

Article

Dynamically-Tunable Dataflow Architectures based on Markov Queuing Models

Mattia Tibaldi ¹ , Gianluca Palermo ¹ , and Christian Pilato ^{1,*} ¹ Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria, Milano, Italy

* Correspondence: christian.pilato@polimi.it

Abstract: Dataflow architectures are fundamental to achieve high performance in data-intensive applications. They must be optimized to elaborate input data arriving at an expected rate, which is not always constant. While worst-case designs can significantly increase hardware resources, more optimistic solutions fail to sustain execution phases with high throughput, leading to system congestion or even computational errors. We present an architecture to monitor and control dataflow architectures that leverage approximate variants to trade-off accuracy and latency of the computational processes. Our microarchitecture features online prediction based on queuing models to estimate the response time of the system and select the proper variant to meet the target throughput, enabling the creation of dynamically-tunable systems.

Keywords: Dataflow; Markov Queue; Hardware Accelerator

1. Introduction

Modern applications require the elaboration of massive amounts of data, e.g. in real-time video streaming for entertainment or surveillance applications, or network communications [1,2]. To achieve high performance, such applications demand heterogeneous System-on-Chip (SoC) architectures with specialized hardware components. Thanks to customization, these architectures can significantly minimize the cost, while hardware parallelism can optimize the execution time [3].

Due to their distributed nature, modern applications may need to support a variable behavior, where input data are not always available at the same speed [4]. In such cases, designers must guarantee not only a high-quality of the result (e.g., a nice video experience) but also a continuity of the service (e.g., continuous streaming in surveillance). Latency-insensitive protocols can be used to ensure correct execution in case of changes in the surrounding behavior, for example by stalling the execution of the component when the data are not available [5]. However, hardware accelerators have limited flexibility. Their entire behavior must be defined and implemented at design time. After that, they cannot implement a new functionality. Also, the execution time is fixed and depends on the microarchitecture. In case of variable behaviors, one can design the components considering the fastest speed that can support all behaviors (“worst-case” approach) but the resulting component would be underutilized in most of the cases or can even create congestion on the next components since it produces the results too fast. Targeting an average speed, instead, would lead to congestion on the inputs when the component cannot keep pace with the input data. These situations have been exploited to reduce the power consumption with dynamic frequency and voltage scaling (DVFS) [6] and they can be used also for implementing adaptive behaviors.

In software, designers can achieve adaptivity by *approximating* the execution of some phases, provided that the application designers can accept a minimal degradation of the outputs [7]. When multiple approximate alternatives are available for the same code, the system can select dynamically which version to be executed. Such systems are called *multi-variant* [8]. When applied to hardware, approximation can either save resources (i.e.,



Citation: Tibaldi, M.; Palermo, G.; Pilato, C. Dynamically-Tunable Dataflow Architectures based on Markov Queuing Models. *Electronics* **2021**, *1*, 0. <https://doi.org/>

Received: December 31, 2021

Accepted: February 9, 2022

Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

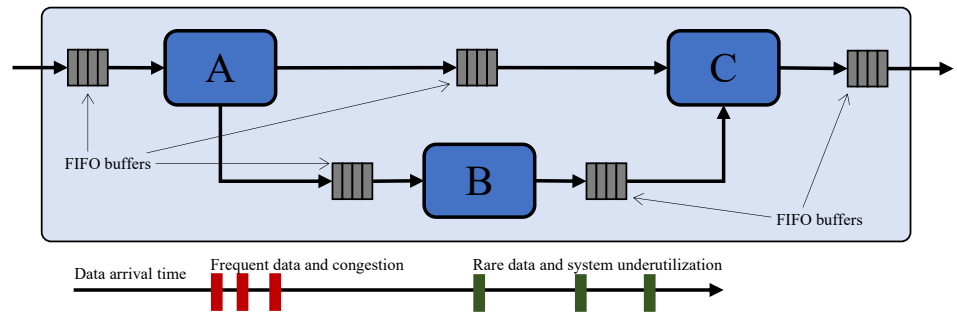


Figure 1. Dataflow architecture with variable input rate.

less logic is used to perform the same computation) or improve the performance (i.e., some computation can be executed faster, improving the hardware microarchitecture or performing the operations in a different way). While many software approximation techniques can be easily applied to hardware accelerators (e.g., variable-to-constant optimizations [9]), multi-variant hardware systems are more difficult to be designed since they need to (1) design efficient hardware modules able to support all variants, and (2) detect the proper variant efficiently and correctly based on the given workload.

In this paper, we focus on the second aspect of the problem, assuming that the designer applies existing approximation techniques to generate *multi-variant hardware components*. With our control approach, we enable the creation of *dynamically-tunable* dataflow architectures by managing multi-variant accelerators that can dynamically adapt their execution speed to the surrounding conditions. This system allows modulating the hardware to use the approximate versions of a given functionality only when strictly necessary. We start from components that implement multiple variants trading off accuracy and latency. Such variants (called also *configurations*) can be generated with different approximation solutions and merged to reduce area overhead. We extend the multi-variant hardware module with a microarchitecture to automatically select the proper configurations based on the system workload. Such microarchitecture monitors the input data, estimates their arrival rate based on queueing models, and accordingly adjusts the speed of the component. Our main contributions are:

- **A microarchitecture for online predictions of system workload based on queueing models (Section 4);**
- **A framework for the creation of dynamically-tunable dataflow architectures that integrate a hardware implementation of such prediction model (Section 5);**
- **An evaluation of the proposed method in different workload conditions (Section 6).**

Our systems can efficiently reach a target throughput with less error than using preset configurations using minimal additional hardware resources.

The remainder of this paper is structured as follows. Section 2 provides a simple example that motivates our effort, while Section 3 briefly describes the related works on the paper topic. Sections 4 and 5 introduces the proposed method describing the different phases involved at run-time and for the module hardware module generation. Section 6 reports the experimental results obtained adopting the proposed method with respect to state of the art techniques. Finally, Section 7 concludes the paper.

2. Motivating Example

Dataflow architectures are widely used to implement hardware systems that can elaborate a set of incoming data to produce the corresponding results. They are based on a set of concurrent hardware modules that communicate through First-In-First-Out (FIFO) buffers with a producer-consumer paradigm. An example is shown in Figure 1. Such buffers implement a latency-insensitive protocol [5] that guarantees correct computation when the *producer* is not able to provide enough data (leaving the buffers empty) or when the *consumer* is not able to consume enough data (leading to data accumulation in the queues).

Both these cases can lead to system congestion or poor performance. A traditional solution is to design the accelerator by considering the worst-case scenario, aiming at supporting the fastest input rate. In many cases, it is impossible to optimize the accelerator in this way and the designer needs to use approximated implementations. While approximated solutions are fast and can avoid system congestion for the input buffers, they introduce errors in the output results. Also, since the execution becomes faster than before, the congestion can move to the output buffers. We thus need a smarter way to create dynamically tunable accelerators, i.e. architectures that can dynamically change the execution speed (and corresponding error) based on the current workload conditions.

Example: Consider a moving-average filter as a case study. The size of the sampling window can be dynamically adjusted to read more or fewer values, leading to different execution times and errors in the computation of the average. Our goal is indeed to understand how to adjust the window size to achieve a given throughput while minimizing the approximation error. In this case, using the fastest solution for the entire computation leads to achieve the given throughput but the error is around 90%. An alternative approach uses a threshold control system that determines the best configuration for the accelerator based on the number of elements in the buffers. For example, we can use a system where:

$$W_1 = \text{size}_{buffer} / N_{conf}$$

$$W_i = W_1 * (i)$$

where size_{buffer} is the size of the input buffer, N_{conf} is the number of available configurations, W_i is the maximum number of elements that are allowed in the buffer for configuration i (threshold), and i ranges between 1 and N_{conf} . When the buffer number exceeds a threshold, the accelerator is moved to the next (and faster) configuration. This system reduces the error but it does not guarantee that the constraint on the response time is respected. Also, an accelerator can continuously change its configuration when the number of buffer elements fluctuates around one of the thresholds (hysteresis loops). ■

In this paper, we aim at modeling the problem as a Markov Decision Process to correctly set the controller's thresholds while minimizing the approximation error. In queueing theory, an M/D/1 (Markov/Deterministic/1) queue [10] represents a single server queueing process in which the jobs arrive with a Poisson distribution and the overall service time is deterministic. The jobs are served in their order of arrival (like in FIFOs) and the successive job forms a m -state Markov chain $\{0,1,2,3,\dots\}$ where the value corresponds to the number of entities in the system (the configurations in our case). So, arrivals move the process from position i in the chain to position $i + 1$. Queues based on Markov processes may occur in practice when a service adjustment is required (like in case of inputs arriving with variable rate). If we count the service time of a job and its time in the system, the different service times correspond to transitions in the Markov chain (i.e., our *configuration changes*).

3. Related Work

Approximate systems are widely used to reduce the area, power consumption, or latency of a circuit, when the given application can tolerate a certain computational error. Approximate systems are created both at hardware and software levels [7,16]. Hardware approximation can achieve larger benefits, like the generation of smaller circuits. On the contrary, software approximation is more flexible and can be tuned more easily based on application requirements.

Software-level approximation trades off accuracy and performance [17–19]. Memoization speeds up computation by storing results of expensive function calls with the same inputs [20]. Skipping some iterations of a loop (*loop perforation* [17]) or even entire tasks (*task skipping* [21]) can significantly reduce execution time. Software approximation enables the creation of multiple variants (e.g., alternative codes) that can be dynamically selected based on the workload conditions and the application requirements (*multi-versioning* [22]). This

Table 1. Overview of the main characteristics for the related works.

Systems	Description	Advantages	Disadvantages
Gao <i>et al.</i> [11], 2017	data-precision manipulation	A run time system. Approximate the minimum necessary. Great energy savings.	Controller at software level.
Vayerka <i>et al.</i> [12], 2016	Genetic programming to create a library of approximate components	Automatic generation of optimized hardware libraries. Work at HLS level	No run time support. Static approximation. Unfeasible on entire circuit
Lee <i>et al.</i> [13], 2017	Rescheduling operations to reduce circuit latency	The transformations are at the HLS level. Much faster than other rescheduling methodologies. Higher energy savings.	No run time support. Near-optimal result. Not applicable directly in hardware
Nepal <i>et al.</i> [14], 2014	Greedy approach to generate approximate implementations	Automatically discovers approximate design. Area and power savings. Applicable to generic circuits. Is transparent to the design flow	No run time support. Static approximation
Li <i>et al.</i> [15], 2015	Solution for precision optimization at HLS level	Simple. Provide an ILP formulation for precision optimization	Not applicable to cases which are both data-intensive and control-intensive. No run time support. Unfeasible on complex systems
Mantovani <i>et al.</i> [6], 2016	A controller for exploiting dynamic voltage and frequency	Enables pre-silicon tuning and design exploration. Implemented in hardware. Higher power savings	Create unnecessary configuration changes due to thresholds. Focus on power optimization rather performance.

technique is hard to be directly applied in hardware since it requires additional resources for each variant.

Customizing data precision is a popular approximation to create smaller components. For example, Gao *et al.* [11] determine the effects of data-precision manipulation on outputs. Vayerka *et al.* [12] use genetic programming to create a library of approximate components (e.g., adders and multipliers) to be used in HLS. However, approximating an entire circuit with this method is unfeasible due to its exponential complexity. Lee *et al.* [13] leverage a HLS-based method to reduce the circuit latency by eliminating or rescheduling operations (similar to task skipping). Nepal *et al.* [14] use a greedy approach on the hardware behavioral specification to generate a Pareto-optimal set of alternative approximate implementations. Li *et al.* [15] present a comprehensive solution for precision optimization, scheduling, and resource assignment during HLS. Any approximation method requires to estimate the error that can be obtained with statistical estimations [23] or with RTL simulations. All these approaches can be used to create the approximate configurations. However, since such implementations are often structurally similar, datapath merging methods enable the creation of multi-variant hardware components [24,25].

Finally, dynamically changing the “speed” of hardware components to reduce congestion requires online monitors and controllers. For example, Mantovani *et al.* [6] use a local controller for exploiting dynamic voltage and frequency scaling (DVFS) in NoC-based

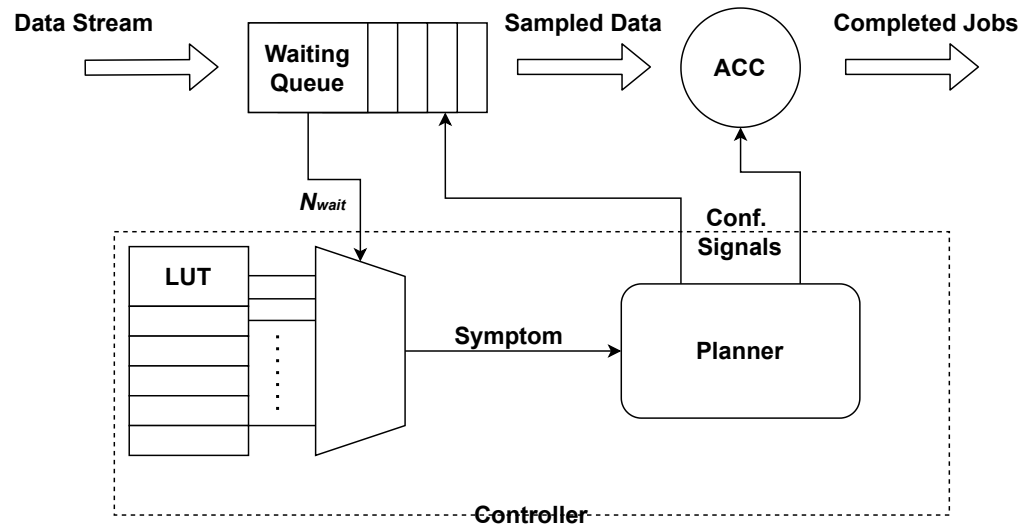


Figure 2. Proposed hardware architecture of our dynamically-tunable accelerators.

architectures. We use a similar approach to analyze the “congestion” on the communication buffers and determine when the component can change implementation, and thus approximation level. However, as discussed in Section 2, this threshold-based approach is inefficient because it can create unnecessary configuration changes. Table 1 summarizes the advantages and disadvantages of the presented works. We aim at implementing a smarter approach based on queue models, which have been successfully used for runtime resource allocation in multicores [10]. This paper describes how to create the corresponding hardware microarchitecture that efficiently changes the accelerator’s configuration.

4. Hardware Architecture and Model for Online Predictions

We assume to have a hardware dataflow accelerator like the one in Figure 1. The accelerator has input and output FIFO buffers to decouple computation and communication with latency-insensitive protocols [5]. We also assume that each dataflow accelerator supports **dynamic tuning**, i.e., it has a set of input parameters that can be used to select an *approximated configuration*. Each accelerator has K configurations ($k = 1, 2, 3, \dots, K$). Each configuration k is characterized by an execution time τ_k and an execution error ϵ_k . The entire set of configurations can be obtained by combining approximation techniques and design space exploration as discussed in Section 3. Execution time and error are known at design time and can be obtained analytically or with RTL simulation.

4.1. Key Idea and Architecture

We associate a *controller* to the dataflow accelerator to be dynamically tuned. The controller selects the configuration to be executed and provides the corresponding identifier to the component. In case of dataflow accelerators composed of multiple sub-components that can be individually tuned, the configuration is characterized by a set of parameters to be provided at the same time to each sub-component. The designer needs to identify in advance configurations (i.e., a combination of parameters) that are inefficient (e.g., due to large errors) and exclude them from the list. However, this process is part of the design exploration process that selects the Pareto set configurations. We add logic to delay the selection until the start of the component’s iteration (i.e., when it reads data from the input FIFOs) to avoid inconsistencies during the computation. The approach is valid for both ASIC and FPGA implementations. In case of FPGA implementations, this approach is much faster than partial dynamic reconfiguration since the hardware is deployed on the configuration logic only once and it does not require any further changes during execution.

The controller includes the logic to detect congestion and to “speed-up” the computation, and it is parametric with respect to the number of configuration in the Pareto set

of the supervised component. To monitor the execution, each controller is connected to the input FIFO buffers with *full*, *almost-full*, and *empty* signals. The status of the input queues is monitored at regular interval (called *observation time*). When one of input FIFOs is *almost-full*, the controller selects a *faster* configuration for the component to facilitate emptying the queue. Instead, when the input queue becomes *empty*, the controller can select a more accurate but slower configuration to improve the accuracy. Figure 2 shows an example of the resulting hardware architecture.

The controller does not require to have specific information on the configurations because it assumes they are ordered from fastest to slowest (e.g., configuration i is the fastest with no approximation error). This approach is similar to the use of fine-grained DVFS with integrated voltage regulators [6]. The selection of the next configuration requires to avoid hysteresis loops around the buffer thresholds (see Section 2). For this reason, our controller is based on queue models.

4.2. Queue modeling for predicting the response time

The proposed model aims at providing a suitable run-time policy for configuring the accelerator to minimize the approximation error while meeting a specified constraint on the response time. In this work, we model average response time R using the theory of *queuing networks*. We model the accelerator as a single resource service station (see Figure 3). The accelerator is the resource that serves the transaction, while its queue is modeled as the waiting line of the station. The service time of the station is modeled as the execution time t_k in each one of the different configurations k , while the **expected service rate** is calculated as $\mu_k = 1/\tau_k$. Given the balance equation, the job arrival rate λ of an application represents the throughput required by the user. It is measured in Job/s and it depends on the activity to be monitored at run-time.




Figure 3. Queuing system for single accelerator.

To enable run-time management as described previously in the paper, the controller have to maintain and dynamically evaluate the expected average response time. If we consider that the job arrival times can be modelled as a continuous-time Markov process, and in particular job inter-arrival times are exponentially distributed with mean $\lambda = 1$, we can produce a prediction model for R by modeling the problem as an M/D/1 Markov process. I.e. arrivals are determined by a Poisson process (M), job service times are deterministic (D) and single resource service station (1).

In the M/D/1 model, the expected number of jobs in the system (either waiting in queue or being served) in the steady state is given by:

$$\rho(\mu_k, \lambda) = \lambda / \mu_k \quad (1)$$

$$N(\mu_k, \lambda) = \rho(\mu_k, \lambda) + [\rho(\mu_k, \lambda)]^2 / 2[1 - \rho(\mu_k, \lambda)] \quad (2)$$

where ρ is the system utilization, i.e., the fraction of time in which the system is busy. Given Equations (1) and (2), we can build a prediction model for R by using *Little's law*:

$$R(\mu_k, \lambda) = N(\mu_k, \lambda) / \lambda \quad (3)$$

where R only depends on the arrival rate λ and the estimated service rate μ_k .

We use this model to find the maximum arrival rate that guarantees a response time under a user-defined bound in each configuration. So, from configuration k we derive the

estimated $\lambda_{max,k}$ that matches to a specific amount of jobs (elements) W_k in the waiting queue:

$$W_k = \lambda_{max,k} * obs_{time} \quad (4)$$

Where $\lambda_{max,k}$ is predicted from the inverse of the equation 3 and obs_{time} is the *observation time* of the queue, i.e. the frequency in which the controller samples the status of the queue. W_k represents the maximum number of elements that can be stored in the queue for which the system is able to achieve the response time R by using configuration k . Once the values are computed for each configuration k , we can generate a lightweight logic that changes the accelerator configuration k to $k + 1$ when the number of elements detected in the input buffer exceeds the corresponding bound W_k (and vice-versa).

5. Generation Methodology for the Online Controller

Our framework requires the user to specify the characteristics of the K individual configurations along with the execution time τ_k for each of them. It also needs the sizes of the input buffers and the required observation time obs_{time} of the controller. Finally, it requires the description of the accelerator configuration ports and the control signals to correctly apply the decisions.

From these data, the framework computes the M/D/1 model, i.e. the values W_k for each configuration. The solution for configuration k is admissible if the corresponding value W_k is smaller than the size of the input buffers. Otherwise, it means that the input buffers cannot store enough values to achieve the response time. If all models can be computed (i.e., all values W_k , our generator automatically produces the Verilog description of the corresponding controller as shown in Figure 2. In particular, the input buffers are extended with logic to count the number of elements currently in the queue (N_{wait}). The resulting structure is called *smart waiting queue*.

At run time, the controller samples the waiting queue at regular intervals defined by the observation time. It reads the amount of data stored inside the queue and automatically determines the next configuration for the accelerator, given the current configuration and the number of elements in the queue. The values W_k are stored in a look-up table. Assuming the controller is in configuration k , we have three possible cases (represented by the signal *symptom* in Figure 2:

1. $N_{wait} < W_{i-1}$, i.e. the accelerator is emptying the queue and it can slow down (going to configuration $i - 1$ with more precision);
2. $N_{wait} > W_i$, i.e. the accelerator is not able to consume the elements in the queue, which are accumulating, and must accelerate (going to configuration $i + 1$);
3. $W_i < N_{wait} < W_{i-1}$, i.e. the accelerator can stay in the current configuration.

The *planner* can apply the decision to the control signals of the accelerator right before the next iteration starts (i.e. the next value is read from the input buffers). Since every iteration is independent of the previous ones, this mechanism ensures the correct execution of the acceleration when changing the configurations.

We implemented the generator of the controller in PyVerilog [26]. It receives a `json` configuration file as input with all necessary information about the target accelerator, and extends the original design with the corresponding controller, directly generated in Verilog.

6. Experimental Results

To evaluate our solution, we applied this method to the five accelerators described in Table 2. We use two signal processing benchmarks (DSP) and two image processing benchmarks (IP) [9]. The fifth benchmark is a combination of two other accelerators. We used this example to show how the methodology can be applied to a complex accelerator composed of subcomponents. The same table describes also the input stimuli and the quality metric used to evaluate the accuracy of the output results. In this work, we use

Table 2. Benchmark Characteristics

Circuit	Description	Application	Input Stimuli	Quality Measure
AVE8	Moving Avg Calculator	Signal processing	1,000 random integers	MAPE
FIR	Finite Impulse Response Filter	Signal processing	1,000 random integers	MAPE
SOBEL	Sobel Edge Detector	Image processing	1920x1080 image	PSNR
GS	Grey Scale Filter	Image processing	256x256 image	PSNR
GS+SOBEL	Two accelerators in series	Image processing	256x256 image	PSNR

Mean Average Percentage Error (MAPE) for the DSP applications and Peak Signal-to-Noise Ratio (PSNR) for the image processing ones. We compute MAPE as

$$MAPE = 1/N \sum_{i=1}^N (GO_i - AO_i) / GO_i \quad (5)$$

where GO_i is the golden output (i.e., the correct/original one), AO_i is the output of the approximate description, and N is the total number of inputs. We compute PSNR as

$$PSNR = 10 * \log(255^2 / MSE) \quad (6)$$

where MSE (Mean Square Error) is defined as


$$MSE = \frac{1}{N} \sum_{i=1}^N (GO_i - AO_i)^2 \quad (7)$$

We created six different configurations for each benchmark, where CONF0 is the slowest (with no approximation) and CONF5 is always much faster than the target response time.

To test the system under dynamic conditions, we created three workload situations, with *highly-congested*, *congested*, and *uncongested* traffic. In our experiments, we set the response time R to 1.5us, which has been kept constant across all the benchmarks, and an observation time of 1.28us, which is the minimum time to fill half of the input buffers. The given response time is set to a value that can be never achieved with CONF0 (i.e., the precise configuration), demanding to introduce approximations to achieve it.


In our experiments, we evaluated FPGA implementations, while ASIC ones are completely equivalent. For each design we used Xilinx Isim 2018.3 to evaluate the performance and the approximation error, and Xilinx Vivado 2018.3 to target a Xilinx VC707 board (equipped with a Virtex-7 XC7VX485T FPGA) with a target frequency of 100 MHz to evaluate the resource overhead introduced by our controller.

Figures 4 to 8 show the evolution of the response time of the different accelerators over time, along with the quality metrics in case of preset configurations for the accelerators (from CONF0 to CONF5) or when they use our method (ADA). For clarity reason, we show only the extreme cases: highly congested and uncongested situations. In congested systems, the response time grows constantly overtime to a maximum value. This value is an intrinsic characteristic of the system, and it depends on factors like the size of the input queue(s), the source of the arrival packets, and the processing speed. Conversely, in uncongested systems, the response time has a trend with peaks. These two behaviors depend exclusively on the inbound traffic. In the former, there is a continuous flow of packets into the system that stops solely when the queue saturates. In the latter, the traffic is sporadic and it always



(a) AVE8: Response time - highly congested


(b) AVE8: MAPE - highly congested



(c) textscave8: Response time - uncongested


(d) AVE8: MAPE - uncongested

Figure 4. AVE8: Response time and MAPE (low is better) trends. R is the target response time.



(a) FIR: Response time - highly congested

(b) FIR: MAPE - highly congested



(c) FIR: Response time - uncongested

(d) FIR: MAPE - uncongested

Figure 5. FIR: Response time and MAPE (low is better) trends. R is the target response time.

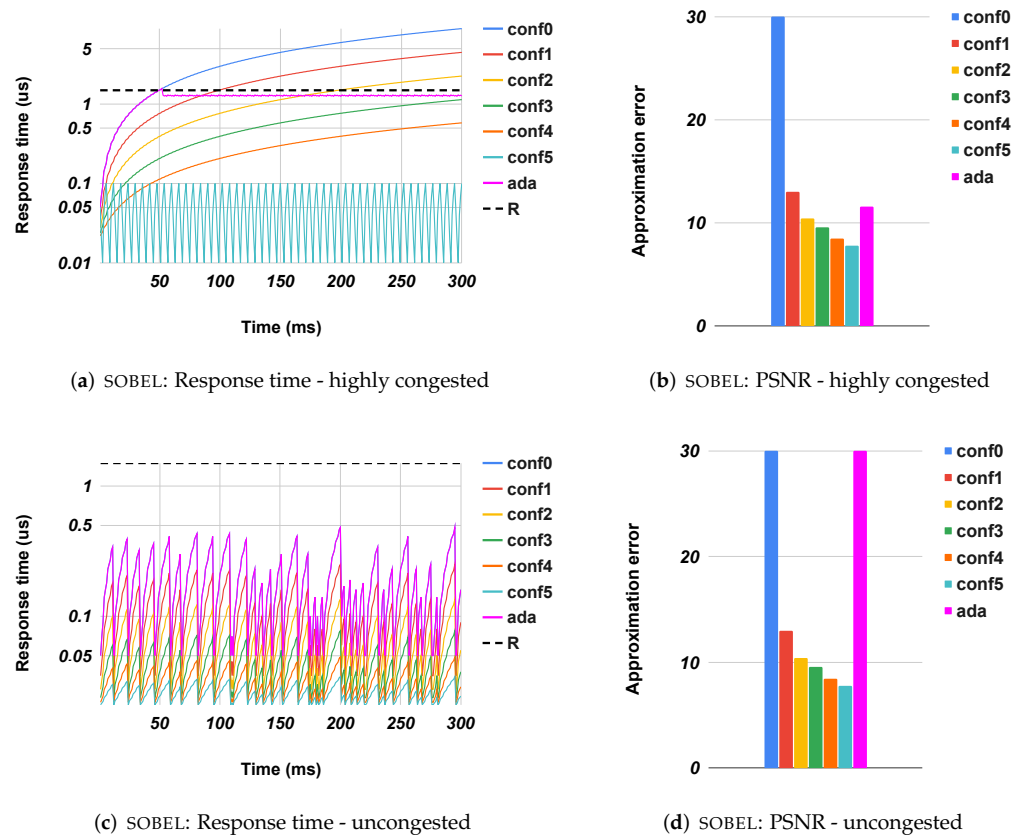


Figure 6. SOBEL: Response time and PSNR (high is better) trends. R is the target response time.

allows the system to empty the queues. The results show our controllers correctly configure the accelerators to satisfy the response time with a quality metric better or an error less than the ones obtained with fixed configurations. In all cases, the response time is close to the expected one (exploiting the speed of approximate configurations (like CONF4 and CONF5), while limiting the error as configurations CONF0 and CONF1).

Tables 3 and 4 show the corresponding metrics in all three scenarios. The error is always less than the one obtained from preset configurations except for GS in uncongested cases where the controller unnecessary speeds up the execution of the accelerator. This situation happens because, even if the traffic is sporadic, some packet flows re-fill the input queue. In the graph of Figure 7(c), we can see where some CONF0 peaks exceed the target response time. The controller interprets these events by preparing the accelerator for a continuous arrival of packets, but this does not happen. So for a time window that corresponds to the observation time, the system will go faster than needed, leading to a slight degradation in the quality of the results. In general, the system adjusts itself based on the workload conditions, trading off accuracy and speed as needed. For example, in the highly-congested scenario (Figures 6(a) and 6(b)), the overall SOBEL PSNR has a value close to the one obtained with CONF1 while only CONF4 meets the target response time but with much worse PSNR. Conversely, in the uncongested scenario Figure 6(c), the system adjusts itself to the most precise configuration (CONF0) without any metric degradation.

In the GS+SOBEL Figure 8, we tested a system with two accelerators in series: gs followed by sobel. This experiment aims at showing how to apply our controller to a multi-module system composed of modules that can be approximated independently. In this benchmark, we used only four configurations because some of them had W_i larger than the size of the buffers, making them unfeasible. Also in this case, the controller allows the accelerator to meet the given response time while improving the final error. Note that

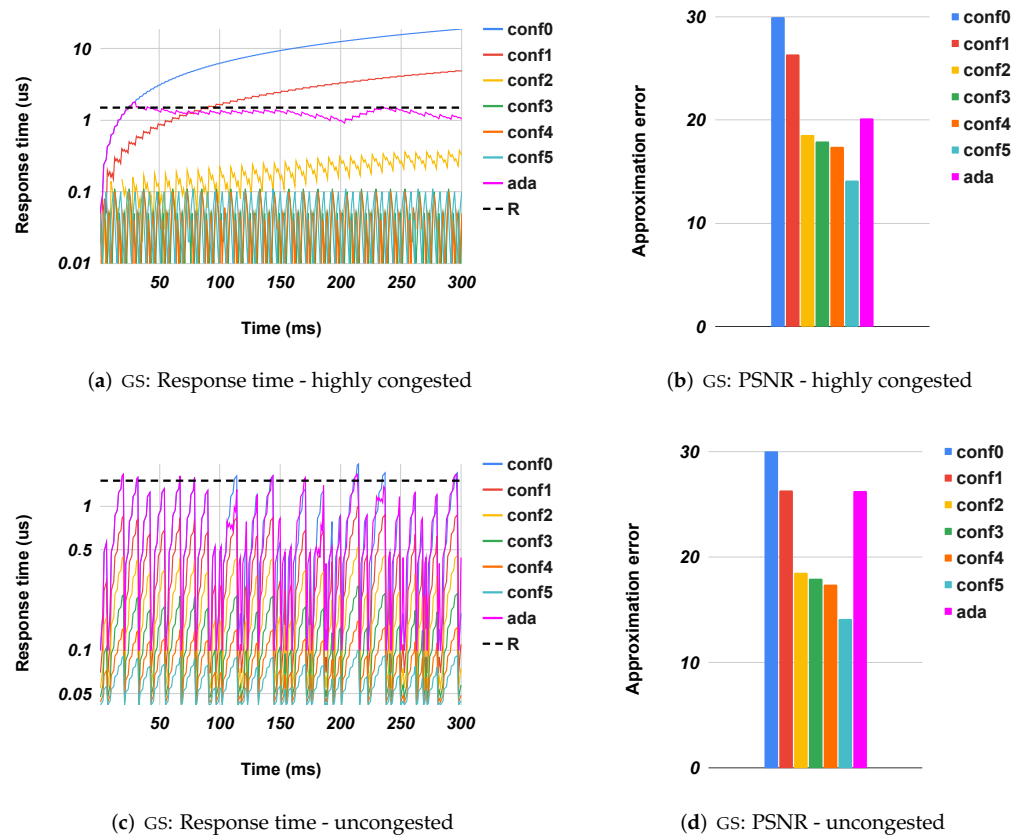


Figure 7. GS: Response time and PSNR (high is better) trends. R is the target response time.

Table 3. DSP Results

Test	Highly Congested		Congested		Uncongested	
	R	MAPE	R	MAPE	R	MAPE
AVE8	+30.9%	-27.6%	+148.1%	-92.3%	-1.2%	-40.7%
FIR	+16.9%	-16.3%	+6.5%	-21.3%	-36.5%	-33.8%

Table 4. IP Results

Test	Highly Congested		Congested		Uncongested	
	R	PSNR	R	PSNR	R	PSNR
SOBEL	+113.2%	-21.4%	-1.4%	-3.5%	-	-
GS	+10.9%	-12.5%	+9.8%	-30.2%	-3.4%	+0.4%
GS+SOBEL	+80.3%	-59.8%	+53.2%	-61.1%	+12.8%	-62.8%

the overall PSNR degradation is large due to an intrinsic characteristic of the benchmark rather than a problem in our methodology.

From the synthesis viewpoint, Table 5 shows that the enhanced accelerator takes a negligible overhead and, as expected, its impact decreases as the complexity of the target module increases. This overhead is significantly less than the one obtained in previous approaches because our method is based on a simple look-up table rather than complete state machines like in [6].

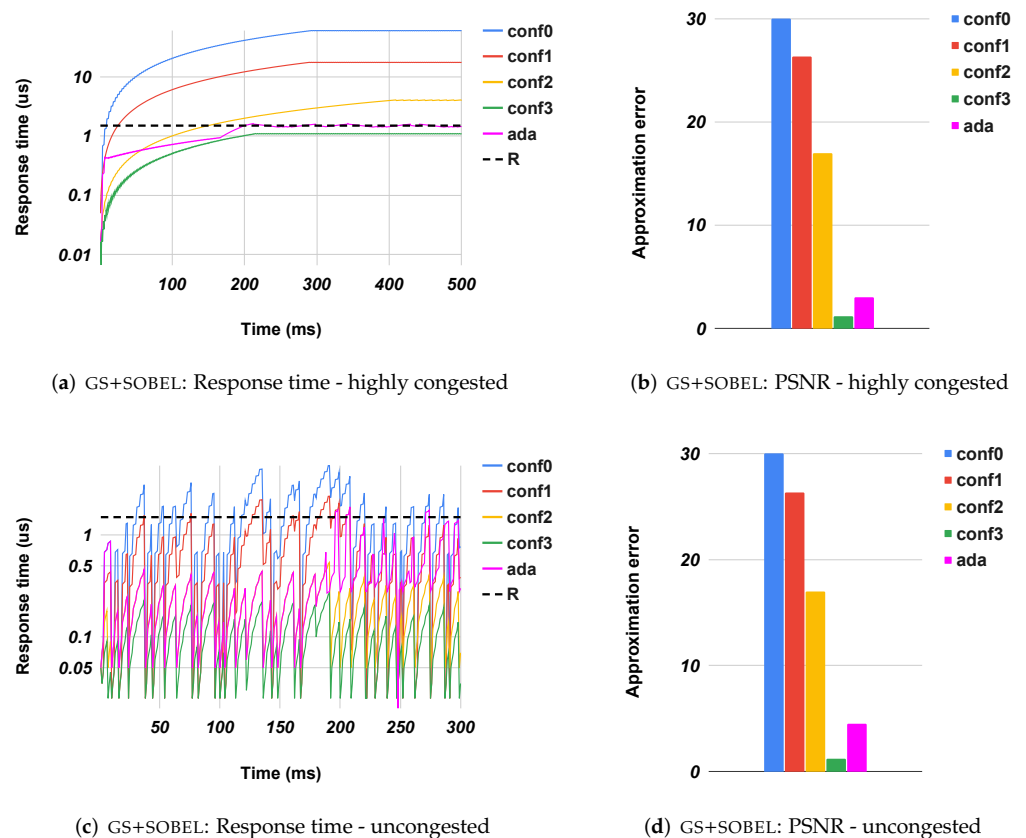


Figure 8. GS+SOBEL: Response time and PSNR (high is better) trends. R is the target response time.
Table 5. Controller and Smart Logic Area Overhead

Circuit	Slice LUTs	Slice Registers
AVE8	+2.2%	+5.5%
FIR	+2.1%	+6.4%
SOBEL	+3.4%	+8.7%
GS	+0.3%	+0.2%
GS+SOBEL	+0.3%	+0.2%

7. Conclusions and Future Work

We presented a solution to create dataflow accelerators that can dynamically trade-off execution latency and quality of results to meet a given response time in case of inputs arriving at an unpredictable rate. We modeled the system as a Markov queueing model to predict the response time and dynamically adjust the speed of the accelerator. Our solution can meet the response time with a final error that is lower than the one obtained by always using the fastest implementation. Our solution has also minimal hardware overhead. In the future, we will work on methods to scale our solution to accelerators composed of multiple modules, considering the case of queues between individual modules, both on the design of the single components and the controller(s), and on upgrading the controller to manage multiple approximation techniques. These extensions contain many challenges, from selecting approximations and studying their impact on the whole system to coordinating multiple controllers to ensure correct execution.

Funding: This work has been partially funded by the Horizon 2020 EU Research & Innovation Programme under grant agreement No. 957269 (EVEREST project).

Author Contributions: “Conceptualization, G.P. and C.P.; Methodology, M.T., G.P., C.P.; Software, M.T.; Writing—Original Draft Preparation, M.T. and C.P.; writing—review and editing, G.P. and C.P.; funding acquisition, G.P. and C.P.; All authors have read and agreed to the published version of the manuscript.”

Conflicts of Interest: “The authors declare no conflict of interest.”

References

1. Wu, S.; Gutgutia, S.; Alioto, M.; Baas, B. Display Stream Compression Encoder Architectures for Real-time 4K and 8K Video Encoding. Proceedings of the Asilomar Conference on Signals, Systems, and Computers (ACSSC), 2018, pp. 251–255. doi:10.1109/ACSSC.2018.8645369.
2. Pilato, C.; Bohm, S.; Brocheton, F.; Castrillon, J.; Cevasco, R.; Cima, V.; Cmar, R.; Diamantopoulos, D.; Ferrandi, F.; Martinovic, J.; et al. EVEREST: A design environment for extreme-scale big data analytics on heterogeneous platforms. Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE), 2021, 2021, pp. 1320–1325. doi:10.23919/DATE51398.2021.9473940.
3. Mantovani, P.; Giri, D.; Di Guglielmo, G.; Piccolboni, L.; Zuckerman, J.; Cota, E.G.; Petracca, M.; Pilato, C.; Carloni, L.P. Proceedings of the ACM/IEEE International Conference on Computer-Aided Design (ICCAD), 2020. doi:10.1145/3400302.3415753.
4. Babcock, B.; Babu, S.; Datar, M.; Motwani, R.; Widom, J. Models and Issues in Data Stream Systems. Proceedings of the ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS), 2002, p. 1–16. doi:10.1145/543613.543615.
5. Carloni, L.P. From Latency-Insensitive Design to Communication-Based System-Level Design. *Proceedings of the IEEE* **2015**, *103*.
6. Mantovani, P.; Cota, E.G.; Tien, K.; Pilato, C.; Di Guglielmo, G.; Shepard, K.; Carlon, L.P. An FPGA-based infrastructure for fine-grained DVFS analysis in high-performance embedded systems. Proceedings of the ACM/EDAC/IEEE Design Automation Conference (DAC), 2016. doi:10.1145/2897937.2897984.
7. Mittal, S. A Survey of Techniques for Approximate Computing. *ACM Computing Surveys* **2016**, *48*, 1–33.
8. Cherubin, S.; Agosta, G. libVersioningCompiler: An easy-to-use library for dynamic generation and invocation of multiple code versions. *SoftwareX* **2018**, *7*, 95–100.
9. Chowdhury, P.; Carrion Schafer, B. Unlocking Approximations through Selective Source Code Transformations. Proceedings of the ACM Great Lakes Symposium on VLSI (GLSVLSI), 2021, p. 359–364.
10. Mariani, G.; Palermo, G.; Zaccaria, V.; Silvano, C. ARTE: An Application-specific Run-Time management framework for multi-cores based on queuing models. *Parallel Computing* **2013**, *39*.
11. Gao, M.; Qu, G. Energy efficient runtime approximate computing on data flow graphs. Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD), 2017, pp. 444–449. doi:10.1109/ICCAD.2017.8203811.
12. Vaverka, F.; Hrbacek, R.; Sekanina, L. Evolving component library for approximate high level synthesis. Proceedings of the IEEE Symposium Series on Computational Intelligence (SSCI), 2016, pp. 1–8. doi:10.1109/SSCI.2016.7850168.
13. Lee, S.; John, L.K.; Gerstlauer, A. High-level synthesis of approximate hardware under joint precision and voltage scaling. Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017, 2017, pp. 187–192. doi:10.23919/DATE.2017.7926980.
14. Nepal, K.; Li, Y.; Bahar, R.I.; Reda, S. ABACUS: A technique for automated behavioral synthesis of approximate computing circuits. Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE), 2014, pp. 1–6. doi:10.7873/DATE.2014.374.
15. Li, C.; Luo, W.; Sapatnekar, S.S.; Hu, J. Joint precision optimization and high level synthesis for approximate computing. Proceedings of the ACM/EDAC/IEEE Design Automation Conference (DAC), 2015. doi:10.1145/2744769.2744863.
16. Sampson, A. Hardware and Software for Approximate Computing. PhD thesis, University of Washington, 2015.
17. Sidiroglou-Douskos, S.; Misailovic, S.; Hoffmann, H.; Rinard, M. Managing Performance vs. Accuracy Trade-Offs with Loop Perforation. Proceedings of the ACM SIGSOFT Symposium and the European Conference on Foundations of Software Engineering (ESEC/FSE), 2011, p. 124–134. doi:10.1145/2025113.2025133.
18. Sampson, A.; Dietl, W.; Fortuna, E.; Gnanapragasam, D.; Ceze, L.; Grossman, D. EnerJ: Approximate Data Types for Safe and General Low-Power Computation. Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), 2011, p. 164–174. doi:10.1145/1993498.1993518.
19. Samadi, M.; Lee, J.; Jamshidi, D.A.; Hormati, A.; Mahlke, S. SAGE: Self-tuning approximation for graphics engines. Proceedings of the Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), 2013, pp. 13–24. doi:10.1145/2540708.2540711.
20. Tziantzioulis, G.; Hardavellas, N.; Campanoni, S. Temporal Approximate Function Memoization. *IEEE Micro* **2018**, *38*, 60–70. doi:10.1109/MM.2018.043191126.
21. Rinard, M. Probabilistic Accuracy Bounds for Fault-Tolerant Computations That Discard Tasks. Proceedings of the Annual International Conference on Supercomputing (ISC), 2006, p. 324–334. doi:10.1145/1183401.1183447.
22. Zhou, M.; Shen, X.; Gao, Y.; Yiu, G. Space-Efficient Multi-Versioning for Input-Adaptive Feedback-Driven Program Optimizations. Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages & Applications (OOPSLA), 2014, p. 763–776. doi:10.1145/2660193.2660229.
23. Su, S.; Wu, Y.; Qian, W. Efficient Batch Statistical Error Estimation for Iterative Multi-Level Approximate Logic Synthesis. Proceedings of the ACM/EDAC/IEEE Design Automation Conference (DAC), 2018. doi:10.1145/3195970.3196038.

24. Souza, C.C.d.; Lima, A.M.; Araujo, G.; Moreano, N.B. The Datapath Merging Problem in Reconfigurable Systems: Complexity, Dual Bounds and Heuristic Evaluation. *Journal of Experimental Algorithmics* **2005**, *10*.
25. Fanni, T.; Sau, C.; Meloni, P.; Raffo, L.; Palumbo, F. Power and Clock Gating Modelling in Coarse Grained Reconfigurable Systems. Proceedings of the ACM International Conference on Computing Frontiers (CF), 2016, p. 384–391. doi:10.1145/2903150.2911713.
26. Takamaeda-Yamazaki, S. Pyverilog: A Python-Based Hardware Design Processing Toolkit for Verilog HDL. Proceedings of the International Symposium on Applied Reconfigurable Computing (ARC), 2015, pp. 451–460.