LP-HLS: Automatic power-intent generation for high-level synthesis based hardware implementation flow

Original

LP-HLS: Automatic power-intent generation for high-level synthesis based hardware implementation flow / Qamar, Affaq; Muslim, Fahad Bin; Iqbal, Javed; Lavagno, Luciano. - In: MICROPROCESSORS AND MICROSYSTEMS. - ISSN 0141-9331. - 50:(2017), pp. 26-38. [10.1016/j.micpro.2017.02.002]

Availability:

This version is available at: 11583/2666625 since: 2018-05-24T13:28:55Z

Publisher:

Elsevier

Published

DOI:10.1016/j.micpro.2017.02.002

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2017. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.micpro.2017.02.002>

(Article begins on next page)

LP-HLS: Automatic Power-Intent Generation for High-Level Synthesis Based Hardware Implementation Flow

AFFAQ QAMAR, Politecnico di Torino

FAHAD BIN MUSLIM, Politecnico di Torino

JAVED IQBAL, Politecnico di Torino

LUCIANO LAVAGNO, Politecnico di Torino

The abstraction level for digital designs is rising from Register Transfer Level (RTL) to algorithmic untimed or transaction-based, followed by an automated high level synthesis (HLS) flow. However, it is still a significant challenge for chip architects and designers to describe low-power design decisions at the system level. Nowadays, low power design techniques for digital blocks are applied at RTL and there exists no commercial tool or methodology that can automatically derive the power intent from the system level description. The process requires considerable amount of human intervention and various lower-level details that are needed to implement low power schemes at RTL. This research aims to integrate low power techniques, specifically Power Shut-Off (PSO), within a model based hardware flow and to derive an automated Low Power-High Level Synthesis (LP-HLS) methodology. The methodology aims at minimizing the design effort for low power design by deriving low-level power intent automatically for model-based designs, while using high level synthesis to achieve a broad set of target system implementations. All the information that is needed to implement low power techniques is automatically derived from the system-level design using a set of pragmas and a directives file. To illustrate the methodology, three model designs, ranging from simple designs to medium complexity hardware accelerators, are considered. Finally, the power saving results for the design scenarios validate the effectiveness of our LP-HLS methodology.

• **Hardware-Best practices for EDA** • **Hardware-Software tools for EDA** • **Hardware-Hardware description languages and compilation** • **Hardware-Modeling and parameter extraction**

Additional Key Words and Phrases: Hardware accelerators, HLS, RTL, design space exploration, common power format, low power designs, design automation.

1. INTRODUCTION

System-on-chip (SoC) designs are becoming increasingly heterogeneous as they combine multicore architectures with a variety of hardware accelerators to carry out dedicated computational tasks. These hardware accelerators offer several orders-of-magnitude higher power and timing efficiency than a corresponding software implementation [Horowitz 2014]. However, the presence of accelerators aggravates the complexity of SoC design. Moreover, digital designers aim at developing designs that optimize both timing and power, which generally are two conflicting performance objectives. Pareto-optimality can be comfortably achieved only if power specification is considered at the system-level. This is because major architectural decisions related to cost (area), performance and energy utilization can be made only at higher abstraction levels [Muslim et al. 2015]. For example, parallelism versus voltage and clock frequency is much easier to trade off at system-level. Early consideration of effective power management helps relieve the power bottlenecks of today's VLSI design, enabling sustained high-performance operation with desired power consumption levels.

1.1 Model-based designs

As far as behavioral description for the hardware design is concerned, the abstraction level is rising from RTL to algorithmic untimed or transaction-based, followed by an automated high level synthesis (HLS) flow [Liu and Carloni 2013]. HLS takes as input the model-based description of the design, specified in some high level language such as C, C++, SystemC or Simulink, and synthesizes it to generate RTL. By elaborating different sets of constraints, HLS tools allow designers to evaluate multiple implementation alternatives, a process known as Design Space Exploration

(DSE) [Ravi et al. 2014; Cong 2014]. DSE with HLS is already a major leap from DSE with Logic Synthesis, since the former can be achieved by simply changing HLS tool directives, while the latter usually requires one to manually change a detailed hardware description expressed in the form of Verilog or VHDL. Such hardware descriptions at a lower abstraction level are often tried only for one or two architectural options, due to the much slower design and verification cycle [Daoud et al. 2014]. Also, high design productivity requires that complex SoCs consist of 90% reused components [Liu et al. 2012]. This in turn requires soft IP components, which are designed once using high-level languages and implemented into various instances to meet the changing design requirements [Cesário et al. 2002].

1.2 Low power design flow

As far as functional description of hardware is concerned, it has evolved to the extent that model-based design approaches using automated HLS flows have gained popularity within the design community [Qamar et al. 2015]. However, it is still a significant challenge for chip architects and designers to describe low-power design decisions at the system level. This is because; system architects have little or no visibility of the lower-level details that are needed to implement low power schemes at RTL [Zhang 2015]. Similarly, a digital backend designer needs to interact intensively with the system architect as well as with the verification team to formulate an appropriate hardware platform and a low power scheme in order to meet the requirements. This process still involves a considerable amount of human intervention, because nowadays the low power techniques for a digital system are applied at register transfer level (RTL). These techniques include supply voltage and clock control technologies, such as power shut-down, clock gating, dynamic voltage and frequency scaling, and adaptive voltage scaling [Kurimoto et al. 2013]. The power shut-down technology is especially vital to leakage power reduction for battery-driven devices.

1.3 Contribution

This research aims to integrate the aforementioned low power techniques within a model based hardware flow and to derive an automated Low Power High-Level Synthesis (LP_HLS) methodology. This methodology enables the specification of both the behavioral functionality and the power intent at a level of abstraction higher than RTL and the use of high-level synthesis (HLS) tools at the front-end of the ensuing design flow. The behavioral description of the design can be captured using SystemC or C/C++, although in this paper we use SystemC for the sake of illustration. In our proposed flow, the *power intent*, i.e. the set of power-saving techniques to be used and all the information that is needed to implement them, is automatically derived from the system-level design using a set of pragmas and a directives file. Combining this information with technology dependent parameters, a tool that we developed derives low power intent written in the Common Power Format (CPF), or Unified Power Format (UPF), which are widely used standards supported by several commercial EDA tool vendors.

In this work, we use CPF for the sake of illustration, and consider power shut-off (PSO) along with clock gating (CG) to achieve power saving [Benini et al. 2001; Verma et al. 2015]. To illustrate the methodology, we use three examples ranging from simple designs to medium complexity hardware accelerators. These include a 32 bit Ripple-Carry Adder (RCA), a general purpose ALU performing arithmetic and logical operations, and an IDCT block often used in image and video processing. Our methodology aims at minimizing the design effort for low power design by deriving low-level power intent automatically for model-based designs, while using high level

synthesis to achieve a broad set of target system implementations. In particular, we propose:

- A generic power management module description at the system-level, to specify the design context of a hardware block;
- A tool that, for a given system design context automatically generates the low power design directives needed to implement the PSO technique in the back-end flow.

2. RELATED WORK

Most work on low power techniques using CPF and UPF focuses on the RTL, rarely considering system-level implications. An example is [Mathur and Wang 2009; Varma 2015], which focuses on the manual description of the power intent using CPF for RTL implementations generated using HLS. In this work, on the other hand, we propose a methodology to *automatically generate the CPF power specification from a high-level model* just like an HLS tool generates the RTL from a behavioral description.

Benini et al. [2001] and Lin et al. [2015] suggest various energy optimizations starting from software down to circuit-level for electronic systems. Their strategy starts from power optimizations at system-level including energy-aware task scheduling, hardware/software partitioning, power aware architectures using dynamic power management (DPM) policies and code morphing. Similarly, they emphasize adopting low power schemes such as MSV, PSO, CG, and DVFS at the RTL, while using gates with varying transistor widths to achieve additional energy saving at the gate-level.

Schirrmeister [2009], and Benini & Micheli [2000] provide a review of the different abstraction levels for energy measurement and estimation as well as common techniques to optimize a design for low power above RTL. They use a typical JPEG decoder as a test case. The authors advocate the use of a system-level solution early in the design cycle for data-flow dominated blocks and their associated memories.

Zhang et al. [2015] argues that in order to meet strict power requirements, modern day designers may still have to perform manual optimizations on an RTL design described using Verilog or VHDL by applying numerous low power techniques considering functional, structural, temporal and spatial information together. It is extremely difficult to achieve these goals manually in such a complex multi-dimensional space within a limited time. In addition, power scaling requires designers to evaluate and optimize the system architecture as early as possible in the design flow. Certainly, the solution is to raise the level of abstraction beyond RTL to achieve faster power optimization and use automated RTL synthesis tools.

Thus, there exists a general agreement among the designers regarding the significance of system-level power optimization techniques [Ahmed et al. 2014; Bezati et al. 2014]. Commercial efforts in producing pre-RTL power optimization tools have resulted in tools such as Vista Architect from Mentor Graphics and Chip vision's PowerOpt™. Two new working groups were formed by IEEE at the end of 2014 to standardize system-level power modeling for SoC devices [Shuang 2015]. The prospect of considering power intent at the system-level and applying it in combination with HLS thus is an important but rather sparsely explored area and hence a theme of this work.

3. LP-HLS METHODOLOGY

The earlier energy savings are considered in the design cycle, the more savings can be achieved, as shown in Fig. 1. This is because architectural decisions related to cost (area), performance and energy utilization are made at this level [Sinha and

Srikanthan 2014]. The prospect of a completely automated low power flow, where both the logic and power intent can be incorporated into the design at the system-level, seems very promising. In this section, we present a complete description of our LP-HLS methodology, which comprises an HLS flow, a CPF generator to automatically produce CPF power intent, and the final integration into the backend flow.

Fig. 1. Power reduction as a function of various stages of design cycle

3.1 Model-based hardware design

The overview of the proposed design flow is shown in Fig. 2. The power intent is derived by extracting the design related information from the system-level description. This information is read together with the power rules specifications for the PSO, and the technology related parameters to automatically generate the CPF file. The hardware description for low power design is then integrated with the RTL during the later stages of backend flow.

Luckily, even when the RTL is generated automatically from the system-level code written in SystemC, the naming convention of the design hierarchy, i.e. instances, signals and ports is preserved while the design undergoes high-level synthesis, or names for new objects (e.g. fields of SystemC struct-typed I/O signals) can be easily derived automatically from the high level model.

In our flow, the power intent is also derived from annotations in the SystemC code and a design configuration file that refers only to high-level information. This makes it easier for the design architect to select among the underlying power optimization choices, thanks to the higher level of abstraction.

Fig. 2. Overview of the low power design model

3.2 CPF generation tool

Our automated CPF generation tool relies upon the use of `#pragma` directives in the system-level design file/s as well as a designer-written intent specification file, to identify power related information. Figure 3 shows the general structure of the tool.

Fig. 3. Basic structure of CPF generation tool

The technology specification (tech. spec. in the Fig. 3) file, as the name implies, contains technology related information using directly the corresponding CPF syntax.

In particular, it contains; (i) the set of libraries (typically the worst, best and nominal cases), (ii) the information about the power nets that will be created during physical design, and (iii) the nominal conditions (e.g. voltage levels for various power nets) which will be used by the various power modes afterwards.

```
***** INPUT File 1: system_level_model.cc *****
#PRAGMA default-domain
SC_MODULE (topmodule)
{
    sc_in< bool > clk, reset;
    sc_signal <bool> pse_sig, iso_sig;
    .....
#PRAGMA power-domain-instance MY_DOMAIN
    idct_module XLXI_3;
    .....
    .....
    Power_control XLXI_2; // POWER MANAGEMENT UNIT
    SC_CTOR (topmodule)
    {
        .....
#PRAGMA shutoff-condition
        XLXI_2.pse(pse_sig);
        .....
#PRAGMA isolation-condition
        XLXI_2.iso(iso_sig);
        .....
    }
};
```

```
***** INPUT File 2: spec_file *****
#PRAGMA isolation-cells
cells ISO_FENCE0N_X* EN
isolation_rule MY_DOMAIN always_ON low
```

```
***** OUTPUT File: automate.cpf *****
set_design topmodule

## Specify power domain ##
create_power_domain -name PD_default -default
create_power_domain -name MY_DOMAIN -instances {XLXI_3} -
shutoff_condition {!XLXI_2/pse}

## Isolation Cells ##
define_isolation_cell -cells "ISO_FENCE0N_X*" -enable EN /
-valid_location to

## Isolation rule ##
create_isolation_rule -name iso_1 -from MY_DOMAIN /
-to always_ON -isolation_condition {XLXI_2/iso_en} /
-isolation_output low -isolation_target from
```

Fig. 4. A simple example of input pragmas and corresponding CPF output

The specification (spec.) file contains pragma directives for the power rules that are needed to successfully implement PSO strategy. These rules depend upon design related parameters. This includes rules such as operating corners for multi-mode-multi-corner (mmmc) analysis, analysis views for power domains and the definition of power rules such as isolation, state-retention and power switch rules.

The CPF generator needs to extract from the system-level model file information such as, the module name of the switchable power domain(s), and the instance name of the signal that defines the shut-off condition, and that will be used to drive the power switches. This is done by inserting a `#pragma` directive before the respective instance name. The tool searches for those unique pragmas, creates token strings for the instances and stores the information for later processing. The name of the power domain can either be assigned by the designer (e.g. `MY_DOMAIN`), or be generated by the parser (e.g. `switchable_domain_<n>`).

Once all the information from the input files is gathered into custom data structures, the tool generates power rules in CPF format corresponding to each pragma directive, using the information from the configuration files.

Figure 4 shows an example of how two pragma directives in the SystemC source and some information in the specification file are combined to generate a simple CPF output file.

3.3 LP-HLS flow

After discussing the model-based design approach using HLS and describing our tool that automates the power intent generation, we now discuss the complete LP-HLS flow.

Figure 5 summarizes the flow of operations that are required to go from a system-level specification to a power-optimized gate-level implementation, using the proposed Low Power High-Level Synthesis (LP-HLS) methodology and a standard RTL-to-gates flow. Depending on the application, different constraints (e.g., performance, area cost, and power) must be satisfied during the various phases of the flow. Macii et al. [1998] suggests that, “when the target is a low-power application, the search for the optimal solution must include, at each level of abstraction, a design improvement loop”. In such a loop, a power analyzer/estimator (shown in gray in Fig. 5) ranks the various design, synthesis, and optimization options, and thus helps in selecting the one that is potentially more effective from the power standpoint. However, this methodology requires the availability of power estimators, as well as synthesis and optimization tools, which provide accurate and reliable results at various levels of abstraction.

While adapting a purely algorithmic model to be the input of HLS, the designer must instantiate a Power Management Block (PMB) that decides exactly when the power can be switched ON or OFF, and if necessary keeps the design waiting (e.g. by gating the clock) while the power stabilizes. Of course, this PMB must be part of the always ON domain. Thereafter, the RTL is generated through HLS by performing the necessary steps i.e. by specifying the target technology, the micro-architecture choices and the scheduling constraints. In parallel to HLS, our tool processes the system-level power intent and generates the CPF file, as described above.

The switching activity information for accurate power analysis is extracted by simulating the RTL with the original SystemC test-bench, by using the automatically generated SystemC wrapper. Power aware logic synthesis is then performed by first reading the target libraries, as specified in the CPF file, and by enabling the application of coarse-grained clock gating logic (both in HLS and logic synthesis).

As in the standard power-aware flow, the RTL design is then read and elaborated, followed by technology mapping of the cells to be used as clock-gated integrated cells

(CGIC). The power intent is then read, followed by setting the timing constraints and synthesizing the design. The switching activities are annotated and the power intent is applied, followed by the verification of the power structure to check if the low power cells have been correctly inserted in the design according to the rules specified in the power intent file. Incremental optimization is performed and finally the gate-level netlist is obtained and checked for logic equivalence against the input RTL.

Fig. 5. Low Power High-Level Synthesis LP-HLS methodology flow

4. DESIGN SCENARIOS AND TEST CASES

The example cases that are adopted are simple, yet interesting enough to illustrate the proposed methodology. We start with a simple design of a hierarchical 32 bit ripple-carry adder (RCA) structure to better understand and apply the investigated methodology. The module processing the 16 most significant bits (MSB_RCA₁₆₋₃₁) is selected to be placed in a switchable power domain, to enable low-power processing of 16-bit numbers. Then we move to a larger design, namely an ALU processor comprising eight modules to perform arithmetic as well as logical operations. Here,

we choose to power down the division and multiplication modules when they are not needed. Finally, we consider a more complex design, an inverse discrete cosine transform (IDCT) module, which finds its application in JPEG decoders.

Since our target is to test and validate the flow, we formulate synthetic application scenarios for our example designs. In this section we discuss very briefly, the test bench setup as well as the PMB design for each.

Fig. 6. General structure of test bench for design cases

4.1 Test bench structure

The test bench setup includes a stimulus generator that drives the input control and data signals to the design under test (DUT). A monitor block logs the outputs and checks their validity, as depicted in Fig. 6.

In the case of the RCA and the ALU, the stimulus generator uses pseudo-random number generators (PRNGs) to generate input streams as well as random values of control logic to select between the power ON and OFF conditions for the switchable power domains. Of course, the probabilistic models should reflect the behavior of the real application scenarios. However, this is a well-studied problem that is outside the scope of this paper [Li et al. 2009; Tschanz et al. 2003; Brooks et al. 2000].

For the RCA and the ALU, we use both a synthetic switching activity profile in which the switchable domain stays ON for 10% of the time, one in which it is ON for 90% of the time, and one in which it is on for 50% of the time. For the IDCT design, on the other hand, we consider its real-life usage inside a JPEG decoder.

Please note that the PMB could be made more complex, in order to require some minimum number of idle cycles before shutting off the power, but again these considerations are outside the scope of the paper [Bartolini et al. 2014; Ahmed et al. 2015]. We will see later that each transition between power states requires at least four clock cycles, which would suggest considering a threshold to trigger the transition to be at least four cycles.

4.2 Design under test (DUT)

An important issue worth considering while employing power gating is to prevent floating states from propagating from the Power Switchable Domain (PSD) to the default domain. It is also necessary to save the state of some flip-flops in the design before switching off a part of the design. Isolation cells (ISO) are responsible for isolating the always-on units from the floating values of outputs from the power-gated units. They are typically placed on the outputs of the shut-off power domain during the physical placement stage [Chadha and Bhasker 2013], as shown in Fig. 7.

RET in Fig. 7 represents the state retention cells, which have the ability to retain their states even if the primary power is shut-off. The retention cells are optional and are needed only if the state of some sequential logic in the power switchable domain

must be preserved. The header power switches which are inserted during the physical implementation phase provide the ability to cut-off the supply voltage to the switchable domain. The rules for the insertion of these low power cells are stated in the power intent file provided to the logic synthesis tool.

Fig. 7. General structure of DUT

In order to ensure no loss of data during the power-up process, we chose to add a first-in-first-out (FIFO) module at the input of the DUT, as depicted in Fig. 7. This buffers the input values while the module is still in the sleep mode and the data input is coming in during the power-up process. Moreover, clock gating is introduced by using the SystemC clock gating capability provided by the HLS tool [Cadence 2013].

4.3 Power management block (PMB)

Figure 7 also shows the PMB in the default domain. This is added as a SystemC module to produce the control signals in the correct sequence to power gate the module instances of the power switchable domain(s), as indicated in the Fig. 8. The power control flag in Fig. 8 represents the signal in the functional model that activates the power up/down process. This signal must be identified manually by the designer, e.g. by using activity profiling as discussed in section 4.2, and it captures the inactivity intervals of the design. During the power down process, isolation must happen before state retention, followed by power shut-off, while the reverse sequence must be followed during the power up process. The signals for isolation, state retention and power shut-off as provided by the power control module are `iso_en`, `ret_en` and `pse` respectively, as shown in Fig. 8.

In this case the power up/down sequences consume four clock cycles, plus the number of cycles that are required to bring the power rails to the required supply voltage. Many surveys [Benini and Micheli 2000; Verma et al. 2015] suggest combining PSO with CG in order to maximize power savings. When the switchable domain is powered down, we gate its clock nets at the time of power down. The SystemC pseudo code of our (purely illustrative) PMB block is presented as algorithm1.

Fig. 8. Low Power High-Level Synthesis LP-HLS methodology flow

Algorithm 1. PMB algorithm

```

// ***** Initialization phase
*****
1: Initialize signals (isolation, retention, power switch,
clock gating);
2: wait(); // wait for one clock cycle
// ***** PMB logic
*****
3: while (true){
4:   if (Powershut-off == Enable){
5:     isolation = ON;
6:     wait();
7:     retention = ON;
8:     wait();
9:     power_switch = ON;
10:    clk_gating = ON;
11:  }
12:  else{
13:    clk_gating = OFF;
14:    power_switch = OFF;
15:    wait();
16:    retention = OFF;
17:    wait();
18:    isolation = OFF;
19:  }
20: }
9: end //infinite while loop

```

4.4 Hardware accelerators – Test cases

This section briefly overviews the example cases used to validate the proposed methodology.

4.4.1 32 bit ripple carry adder (RCA)

In order to apply the PSO technique to the RCA, it is modeled in SystemC as a hierarchical module, comprising of two 16 bit RCAs, called MSB_RCA₁₆₋₃₁ and LSB_RCA₀₋₁₅ in Fig. 9. Both modules are functionally identical, but the MSB_RCA is assigned to the switchable power domain, and hence it has a few additional ports to drive power management operations. The output multiplexers also use $P_{\text{shut-off}}$ as a select signal to choose between the valid outputs.

Fig. 9. 32 bit ripple-carry adder (RCA)

4.4.2 ALU processor

The ALU processor, also modeled in SystemC, is capable of performing arithmetic as well as logical operations. The encoder is driven by the control logic (SEL), from the stimulus generator and selects between unique opcodes assigned to each function. Since the multiplication and division blocks consume most of the hardware resources, they are also the most power hungry blocks [Hoang et al. 2013]. Thus, they become ideal candidates to be assigned to two separate switchable power domains as shown in Fig. 10. The power ON/OFF sequence for the respective domain is triggered based on the SEL signal.

Fig. 10. Arithmetic logic unit (ALU) processor

4.4.3 IDCT

IDCT is a well-known algorithm used in data compression standards (e.g. JPEG). A JPEG decoder performs various operations like variable length decoding (VLD), zigzag scanning (ZZ), de-quantization (DQ), inverse discrete cosine transform (IDCT), color conversion, and reordering on the compressed image. A typical JPEG decoder architecture depicting all these operations is shown in Fig. 11.

This work deals with a synthesizable SystemC implementation of a JPEG IDCT decoder. A 2D-IDCT is performed by first performing 1D-IDCT on each of the columns in the matrix followed by 1D-IDCT on each row. The design architecture consists of concurrent processes with the communication between them taking place at transaction level.

IDCT is the major contributor to the overall complexity of a JPEG decoder [Sona et al. 2014]. This gives us a strong reason to power gate the IDCT unit of the JPEG decoder which will also serve in validating our LP-HLS methodology. To employ power gating, activity profiling of the IDCT design is performed first to identify the idle periods in the design so as to correctly apply the power optimization strategies [Muslim et al. 2015]. This is done by simulating the design with a dedicated SystemC test-bench. This enables us to identify the idle periods in the design as well as the signal in the design that can be used to trigger the power control mechanism. Thus, the power control signals in this design, unlike the previous design examples (refer to sec. 4.4.1 and 4.4.2), that use synthetic data to produce the power control signals, come from real-time simulation of a JPEG-IDCT decoder. This provides us with a more realistic scenario to validate the LP-HLS methodology.

Fig. 11. JPEG decoder using IDCT module

The power control signals in the correct sequence are being provided by the PMB, which is added as a separate SystemC module, as depicted in Fig.11. These signals are used to power gate the IDCT module using the power intent captured in the CPF file. The clock gating of the design is performed by adding a separate SystemC clock gating module. This module performs coarse-grain clock gating and since it is enabled by the same signal used to trigger power gating, it clock gates all those instances that are being power gated as well. Without clock gating, the clock buffer tree continues to propagate the system clock even to the power gated modules, thus consuming dynamic power. Thus, it is a general practice to apply clock gating along with the power gating technique.

It should be noted that the power up/down sequences take four clock cycles to complete. We thus need to prevent a loss of data, which is being fed to the hardware accelerator during the power-up process, by incorporating a first-in-first-out (FIFO) buffer into the design as a separate SystemC module. In addition to the isolation cells and the power switch cells, the state retention cells already explained in previous sections, are also added in the power switchable domain here.

5. RESULTS

The experimental setup consists of a SystemC description of the design-under-test (DUT) i.e. 32 bit RCA, ALU processor, and IDCT, respectively. The entire design flow is carried out using tools provided by Cadence Design Systems. This mainly includes Cadence C-to-Silicon Compiler [Cadence 2014] for high-level synthesis and RTL Compiler [Cadence 2013] to implement the backend flow. Several different implementations have been used to validate the results, namely the design without any power optimization techniques, and the design with clock gating and power gating. The general experimental setup uses the LP-HLS flow to automatically derive RTL for design behavior and CPF for power intent description.

The power computation is obtained from the power model included in the standard cell libraries. In this work, where we are interested in validating our proposed LP-HLS methodology, we used the 45nm NanGate Open Cell Library which supports low power cells for ASIC implementation. The power analysis using library power models relies upon the expected state of the signals at the standard cells boundary and their transition activity.

In order to verify our LP-HLS methodology, we sweep the values of both static probability and toggle rate of the power control pins of the PMB. The static probability, which dictates leakage power, accounts for the total operation workload. The toggle rate, which is associated with the dynamic power, determines the number of transitions per unit time for the power control signals.

For the RCA and ALU examples, where we are relying on synthetic input vectors, we control the utilization (static probability) of the PSD by specifying a usage percentage w.r.t the rest of the design. The transition rate for the ON/OFF states of the power control signals is based on the utilization factor. On the other hand, the IDCT uses real-time input vectors from the JPEG decoder, so we keep the utilization fixed while we sweep the toggle rate to observe changes in dynamic power.

The power optimization result for the RCA example is presented in Table. I. The MSB_RCA comprises of 31% of the total size of the design, hence it is a good candidate for the power gating. The rest of the design includes the PMB, LSB_RCA and two multiplexers to compute sum, as mentioned in Fig. 9. The power analysis is performed for RCA without power optimization at 50% usage (i.e. the MSB_RCA is used in 50% of the clock cycles). From the graph in Fig. 12a, it is clear that the total power consumption is almost double than the power consumption with power optimization even when MSB_RCA is active for 90% of the time, mainly because of clock gating.

Table I. Power optimization of RCA w.r.t. MSB_RCA (power switchable domain)

Operations Workload (MSB_RCA)	P _{static} (μ W)	P _{dynamic} (μ W)	Cell Area (μ m ²)
Without power optimization @ 50%	70	255	Total $\rightarrow$ 3362 MSB_RCA $\rightarrow$ 1051 (31%)
90%	35	104	
70%	31	97	
50%	28	93	
30%	25	90	

Similarly for the ALU, we assign the division (DIV) and multiplication (MULT) operations to two separate PSDs. Together they contribute 48% of the total design area as mentioned in Table. II. We then analyze the power consumption by assigning different utilization to DIV and MULT. The choice of 1% usage for DIV and 10% for MULT corresponds to a hypothetical integer workload, while the sweep between 30-10% for DIV and 60-40% of MULT can be regarded as a DSP workload. Fig. 12b shows the leakage as well as dynamic power results for the ALU. Since the cell area of the power switchable logic is almost half of the total logic, the power saving is even greater than in the previous example.

Table II. Power optimization of ALU w.r.t. DIV and MULT (power switchable domains) operations

Operations Workload (DIV – MULT)	P _{static} (μ W)	P _{dynamic} (μ W)	Cell Area (μ m ²)
Without power optimization @ 10% – 40%	190	1072	Total $\rightarrow$ 27361 DIV $\rightarrow$ 8219 (30%) MULT $\rightarrow$ 4790 (18%)
30% – 60%	140	341	
20% – 50%	133	283	
10% – 40%	129	237	
1% – 10%	86	160	

The IDCT design scenario uses vectors from the JPEG decoder to obtain power results with and without power optimization, as indicated by the first two rows of Table. III, respectively. The next three rows were obtained by increasing the toggle rate of the power shut-off signals to 4 times, 8 times and 32 times of the original toggle rates that were obtained by simulating the RTL with a realistic usage scenario. More toggling would result in the IDCT module being powered on and off

more frequently and this would result in an increase in the dynamic power consumption of the design while the static power remains the same. The reason for no change in the static power is that the static probabilities of these pins remain the same for all the cases. The study of this behavior with multiple toggle rates is helpful in estimating the extent of switching beyond which any power savings would be overshadowed by the resulting dynamic power consumption. A visual comparison of the effect of power optimization for the IDCT is provided in Fig. 12c.

Table III. Power optimization of IDCT (power switchable domain)

Operations Workload (IDCT)	$P_{static}(\mu W)$	$P_{dynamic}(\mu W)$	Cell Area (μm^2)
Without power optimization	572	12570	Total $\rightarrow$ 44124 IDCT $\rightarrow$ 42271 (96%)
32x Toggling of enable	240	849	
8x Toggling of enable	240	726	
4x Toggling of enable	240	680	
JPEG usage	240	655	

b. Power consumption for RCA

a. Power consumption for ALU

c. Power consumption for IDCT

Fig. 12. Static and dynamic power consumptions for design scenarios

6. CONCLUSIONS

This work proposes an LP-HLS methodology that derives power intent from the system-level description of the digital design. The framework is based on (1) a generic power management module description at the system-level, to specify the design context of a hardware block, and (2) a tool that, for a given design context automatically generates the low power design directives needed to implement the PSO technique in the back-end flow. To illustrate the methodology, three example hardware accelerators ranging from simple designs to medium complexity were developed in SystemC. These include a 32 bit Ripple-Carry Adder (RCA), a general purpose ALU performing arithmetic and logical operations, and an IDCT design often used in image and video processing. The methodology aims at minimizing the design effort for low power design by deriving low-level power intent automatically for model-based designs, while using high level synthesis to achieve a broad set of target system implementations. Power analysis was carried out for the design scenarios by varying the usage of the designs. The power optimization results at the end validate the accurate derivation of power intent by using our LP-HLS methodology.

REFERENCES

- Horowitz, M. 2014. Computing's energy problem (and what we can do about it). In *IEEE Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 10-14.
- Muslim, F. B., Qamar, A., and Lavagno, L. 2015. Low power methodology for an ASIC design flow based on High-Level Synthesis. In *IEEE 23rd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*.
- Liu, H. Y., and Carloni, L. P. 2013. On learning-based methods for design-space exploration with high-level synthesis. In *Proceedings of the 50th Annual Design Automation Conference, ACM*, 50.
- Liu, H. Y., Petracca, M., and Carloni, L. P. 2012. Compositional system-level design exploration with planning of high-level synthesis. In *Proceedings of the Conference on Design, Automation and Test in Europe*, 641-646.
- Cesário, W. et al. 2002. Component-based design approach for multicore SoCs. In *Proceedings of the 39th annual Design Automation Conference DAC*, 789-794.
- Qamar, A., Muslim, F. B., and Lavagno, L. 2015. Analysis and implementation of the Semi-Global Matching 3D vision algorithm using code transformations and High-Level Synthesis. In *Proceedings of the 81st IEEE Vehicular Technology Conference (VTC Spring)*, 1-5.
- Kurimoto, M. et al. 2013. Verification work reduction methodology in low-power chip implementation. In *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 18(1), 12.
- Mathur, A., and Wang, Q. 2009. Power Reduction Techniques and Flows at RTL and System Level. In *Proceedings of the 22nd IEEE conference on VLSI Design*, 28-29.
- Benini, L., Micheli, G. D., and Macii, E. 2001. Designing low-power circuits: practical recipes. In *IEEE Circuits and Systems Magazine*, 1(1), 6-25.
- Lin, C. Y. et al. 2015. The Design and Experiments of A SID-Based Power-Aware Simulator for Embedded Multicore Systems. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 20(2), 22.
- Schirrmeister, F. 2009. *Design for low-power at the electronic system level*. White paper, Chip Vision Design Systems, v1.1.
- Benini, L., and Micheli, G. D. 2000. System-level power optimization: techniques and tools. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 5(2), 115-192.
- Zhang, Z., Chen, D., Dai, S., and Campbell, K. 2015. High-level Synthesis for Low-power Design. *IPSJ Transactions on System LSI Design Methodology*, 8(0), 12-25.
- Shuang Yu. *IEEE Standards Association*. [online] 2014, retrieved on 29 May 2015, http://standards.ieee.org/news/2014/ieee_p2415_p2416_wgs.html
- Macii, E., Pedram, M., and Somenzi, F. 1998. High-level power modeling, estimation, and optimization. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(11), 1061-1079.
- Chadha, R., and Bhasker, J. 2013. Architectural Techniques for Low Power. In *An ASIC Low Power Primer*, Springer New York, 93-111.
- Cadence design systems, user manual, 2013. *Cadence C-to-Silicon Compiler User Guide Product Version 13.20*.
- Hoang, T. T. et al. 2011. Power gating multiplier of embedded processor datapath. In *IEEE 7th Conference on Ph. D. Research in Microelectronics and Electronics (PRIME)*, 41-44.
- Sonka, M., Hlavac, V., and Boyle, R. 2014. Image processing, analysis, and machine vision (4th. ed.). CL Engineering.

- Cadence design systems, user manual, 2013. *Cadence Low power in Encounter RTL compiler*, product version 10.1.
- Ravi, S., and Joseph, M. 2014. High-Level Test Synthesis: A Survey from Synthesis Process Flow Perspective. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 19(4), 38.
- Cong, J. 2014. From design to design automation. In *Proceedings of the ACM International symposium on physical design*, 121-126.
- Bartolini, A., Hankendi, C., Coskun, A. K., and Benini, L. 2014. Message Passing-Aware Power Management on Many-Core Systems. *Journal of Low Power Electronics*, 10(4), 531-549.
- Ahmed, R., Bsoul, A., Wilton, S. J., Hallschmid, P., and Klukas, R. 2014. High-level synthesis-based design methodology for Dynamic Power-Gated FPGAs. In *24th IEEE International Conference on Field Programmable Logic and Applications (FPL)*, 1-4.
- Ahmed, R., Wilton, S. J., Hallschmid, P., and Klukas, R. 2015. Hierarchical Dynamic Power-Gating in FPGAs. In *Applied Reconfigurable Computing, Springer International Publishing*, 27-38.
- Li, S., Ahn, J. H., Strong, R. D., Brockman, J. B., Tullsen, D. M., and Jouppi, N. P. (2009). McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *42nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO-42*, 469-480.
- Tschanz, J. W., Narendra, S. G., Ye, Y., Bloechel, B., Borkar, S., and De, V. (2003). Dynamic sleep transistor and body bias for active leakage power control of microprocessors. *IEEE Journal of Solid-State Circuits*, 38(11), 1838-1845.
- Brooks, D., Tiwari, V., and Martonosi, M. (2000). Wattch: a framework for architectural-level power analysis and optimizations, ACM, Vol. 28, No. 2, pp. 83-94.
- Daoud, L., Zydek, D., and Selvaraj, H. 2014. A survey of high level synthesis languages, tools, and compilers for reconfigurable high performance computing. In *Advances in Systems Science, Springer International Publishing*, 483-492.
- Verma, G., Kumar, M., and Khare, V. 2015. Low Power Techniques for Digital System Design. *Indian Journal of Science and Technology*, 8(17).
- Bezati, E., Brunet, S. C., Mattavelli, M., and Janneck, J. W. 2014. Coarse grain clock gating of streaming applications in programmable logic implementations. In *Proceedings of the IEEE Conference of Electronic System Level Synthesis (ESLsyn)*, 1-6.
- Sinha, S., and Srikanthan, T. 2014. Dataflow Graph Partitioning for Area-Efficient High-Level Synthesis with Systems Perspective. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 20(1), 5.