vivado HLS [17] this estimate takes the form of a range (min to max latency): the actual value of each single configuration in a reconfigurable case is obtained by running the corresponding mode over the deployed CGR accelerator. The Intel FPGA SDK for OpenCL [4] pursues automatic system-level integration (and parallelization) of hardware accelerators complying the OpenCL computing model. As a result, low-level, cycle-accurate pipeline/scheduling information, typical in other IP-based HLS tools, is sacrificed in favor of an annotated dependence graph describing the generated kernel as a combination of blocks. The estimation provided in the graph refers to the types/sizes of load/store units, stalls and latencies, but a nontrivial calculation of the IP latency is left to designers. Finally, optimization is performed by profiling kernel execution, requiring resynthesizing an equivalent kernel, automatically instrumented with performance counters, to analyze memory access behavior (stall, occupancy, and bandwidth).

The key selling feature of the proposed DFF flow is the ability to compute early and accurate latency¹ estimations *before synthesis*. Such values depend on the critical path length in the DFN and on the maximum number of cycles required for each actor firing. For synchronous DF (SDF) graphs, both can be computed statically. For non-SDF graphs, CAPH estimates latency running the cycle-accurate code generated by its SystemC backend. The latency of the CGR-based DSA generated by the DFF flow depends on the selected mode only, corresponding to the selected stand-alone networks, since merging does not affect latency. As a result, the DSA can be optimized in terms of latency *before synthesis*. To the best of our knowledge, neither [8] nor [15] present this predictability feature, while [9] handles SDF inputs only.

In general, *presynthesis* prediction offers an additional advantage: developers can focus on the application to be

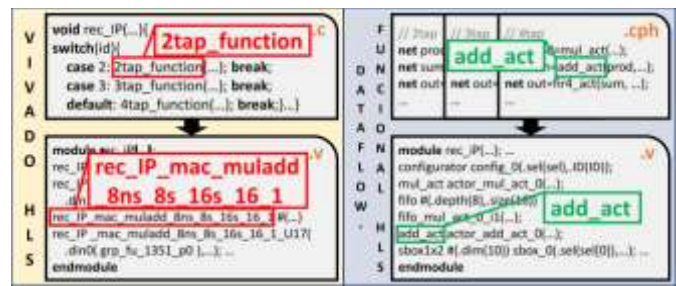


Fig. 2. Proposed flow versus Vivado HLS: starting and generated codes.

accelerated with an *algorithmic-oriented approach* rather than a synthesizer-based one. Imperative-based HLS tools require design space exploration (DSE) to improve latency. Developers have to deal with the mentioned lack of information and to understand how to exploit both the available #pragmas and their combinations toward best performance. Moreover, code refactoring may also be necessary to achieve optimal latency values. For these reasons, developers must have a thorough knowledge of the synthesizer and a considerable effort is required to obtain an optimized architecture. To conclude, when compared to imperative HLS, the DFF nature of the proposed HLS provides earlier explorations, while offering system predictability and guaranteeing performance, as demonstrated in the next section.

In summary, with respect to the context of CGR DSAs, the main benefits/features of the proposed integrated flow are as follows.

- 1) *Custom PE Generation*: Hardware generation considers heterogeneous, HLS-generated PEs for each DF actor, favoring flexibility with respect to [8].
- 2) *Reconfigurability Management*: DF-based mechanism maximizes and controls resource reuse, leveraging on datapath merging techniques. Moreover, reconfiguration management is guaranteed by MDC.
- 3) *Predictability*: Modular specifications facilitate the predictability of system properties. *Before synthesis* latency estimations can be carried out on the basis of the preprocessed CAPH DFNs.
- 4) *Target Independence and Availability*: Both MDC and CAPH are platform-agnostic and open source.
- 5) *Code Readability*: MDC and CAPH preserve the correspondence among DF actors and hardware PEs. This is important for example in case of post-HLS enhancements. Vivado HLS acts differently since it assigns IPs to functional units (see Fig. 2).

IV. EXPERIMENTAL RESULTS

This section compares results obtained with three design flows: 1) the proposed DFF one, using Vivado v2015.2 and Quartus v17.1 to synthesize the CGR DSAs; 2) Vivado HLS v2015.2; and 3) Intel FPGA SDK for OpenCL v17.0. Target devices are Xilinx Virtex 7 and Intel Cyclone V FPGAs. The test case chosen is a 1-D/2-D HEVC 3/5/8-tap hardware accelerator for approximate image interpolation filters, in fixed and reconfigurable configurations (with runtime filter switching) [11]. In the case of Xilinx/Intel HLS tools, a moderate effort has been put into code refactoring for optimization using shift register inference, line buffers, resource reuse when possible, and pragmas (loop unrolling and pipeline). For this reason they involve two reconfigurable designs: baseline (R) and optimized (R*). A comparison among the different flows is proposed in terms of resource utilization and early latency estimation.

¹Number of clock cycles required to compute all outputs as in [17].




Fig. 3. Proposed flow versus (a) and (b) Xilinx (XC7VX485T) and (c) Intel (5CSEMA5) target FPGAs. Intel/Xilinx resources: REG=register/FF, LOGIC=ALUT/LUT, RAM=M10K/BRAM, and DSP=DSP/DSP. Data are shown on a logarithmic scale with base of 10. Latency values in clock cycles are reported in (a).

DFF HLS Versus Vivado IP-Level HLS: Fig. 3(a) shows that for standard implementations (indicated by 3/5/8/R) the DFF flow achieves better latency results. The very large latencies obtained with Vivado HLS are explained by the default area-driven synthesis (which required adding #pragma and code refactoring). Moreover, as stated in Section III, our flow preserves the latency values of the reconfigurable IPs, meaning that merging does not alter latency as opposed to other synthesizers and, on top of that, before synthesis estimations are equal to actual post-synthesis results. With respect to the optimized reconfigurable versions (R*), only for the 1-D filter a lower use of resources [see Fig. 3(b)] and latency has been obtained using Vivado HLS (−72.0% register count, REG; −76.1% logic elements, LOGIC; +0.5% for minimum and 14.8% for maximum latency). In the 2-D filters case our performances are superior. Latency results are more than three times larger in the best Vivado HLS case, and again their value is: 1) not preserved and 2) highly variable in the reconfigurable case.

DFF HLS Versus OpenCL System-Level HLS: As in the Vivado HLS case, a significant effort on code refactoring is needed for optimization, which is driven by an optimization report offered at early design phases (it contains highly inaccurate resource and system-level pipeline latency/stalls estimations, rather than individual kernel latency values). Hence, there is no sane way to accurately calculate kernel latency and have insights into the generated datapath; therefore, no latency results are reported in Fig. 3(a). For resource usage results refer to the kernel instantiation [See Fig. 3(c)]. The support for the OpenCL model even at kernel level (automatic pipeline to support various threads on the fly, memory accesses optimization, replications, etc.) introduces significant resource overhead.

V. CONCLUSION

The findings in this letter point to DFF HLS as a possible substitute for traditional imperative HLS in the construction of flexible and reliable CGR DSAs. In comparison to commercial HLS techniques, competitive resource use is attained, permitting precise datapath delay estimates and presynthesis. Mainstream HLS tools have likely overlooked an important aspect of hardware design that is crucial to embedded systems in their attempt to make hardware design easier for software developers: having accurate information on the generated datapaths to better optimize hardware accelerators through improved DSE. We illustrated the delay problem and suggested a different design process that outperforms conventional HLS tools in this regard. DFF HLS can address these problems by elevating the

abstraction level, allowing low-level performance estimates to be adjusted further on the hardware. For the purpose of conciseness, we have included the entire assessment of a single accelerator in this letter. The suggested analysis will be expanded upon in later works to include additional test cases and system features.

REFERENCES

- [1] J. Backus, “Can programming be liberated from the von Neumann style? A functional style and its algebra of programs,” *Commun. ACM*, vol. 21, no. 8, pp. 613–641, 2007.
- [2] J. Barr. (2017). *EC2 F1 Instances With FPGAs—Now Generally Available*. [Online]. Available: <https://aws.amazon.com>
- [3] K. Compton and S. Hauck, “Reconfigurable computing: A survey of systems and software,” *ACM Comput. Survveys*, vol. 34, no. 2, pp. 171–210, 2002.
- [4] *Intel FPGA SDK for OpenCL*, Intel, Santa Clara, CA, USA, 2018. [Online]. Available: <https://www.intel.com>
- [5] R. Nane *et al.*, “A survey and evaluation of FPGA high-level synthesis tools,” *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 35, no. 10, pp. 1591–1604, Oct. 2016.
- [6] F. Palumbo, T. Fanni, C. Sau, and P. Meloni, “Power-awareness in coarse-grained reconfigurable multi-functional architectures: A dataflow based strategy,” *J. Signal Process. Syst.*, vol. 87, no. 1, pp. 81–106, 2017.
- [7] M. Pelcat, C. Bourrasset, L. Maggiani, and F. Berry, “Design productivity of a high level synthesis compiler versus HDL,” in *Proc. SAMOS*, 2016, pp. 140–147.
- [8] E. Raffin *et al.*, “Scheduling, binding and routing system for a run-time reconfigurable operator based multimedia architecture,” in *Proc. DASIP*, 2010, pp. 168–175.
- [9] C. Sau *et al.*, “Automated design flow for coarse-grained reconfigurable platforms: An RVC-CAL multi-standard decoder use-case,” in *Proc. SAMOS*, 2014, pp. 59–66.
- [10] C. Sau, L. Fanni, P. Meloni, L. Raffo, and F. Palumbo, “Reconfigurable coprocessors synthesis in the MPEG-RVC domain,” in *Proc. Int. Conf. ReConFig. Comput. FPGAs (ReConFig)*, Mexico City, Mexico, 2015, pp. 1–8.
- [11] C. Sau *et al.*, “Challenging the best HEVC fractional pixel FPGA interpolators with reconfigurable and multifrequency approximate computing,” *IEEE Embedded Syst. Lett.*, vol. 9, no. 3, pp. 65–68, Sep. 2017.
- [12] J. Sérot, F. Berry, and S. Ahmed, “CAPH: A language for implementing stream-processing applications on FPGAs,” in *Embedded Systems Design With FPGAs*. New York, NY, USA: Springer, 2013.
- [13] N. Siret, I. Sabry, J. F. Nezan, and M. Raulet, “A codesign synthesis from an MPEG-4 decoder dataflow description,” in *Proc. ISCAS*, 2010, pp. 1995–1998.
- [14] Synflow. *Synflow Studio*. Accessed: May 1, 2018. [Online]. Available: www.synflow.com
- [15] N. S. Voros *et al.*, “MORPHEUS: A heterogeneous dynamically reconfigurable platform for designing highly complex embedded systems,” *ACM Trans. Embedded Comput. Syst.*, vol. 12, no. 3, 2013, Art. no. 70.
- [16] M. Wijnvliet, L. Waeijen, and H. Corporaal, “Coarse grained reconfigurable architectures in the past 25 years: Overview and classification,” in *Proc. SAMOS*, 2016, pp. 235–244.
- [17] *Xilinx Vivado HLS*, Xilinx, San Jose, CA, USA, 2018. [Online]. Available: www.xilinx.com