

Clock-gating of streaming applications for energy efficient implementations on FPGAs

Endri Bezati, Simone Casale-Brunet, *Member, IEEE*
Marco Mattavelli, and Jorn W. Janneck, *Member, IEEE*

Abstract—The paper investigates the reduction of dynamic power for streaming applications yielded by asynchronous dataflow designs by using clock gating techniques. Streaming applications constitute a very broad class of computing algorithms in areas such as signal processing, digital media coding, cryptography, video analytics, network routing and packet processing and many others. The paper introduces a set of techniques that, considering the dynamic streaming behavior of algorithms, can achieve power savings by selectively switching off parts of the circuits when they are temporarily inactive. The techniques being independent from the semantic of the application can be applied to any application and can be integrated into the synthesis stage of a high-level dataflow design flow. Experimental results of at-size applications synthesized on FPGAs platforms demonstrate power reductions achievable with no loss in data throughput.

I. INTRODUCTION

Power dissipation is currently the major limitation of silicon computing devices. Reducing power has also other beneficial effects, it implies less stringent needs for cooling, improved longevity, longer autonomy in the case of battery operated devices and obviously, lower power costs. For all these reasons power also frequently affects the choice of the computing platform right at the outset. For example, Field-Programmable Gate Arrays (FPGAs) imply higher power dissipation per logic unit when compared to equivalent Application-Specific Integrated Circuit (ASIC), but often compare favorably to conventional processors used for the same functional tasks.

For any silicon device, power dissipation can be partitioned into two components: a static and a dynamic component. Static power dissipation, also referred to as quiescent or standby power consumption, is the result of the leakage current of the transistors, also affected by the ambient temperature. By contrast, dynamic power dissipation is caused by transistors being switched and by losses of charges being moved along wires. Power dissipation increases linearly with frequency, due largely to the influence of parasitic capacitances. To counteract this effect, ASIC designers have employed *clock gating* (CG) techniques in the last twenty years [1], [2], [3].

Different strategies for optimizing power consumption on ASICs and FPGAs are discussed in Section II. These papers describe the impact of a chosen technology for a given architecture, but do not describe how to reduce power at

the design abstraction level. As a consequence, adding power controllers at the behavioral description design stage constitutes an additional task that has to be carried-out with care to avoid introducing undesired application behaviors and might reduce the portability of the code (i.e platform is changed during the development process). In addition, it is extremely difficult for HLSs approaches that are based on Imperative Model of Computations (MoCs) [4] to apply power optimization solutions that can be yielded by automatic tools starting from the (imperative) behavioral description. Conversely, dynamic dataflow [5], [6], [7] designs such as for instance the ones expressible using the formal RVC-CAL language possess interesting properties that can be exploited for reducing the power consumption without affecting, by construction, the behavioral characteristics of the application. In RVC-CAL, every actor can concurrently execute processing tasks, executions might be disabled by input blocking reads, and every communications among actors can occur only by means of order preserving lossless queues. As a consequence, an actor may be stopped for a certain period if its processing tasks are idle or its outputs queues (buffers) are full without impacting the overall throughput and semantical behavior of the design. In addition, to higher levels of dynamic behaviors that might be present in a given dataflow design, correspond higher levels of power reduction opportunities. This is not the case for synchronous dataflow designs that always consume and produce a fixed amount of data tokens. Thus, synchronous dataflow design always dissipate a constant amount of power compared to asynchronous dataflow. In this perspective the techniques that transform intrinsically dynamic algorithm into static versions such as the ones that are implemented by static dataflow MoC for deriving analytical guaranteed bounds or other analyzability purposes. In general this transformations are done by introducing dummy tokens guaranteeing constant rates. Thus, in terms of power optimization such approaches are inefficient.

This paper is organized as follows: in Section II previous works on clock-gating are briefly introduced. Section III describes in details the clock-gating strategy and how it is applied on a dataflow design. In Section IV experimental results are presented and conclusions are finally drawn in Section V.

II. RELATED WORK

Globally Asynchronous Locally Synchronous (GALS) based systems consist of several locally synchronous components which communicate with each other asynchronously.

Endri Bezati, Simone Casale Brunet, Marco mattavelli are with the Laboratory SCI-STI-MM of École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland e-mail: (endri.bezati@epfl.ch)

Jorn W. Janneck is with Departement of Computer Science, Lund University, Sweden

Manuscript received April 19, 2014

Works on GALS can be separated into three categories: partitioning, communication devices, and dedicated architectures. Dataflow design modeling, exploration, and optimization for GALS-based designs has been studied previously by several authors. Shen et al. [8] proposed a design and evaluation framework for modeling application-specific GALS-based dataflow architectures for cyclo-static applications, where system performance, e.g. throughput, is taken into account during optimization. Similarly, Wu et al. [9] and Ghavami et al. [10] proposed a method for automatic synthesis of asynchronous digital systems. These two approaches were developed for fine-grained dataflow graphs, where actors are primitives or combinational functions. Related to our work, authors in [11] proposed a multiple clock, domain-design methodology for reducing the power consumption of dataflow programs. Their design objective was to optimize the mapping of an application while still meeting design performance requirements. This optimization was achieved by assigning each clock domain an optimized clock frequency to reduce power consumption.

III. CLOCK-GATING STRATEGY

Current FPGA families support different clock gating strategies and each manufacturer provides its own IP for managing these different approaches. The methodology described here is based on using primitives specific to Xilinx FPGA architectures. However, this methodology can be modified to support other FPGA vendors primitives. In the remainder of this section, it is briefly described how clock gating techniques are implemented on Xilinx FPGAs and how an automatic clock gating strategy within Xronos HLS is realized.

A. Profile guided buffer size

The execution of a dataflow program consists of a sequence of action firings. These firings can be correlated to each other in a graph-based representation using an approach called Execution Trace Graphing (ETG). The graph is an acyclic directed graph where each node represents an action firing, and a directed arc represents either a data or a control dependency between two different action firings. The effectiveness of analyzing a dataflow program using an ETG is demonstrated in [12]. Xronos provides profiling for each firing execution in clock cycles. This is achieved by retrieving the difference of DONE and GO signals for each action firing during RTL simulation [13]. Timing information is added to the ETG for each firing and each dependency arises according to a corresponding time value, thus transforming the ETG into a weighted graph. A close-to-optimal buffer size configuration, in terms of execution throughput and buffer memory utilization, can be obtained through an iterative analysis of the algorithmic critical path evaluated using the weighted ETG. For a detailed description the interested reader can refer to [14].

B. Coarse-grained clock gating strategy

When the output buffer of any actor is full, the clock of this actor should be turned off as the actor is idle. This is because

switching off its clock will not have an impact on design throughput. Even though RVC-CAL dataflow designs are used for the behavioral description, such clock gating strategy is more general and can be applied to systems that represent the execution of a process that communicates with asynchronous FIFO buffers. The queues should be asynchronous for lossless communication when an actor is clock gated and a design has differing input clock domains.

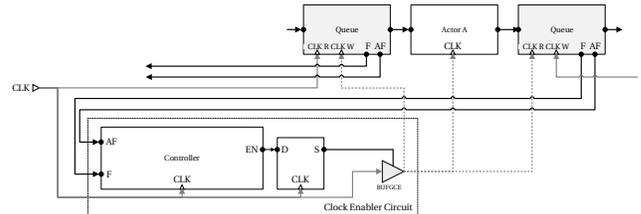


Fig. 1: Clock gating methodology applied for Actor A. The actor A has two outputs one of those have a fanout of two. The Clock enabling circuit takes the *Almost Full* and *Full* signal of each queue and a clock from a clock domain and as a result it is going to activate or deactivate the clock of Actor A.

This strategy consists of adding a *Clock Enabler* circuit for activating the Actors' clock. This circuit contains: a controller for each output port queue of each actor, a combinational logic for the configuration of the output ports, and a clock buffer (which enables the clock). A representation of an actor with a single output port being clock gated is illustrated in Figure 1. As depicted, queues are asynchronous. Queues have two input clocks: one for consuming tokens and one for producing them. Additionally, queues have two output ports: **AF** for almost full, and **F** for full. The actors input clock is connected to the output of the *Clock Enabler* circuit. Finally, the clock buffer *BUFGCE* input clock should be connected with a Flip-Flop for glitch-free clock gating [15].

The Flip-Flop will introduce a one-clock latency when the clock is switched off, but this additional clock cycle will not have an impact on actors that are on the critical path. Those actors are not being clock gated because the TURNUS dimensioning of the FIFO queues is based on critical path analysis. Hence, this approach does not impact overall performance.

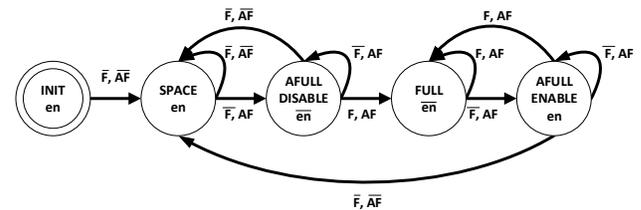


Fig. 2: State machine of the clock enabling controller. The controller has two inputs, **F** for full, **AF** for almost full and one output **en** as the enable signal.

Clock enabling controller: The clock enabling controller is represented in Figure 2. The controller is implemented as a finite state machine having a clock; a reset; input *F*, for

full; input *AF*, for almost full; and output *EN*, for enable. The *AF* input becomes active high when there is only one space left on its FIFO Queue. Its finite state machine (FSM) has 5 states $S = \{INIT, SPACE, AFULL_DISABLE, FULL, AFULL_ENABLE\}$. The controller starts with the *INIT* state and maintains the *EN* output port at active high until *F* and *AF* become active low.

The active high *EN* is maintained during the *SPACE* state. As a queue becomes full, the state changes to *AFULL_DISABLE*. In this state, the *EN* output passes to an active low. A conservative approach is taken in this state as the *BUFGCE* disables the output clock on the high-to-low edge. The clock enables entering the *BUFGCE* should be synchronized to the input clock. Once the queue becomes full, the controller maintains the *EN* at active low. When a token is consumed from the queue, the controller passes to the *AFULL_ENABLE* state, and it activates the clock. Then, depending on whether the buffer becomes full or almost full, the state changes to either the *FULL* or the *SPACE* state.

Strategy: The user can choose a mapping configuration that indicates which actor should be clock gated. To do so, an attribute is given to each actor. If an actor has been selected for clock gating, all of its output FIFO queues, A and AF , are connected to a clock enabler controller. Output queues can be connected through a fanout or directly to a queue. In the first case, the controller results are connected to an AND logic port. This is a safe approach in the case that one of the queues in the fanout is full. In this case, the fanout should command the actor not to produce a token. For the latter case, if an actor's output is connected directly to a queue without a fanout, the result should be connected to an OR logic port as the next actor may need to consume a certain number of tokens to output a token. This may lead the system to lock due to the unavailability of data. In the third case, if there is a combination of outputs with or without a fanout, then an n-input OR logic port is inserted. Figure 3 depicts these configurations.

A pseudo-template of the clock Enabler circuit is given in Template 1. This template generates a Verilog file that takes into account the different cases described previously. These situations are detected and generated automatically as described in the "Always Clause". A flip-flop (created by the always clause) is connected between the *BUFGCE* and the final OR or AND port. Thus, clock glitches are eliminated and the clock enabling is runt free. The last output of the clock gating is a new clock that is connected to the actors, its fanouts, and its queues' write and read clocks (CLK W and CLK R, respectively) as visualized in Figure 1.

IV. EXPERIMENTAL RESULTS

In this section, the power reduction gain of the aforementioned methodology is evaluated by applying it to a video decoder design. In [16], the reader can find a variety of RVC-CAL applications for dataflow programs. One of these applications is the Intra MPEG 4 simple profile decoder. Due to restrictions on the number of clock buffers in Xilinx FPGAs, the design selected was refactored to result in 32 actors.

Template 1: Clock Enabler circuit module creation

```

module clock_enabler
Input : actor
Input : enable
Input : clk_in
Input : reset
Input :  $\forall P_{almost\_full}^{out}$ 
Input :  $\forall P_{full}^{out}$ 
Output: clk_out
for p in  $\forall P^{out}$  do
  | wire [":sizeof(p.fanout)":0] "nameof(p)"_enable;
reg clock_enable;
wire buf_enable;
for p in  $\forall P^{out}$  do
  | for idx in sizeof(p.fanout) do
    | controller c_"nameof(p)"_idx"(
      | .almost_full("nameof(p)"_almost_full["idx"])
      | .full("port.name"_full["idx"]),
      | .enable("port.name"_enable["idx"]),
      | .clk(clk),
      | .reset(reset));
always @(posedge clk) being
  | clock_enable <= for p  $\forall P^{out}$  SEPARATOR "—" do
    | if sizeof(p.fanout) > 1 then
      | for idx in sizeof(p.fanout) SEPARATOR "&" do
        | | nameof(p)_enable["idx"]
    | else
      | | nameof(p)_enable
assign buf_enable = en ? clock_enable : 1;
BUFGCE clock_enabling (.I(clk), .CE(buf_enable), .O(clk_out));
endmodule

```

Test Design: The Intra MPEG-4 SP description contains 32 actors and it is 4:2:0 decoder which is separated into 8 processing blocks: four components for luminance (Y) and two each for chrominance (U and V)). The parser block includes the syntactical bitstream parser and the variable length decoding process which the Tex Y, U, and V blocks (for texture) implements. The residual decoding (AC-DC prediction, inverse scanning, inverse quantization, and IDCT) and the MOT Y, U and V realize the motion compensation stage (framebuffer, interpolation, and residual error addition). Due to the nature of the experiments, the motion compensation blocks contain only the residual error addition actor. By using the TURNUS profiler, a close to minimum queue size [17] for each queue in the decoder is determined.

Experimental Flow: For the experimental evaluation, a Virtex 7 XC7VX485T-2 FPGA (VC707 Evaluation Kit) was used. The HDL code of the decoder was generated by Xronos and synthesized with the Xilinx XST synthesizer. Following synthesis, placing and routing was applied to produce a final netlist. This netlist was then simulated with Modelsim to extract the switching activity information (SAIF file) of the design. The Xilinx XPower analyser was then used to determine power consumption, using the the design netlist, the design constraints, and the simulation activity SAIF as inputs. Also, all of the results given have a high confidence level meaning that at least 97% of the design nets are found within the SAIF file. Table I shows the synthesis results of the Intra MPEG-4 simple profile decoder with and without clock gating. This example demonstrates that the clock gated decoder uses more slices than the non clock gated one. Even

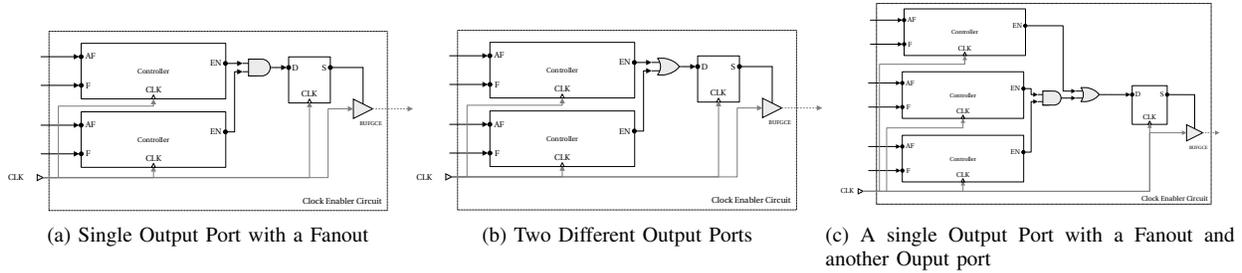


Fig. 3: Clock Enabler Circuit in three different configurations.

Logic Utilization	Non Clock Gated	Clock Gated	Available
Slices	9214	12776	607200
LUTs	21499	25126	303600
BUFGCTRLs	1	32	32
BRAMs	7	7	1030
DSPs	18	18	2850
Max Freq.	109	109	-

TABLE I: Synthesis results of the Intra MPEG-4 simple profile decoder synthesized for Virtex 7 XC7VX485T-2 FPGA, with and without clock gating.

though this represents 28% more slices overall compared to the non clock gated decoder, the clock gating methodology requires only 15% more LUTs. A 50 MHz clock has been given as a synthesis constraint.

Table II depicts the power consumption of the decoder including the circuit of the clock gating methodology. Two test cases were considered: clock gating de-activated and clock gating activated when decoding at maximum throughput.

In Table II, the *Actors Clocks* label only the power consumption of the clock nets of the actor. The *Clocks* cell contains the *Actors Clock* nets, the enabling of clock buffer nets, and the nominal 50Mhz clock net. As a result of clock gating, the *Actors Clocks* consume 26% less power, but due to the decoder running at full speed, the activation rate of the *Logic* and *Signals* nets remain resulting in a total power decrease of 4% (16 mW less).

Clock Gating	Disabled (mW)	Enabled (mW)
Actors clocks	58	43
Clocks	94	80
Logic	25	24
Signals	42	41
Leakage	242	242
Total	403	387

TABLE II: Power consumption of the Intra MPEG-4 SP decoder when the clock gating is disabled/enabled.

A. Power saving efficiency over decoder throttling

As described in Table I the maximum decoder throughput rate is 350 frames per second for a QCIF image (176x144 pixels). For the experiment, the decoder is throttled such that it decodes only 30 images per second for two resolutions QCIF and CIF (384x288 pixels).

Figure 5b reports the power consumption and the activation rate for each actor's clock (found on the Intra MPEG-4

simple profile decoder). The activation rate of the actor's clock demonstrates that some of them have an activation rate of less than 10%. As a result of these activation rates, the power consumption on clocks has drastically fallen by 53.7% for the QCIF resolution and 47.6% for the CIF resolution. As for overall power consumption, the decoder consumes 59mW less for the QCIF resolution and 54mW less for the CIF resolution. Furthermore, out of 31 actors, 15 are almost always on. This means that for the 15 actors, their output buffers never fill up. Further improvements of this methodology could entail detecting which actors do not benefit from clock gating and eliminating the instantiation of unnecessary additional logic.

B. Power saving efficiency over bandwidth demand

In this experiment, the decoder was throttled from 0% to 90% simulating a channel with differing consumption rates. This is an example of clock gating applied not specifically to video decoding applications, but to a general application. Figure 4 represents the power consumption of the decoder

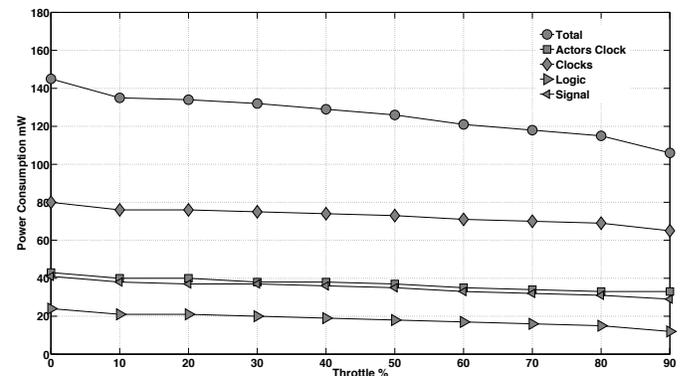
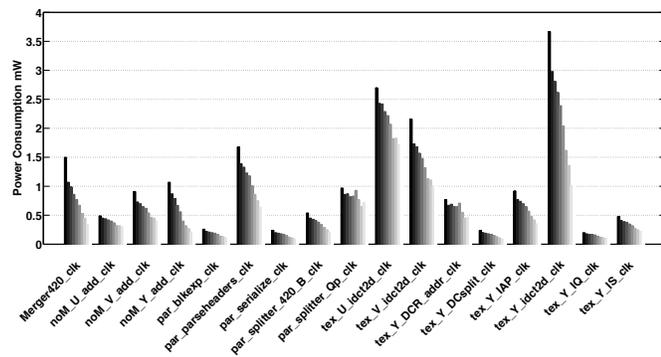
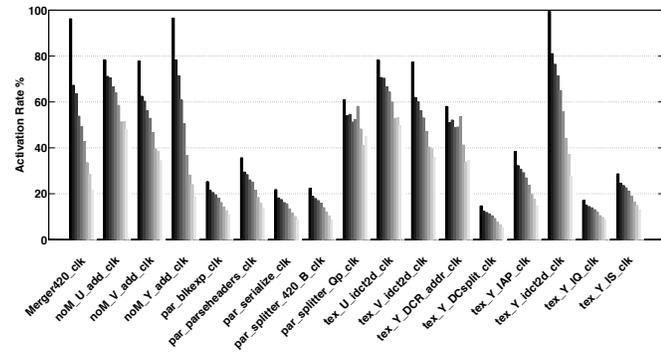


Fig. 4: Power consumption of overall clocks, the signals, logic, and the total dynamic power consumption of the Intra MPEG-4 SP decoder when its output is throttled from 0% to 90%.

when its output is throttled from 0% to 90%. As demonstrated, the total dynamic power consumption has decreased from 145 mW to 106 mW, a power reduction of 27%. Compared to the non clock gate decoder, the dynamic power have been reduced by 34%. Figure 5 reports the power consumption of each clock and their activation rate when throttled. From this graph, the data of 15 actors has been removed due to their activation rate being more than 99%. All actor clock activation



(a) Actors clock power consumption.



(b) Actors clock activation rate.

Fig. 5: Power consumption and activation rate of each clock gated actor clock of the MPEG-4 SP decoder. Median values were retrieved from a MPEG-4 reference QCIF input stimuli (video sequence).

rates decreased while increasing throttle (apart from two cases, `par_splitter_Qp_clk` and `tex_Y_DCR_addr_clk` where the power consumption increased slightly). The decoder used was YUV 420. When it reaches 60%, the decoder throttles the luminance decoding, but the chrominance decoding remains active. This also occurred during a behavioral simulation in Modelsim.

V. CONCLUSION

This paper presents a clock-gating methodology applied to dataflow designs that can be automatically included in the synthesis stage of a HLS design flow. The application of the power saving technique is independent from the semantic of application and does not need any additional step or effort during the "design" of the application at the dataflow program level. The clock gating logic is generated during the synthesis stage together with the synthesis of the computational kernels connected via FIFO queues constituting the dataflow network. Conceivably, these techniques could be extended to other dataflow Methods of Computation.

Experimental results are very encouraging: savings in power dissipation achieved with a slight increase in control logic without any reduction in throughput have been achieved. Unsurprisingly, clock gating is attractive in situations where the design is not used to its full capacity. In these circumstances clock gating is a simple, automatic, and effective

way to recover power otherwise lost in "idle" cycles. As a result, this technique is particularly interesting in applications with dynamically varying performance requirements, when designing to a particular performance point is impossible, and when power consumption is deemed costly.

Further investigations into clock gating should consider more aggressive control logic, whereby control is given to each individual actor, allowing greater flexibility to actor inactivity. Furthermore, it will be necessary to develop tools that partition complex applications onto the limited number of clock domains for more efficient implementations. Lastly, additional considerations could be given to controlling clock speed and, possibly, voltage transitions.

REFERENCES

- [1] Massoud Pedram, "Power minimization in ic design: Principles and applications," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 1, no. 1, pp. 3–56, Jan. 1996.
- [2] Qing Wu, M. Pedram, and Xunwei Wu, "Clock-gating and its application to low power design of sequential circuits," *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on*, vol. 47, no. 3, pp. 415–420, Mar 2000.
- [3] G.E. Tellez, A. Farrahi, and M. Sarrafzadeh, "Activity-driven clock design for low power circuits," in *Computer-Aided Design, 1995. ICCAD-95. Digest of Technical Papers., 1995 IEEE/ACM International Conference on*, Nov 1995, pp. 62–65.
- [4] E. Lee and A. Sangiovanni-Vincentelli, "Comparing models of computation," in *Proceedings of the 1996 IEEE/ACM international conference on Computer-aided design*. IEEE Computer Society, 1997, pp. 234–241.
- [5] Gilles Kahn, "The Semantics of Simple Language for Parallel Programming," in *IFIP Congress*, 1974, pp. 471–475.
- [6] Edward A. Lee and David G. Messerschmitt, "Static scheduling of synchronous data flow programs for digital signal processing," *IEEE Trans. Comput.*, vol. 36, no. 1, pp. 24–35, 1987.
- [7] E.A. Lee and T.M. Parks, "Dataflow process networks," *Proceedings of the IEEE*, vol. 83, no. 5, pp. 773–801, may 1995.
- [8] Syed Suhaib, Deepak Mathaikutty, and Sandeep Shukla, "Dataflow architectures for GALS," *Electronic Notes in Theoretical Computer Science*, vol. 200, no. 1, pp. 33–50, 2008.
- [9] Tzyh-Yung Wu and Sarma B. K. Vrudhula, "Synthesis of asynchronous systems from data flow specification," Research Report ISI/RR-93-366, University of Southern California, Information Sciences Institute, Dec 1993.
- [10] Behnam Ghavami and Hossein Pedram, "High performance asynchronous design flow using a novel static performance analysis method," *Comput. Electr. Eng.*, vol. 35, no. 6, pp. 920–941, Nov. 2009.
- [11] S.C. Brunet, E. Bezati, C. Alberti, M. Mattavelli, E. Amaldi, and J.W. Janneck, "Partitioning and optimization of high level stream applications for multi clock domain architectures," in *Signal Processing Systems (SiPS), 2013 IEEE Workshop on*, Oct 2013, pp. 177–182.
- [12] Simone Casale-Brunet, *Analysis and optimization of dynamic dataflow programs*, Ph.D. thesis, STI, Lausanne, 2015.
- [13] Endri Bezati, *High-level synthesis of dataflow programs for heterogeneous platforms*, Ph.D. thesis, STI, Lausanne, 2015.
- [14] S. Casale-Brunet, M. Mattavelli, and J.W. Janneck, "Buffer optimization based on critical path analysis of a dataflow program design," in *Circuits and Systems (ISCAS), 2013 IEEE International Symposium on*, May 2013, pp. 1384–1387.
- [15] Xilinx, *Analysis of Power Savings from Intelligent Clock Gating*, August 2012, XAPP790.
- [16] "Open RVC-CAL Applications," 2014, <http://github.com/orcc/orc-apps>, accessed 25-February-2014].
- [17] M. Canale, S. Casale-Brunet, E. Bezati, M. Mattavelli, and J. Janneck, "Dataflow programs analysis and optimization using model predictive control techniques," *Journal of Signal Processing Systems*, pp. 1–11, 2015.