Floorplan-based FPGA Interconnect Power Estimation in DSP Circuits

Ruzica Jevtic, Carlos Carreras and Vukasin Pejovic
Dpto. de Ingeniería Electrónica
Universidad Politécnica de Madrid
Madrid, Spain
{ruzica, carreras, vule}@die.upm.es

ABSTRACT

A novel high-level approach for estimating power consumption of global interconnects in data-path oriented designs implemented in FPGAs is presented. The methodology is applied to interconnections between modules and depends only on their mutual distance and shape. The power model has been characterized and verified with on-board power measurements, instead of using low-level estimation tools which often lack the required accuracy (observed errors go up to 350%). The results show that most of the errors of the presented power model lie within 20% of the physical measurements. This is an excellent result considering that in [2] it is shown that there is already a 20% variation in net capacitance due to the different routing solutions given by router for the same placement.

Categories and Subject Descriptors

B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Gate arrays*

General Terms

Algorithms, Design

Keywords

FPGA, low power, interconnects, power estimation

1. INTRODUCTION

FPGAs have become an attractive solution for various embedded designs due to their ability for reconfiguration and significantly lower cost compared to ASICs. As they are aimed to implement many different designs, a large number of routing switches is used in order to obtain flexible interconnections, and look-up tables (LUTs) are used for logic, as they are capable of implementing any logic function for the given number of inputs. However, this type of architecture prevents optimized design implementations because it utilizes an excessive number of additional transistors and routing resources, which in turn, contribute to a significant increase in the power consumption of the design.

Interconnects represent the dominant source of power consumption in FPGAs (45% for Spartan 3 and 65% for Virtex-II for the internal benchmarks in [7, 16]). Circuit routability should be characterized at the earliest possible time, as to avoid implementations that exceed the power constraints. The interconnect power estimation problem has been studied in depth for ASICs. However, FPGAs have a more complex routing structure composed of different wire types and switching matrices that impose serious limitations on the estimation of the wire capacitance.

Complex DSP systems implemented in FPGAs consist of arithmetic components which are connected by data-path buses. Early prediction of interconnect power is necessary, since the programmable switch matrices consume a significant amount of power and their number and position need to be optimized. However, the methods found in the literature target power estimation of both, control-oriented and data-oriented, designs. In order to estimate the power of data buses more accurately, a more centralized model is needed.

In this paper, we present a high-level interconnect power model that is capable of giving fast and accurate estimates for data-path buses in DSP circuits for any given distance between the connected components. There are only two input parameters to the model: the relative position of the components and the ordering of the pins on the components' boundaries. We explore the accuracy of the proposed interconnect model for a wide range of input parameters, signal components and design positions on a chip. The estimates are compared to the real power values obtained from on-board measurements. Results show that, in spite of the inherent variability in the net capacitance due to different router solutions (as large as 20% in average, according to [2]), the accuracy of the proposed model is similar to the accuracy of the post-placement estimation model in [2]. The model in [2] uses a high number of input parameters known only after the placement, whereas the usage of the model proposed here is not limited to this stage of the design flow. It can be used at higher levels of abstraction as well, because the only information that is needed is the design floorplan. Consequently, the proposed model is suitable for integration with floorplan-aware high-level synthesis aimed at power optimization.

The paper is organized as follows. Section 2 highlights the previous work regarding interconnect power estimation. Section 3 describes the measurement system. In Section 4, the

SLIP'09, July 26–27, 2009, San Francisco, California, USA.
Copyright 2009 ACM 978-1-60558-576-5/09/07 ...$5.00.

power estimation model developed for global interconnects is presented in detail. It is followed by the experimental results in Section 5, and conclusions in Section 6.

2. RELATED WORK

Most of the existing interconnect estimation techniques are applied at the post-placement design level [2, 12], as the information on global routes is extremely scarce at the higher levels of abstraction. Furthermore, in [14], it was noted that the delay of some interconnects can change significantly (up to 2 or 3 times) by changing the seed of the placement algorithm for generic logic in FPGAs.

There are only few methodologies that try to predict routing parameters before placement. The methodology described in [8], models power estimation at the post-synthesis level, by using the information about the fanout of the design. The main goal is not to achieve precise power estimation, but to guide the user to identify hot-spots in the circuit. Besides, results are given only for the circuit which was used as a part of the characterization set and therefore, the applicability of the model is not very clear.

A stochastic approach was used in order to predict interconnection lengths of communication links in FPGAs in [13]. The model is applicable to floorplanning, as it depends on the parameters such as area dimensions of the connected regions, the Manhattan distance between them, the number of connections and the number of long lines in a channel. The model is used in order to determine the delay of the connection, which is found to have a linear dependence on the wirelength. This approach has some similarity to the approach for interconnect power estimation presented here, regarding the modelling of the component by its area constraint, while accounting for the local routes separately. However, they assume that no more than two regions are connected, only long lines are used for the routing, and the connected regions are separated by a significant distance on the chip, whereas the work presented here has no such limitations.

A pre-placement methodology for predicting individual wire length and routing demand of each net in designs implemented in FPGAs is presented in [3]. A circuit is presented as a set of nodes, and a heuristic approach is applied in order to simulate the placement of the circuit and determine the bounding box for every single net. The resulting bounding boxes are further used in order to predict the channel width for the routability of the design. The design has to be mapped before applying this methodology.

Some efforts have been made to use Rent's rule to estimate the average wire length, needed for the calculation of dynamic power consumption [4]. Rent's Rule is an empirical metric used to quantify circuit complexity. The method has been widely used for ASICs and microprocessor architectures ([5], [20]). However, the errors introduced by this method when applied to FPGAs can be significant [18], especially for non-hierarchical designs where different parts of the design exhibit different routing congestion, wire length and number of interconnects. In these cases, the characterization of the routing structure of the whole design with only one parameter is hardly possible.

3. MEASUREMENT METHODOLOGY

The first goal of the measurement methodology is to obtain the capacitance values for the global wires used for routing, as this information contains proprietary technology details and it is unavailable to us. In particular, Xilinx Virtex-II Pro devices are the target technology considered in this work. In order to accomplish this goal, we use a common method to extract the effective capacitance ([16, 7]) as follows.

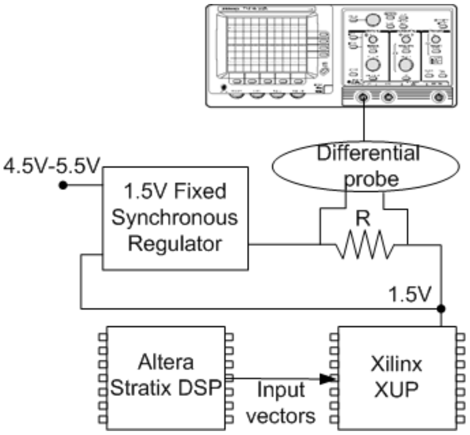

Figure 1: Measurement setup

The measurement setup is presented in Fig. 1. The development of the measurement setup was inspired by the work presented in [6], and it was also described in [11]. The system contains two FPGA boards: a XUP board from Xilinx [19] and a Stratix DSP Development board from Altera [1]. The board from Altera is used for loading the simulation vectors to the XUP board. The XUP board serves for measuring the power of a specific design. We use a resistance at the entrance of the core power supply to the chip, and for each test design, we measure the voltage over this resistor, which enables us to calculate the current going from the supply. The value for this resistance was set to 10 ohms as this value provides maximum measurement precision and ensures the correct functionality of the buck-switching PWM regulator on the XUP board. The functionality of the chip itself is guaranteed by a direct feedback from the chip power supply to the input of the regulator (see Fig. 1). The power is then obtained as the product of the power supply voltage and the measured current.

The power is measured for simple designs that consist of a multiplier or adder core with registered inputs and outputs (in further text both referred to as modules), that is replicated between one and four times in the design, in order to improve the accuracy of the measurements.

We repeat the following set of measurements: static power (when no input signals nor clock are injected), the power of the clock circuitry (when inputs are set to 0's) and various power measurements performed for sets of 10000 input signal vectors with gaussian distributions, with the design in two different positions on the chip: one where the design is placed very close to the I/O pins, and the other where it is placed far from them (see Fig. 2). By subtracting the two values obtained for the dynamic power consumption in the two positions, we are able to obtain the value that corresponds to the power consumption of the interconnect difference between them.

Although the static power varies during circuit operation due to the temperature increase, the designs we have used are small, so it has been assumed that the static power

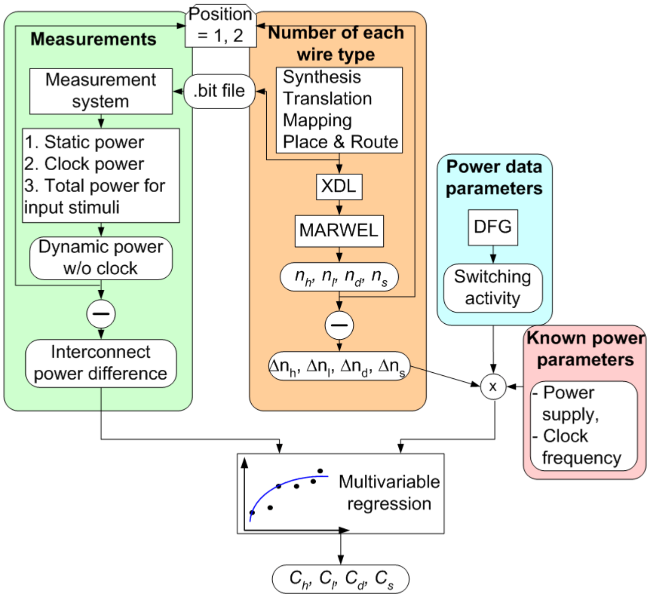

Figure 2: Methodology for effective capacitance extraction

increase would be negligible. In order to confirm this assumption, we have repeated steps 1) - 3) listed on the left-hand side of Fig. 2 for two different frequencies (50Mhz and 100MHz) for several most power-consuming test designs (containing multipliers implemented in LUTs). We isolated the dynamic power for each frequency and observed that, indeed the relationship between the two obtained dynamic power values for each design corresponded to the relationship between these two frequencies.

In the Xilinx Virtex-II Pro device, various routing resources can be identified based on long, hex, double and direct wires. We model the effective capacitance of each resource as the capacitance of the routing wire together with the programmable switch that drives the wire, as in [7]. After placement and routing of a design, the Xilinx tool ISE creates a native circuit description file (.ncd) which represents the physical circuit description of the input design. We have developed a tool in C++ called MARWEL (Measurement of ARchitectural WirE Lengths), based on the Graph Template Library (GTL) [10], which obtains the length and the number of the different wires used from the XDL file. This file is the text version of the .ncd file and is created by the Xilinx Design Language (XDL) tool. Therefore, for each interconnect i that goes from or to I/O pins in the design, we obtain the number of hex, long, double and single wires used for its routing: n_{hi}, n_{li}, n_{di} and n_{si}.

As the inputs and outputs are registered and there is no glitching in the wires that connect I/O pins with inputs and outputs of the modules, we are able to obtain the switching activities sw_i, of the routing wires from simple data flow graph simulations. The value of the switching activity for each interconnect is then multiplied by the corresponding number of wires of each type used for its routing.

Four parameters are required in order to calculate the power of the interconnects. Two of them are known, as the power supply has a value of 1.5V for Virtex-II Pro devices, and the clock frequency is fixed to the value used in our measurements.

As we can not obtain the measured interconnect power value separately from the rest of the design power, we substract the obtained dynamic power for two different positions of modules on the chip. Thus, we eliminate the logic power that is the same for both implementations. This allows us to break-out the power consumed in the interconnects. Therefore, we can express the power difference of a design in the two measured positions as:

$$\begin{aligned} P_1 - P_2 = V_{dd}^2 \cdot f \cdot (C_h \cdot &\sum_{i=1}^{I_1+I_2+O} [(n_{hi}^1 - n_{hi}^2) * sw_i] \\ +C_l \cdot &\sum_{i=1}^{I_1+I_2+O} [(n_{li}^1 - n_{li}^2) * sw_i] \\ +C_d \cdot &\sum_{i=1}^{I_1+I_2+O} [(n_{di}^1 - n_{di}^2) * sw_i] \\ +C_s \cdot &\sum_{i=1}^{I_1+I_2+O} [(n_{si}^1 - n_{si}^2) * sw_i]) \end{aligned} \tag{1}$$

where P_1 and P_2 are the measured dynamic power of the design with the modules in the positions far from and near to the I/O pins respectively, C_h, C_l, C_d, C_s are the effective capacitance of the hex, long, double and single wires respectively, I_1, I_2 are the word-lengths of the two input operands and O is the word-length of the output. The design position is identified through the superscripts 1 (far) and 2 (near). A multivariable regression over a number of measurements for modules with various operand word-lengths is applied, as to obtain the effective capacitance for all types of wires. Once we have these values, we can obtain the power consumption of any interconnect, by using the information about the number of different wire types used for its routing. Beside the interconnect power, the module power is easily obtained by subtracting the power of the interconnects from the dynamic power of the design.

4. HIGH-LEVEL INTERCONNECT MODEL

We have developed a high-level power estimation model (HLM) for the interconnections between n modules, by applying a rectilinear Steiner tree (RST) algorithm to the centers of the module pins. The module pin center is defined as the center of the minimal bounding box that includes all of the module pins connected to the other module. The distance between the modules is computed in unit-lengths. The unit-length is the distance between two neighbouring CLBs.

A detailed analysis of the type of wires used for global routing has demonstrated that three different routing zones can be identified. The first one, corresponds to the minimal distance between the modules, where only direct and double lines are used. The second one, corresponds to distances smaller than some specific distance d_l, where three types of wires are used for routing: direct, double and hex. Finally, the third zone, corresponds to distances larger than d_l, where all four types of wires are used.

During the routing phase, the minimum cost paths are selected to implement the connections. The cost consists of two parts: the one that accounts for the competition between different nets for the same wiring segments, and the one that reflects the routing delay associated with the routing segment [17]. If there is no congestion in the circuit, we assume that the interconnect power per unit-length tends to be constant. In this case, the router will try to minimize the total interconnect delay, which would reflect in the minimization of the interconnect capacitance, resulting

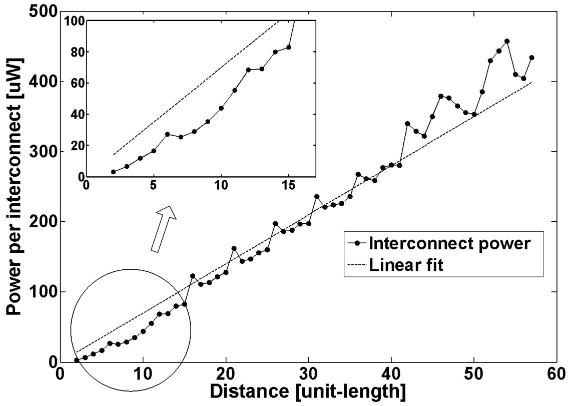

Figure 3: Power per interconnect between the modules A and B

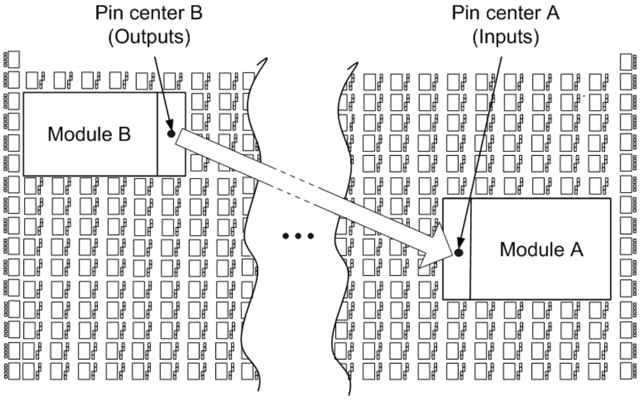

Figure 4: Simulation setup when the interconnects between two modules are considered

also in reduced interconnect power. Although this assumption is quite straightforward in ASICs since the wires have a unique metal capacitance, in FPGAs not only are the routes composed of different wire types, but these wires also pass through a determined number of switch matrices which have a great impact on the total route capacitance.

In order to validate the assumption of linear increase in FPGA interconnect power with the distance, we have plotted the measured power per interconnect between two modules A and B (in this case two multipliers), versus their distance in Fig. 3. The outputs of module B are connected to the inputs of module A as shown in Fig. 4. The position of module A is fixed near the I/O pins on the right-hand side of the chip. The position of module B is varied from the position nearest to module A, to the position near the I/O pins on the left-hand side of the chip, opposite to module A, and further up, along the I/O pins on the left-hand side of the chip. The distance between the modules is computed as the Manhattan distance between the pin centers of the connected modules, which are marked in Fig. 4. The power per interconnect is computed after the place-and-route of the design for each position, by using the effective capacitance of the routing wires (obtained as explained in the previous section), and using the information about the length and the number of different wire types used for the interconnections (obtained from the tool MARWEL). First, the total power of the interconnects is computed and then, it is divided by the number of interconnects to obtain the power per interconnect. The power values are normalized with the switching activity, because data dependencies are not significant for the purpose of this analysis. It can be seen in Fig. 3, that the power per interconnect has almost a linear dependency on the distance between the modules.

Although the dependence seems linear, the linear fit does not describe the interconnect power accurately for the smallest distances. The upper left corner of Fig. 3 shows the power corresponding to the shortest distances. It can be seen that the linear fit overestimates the interconnect power, resulting in increased estimation errors. As the capacity of the long lines is the highest, and the router ceases to use this type of lines beyond distance d_l, this results in a significant decrease in the power consumption of the interconnects below this specific distance.

As a result, we use the following power model for the average power per interconnect:

$$P_{int} = \begin{cases} k_3 \cdot L, & d = d_m \\ k_2 \cdot (d - d_m) + k_3 \cdot L, & d_m < d < d_l \\ k_1 \cdot (d - d_l) + k_2 \cdot (d_l - d_m) & \\ +k_3 \cdot L, & d > d_l \end{cases} \tag{2}$$

where P is the power per interconnect, d_l is the specific distance beyond which the router starts using long lines, d_m is the minimal distance between the module pin centers, L corresponds to the distance between the module pins and their pin center as will be explained next, d is the distance between the modules (the length of the RST), and k_1, k_2, k_3 are the coefficients calibrated by multiple regression analysis over measured power values for different distances between the modules.

The critical distance d_l, has been obtained empirically through MARWEL for all combinations of two different modules: an adder and a multiplier, and it is the same in all cases. For the connection of n modules, assuming that the second routing zone applies to the lines in the proximity of both modules and it can be divided into two equal parts, a new distance limit d_l^{RST}, is computed as follows:

$$d_l^{RST} = n \cdot \frac{d_l}{2} \tag{3}$$

In the case of the two modules described previously, at the minimal distance of one unit-length, the B outputs and A inputs are completely aligned. In real-case designs, the connected inputs and outputs may not be necessarily placed in the same order, specially when considering connections from or to I/O pins, as the I/O pin location also relates to the board design. Besides, as the module pins are located on the boundaries of the module, the distance between the module pin centers d, does not correspond exactly to the length of the interconnects. The parameter L models the limitations that occur due to the module shape and size and is computed as:

$$\begin{aligned} L &= \sum_{k=1}^{n} L_k \\ L_k &= \frac{\sum_{i=1}^{I_{1k}+I_{2k}} l_{k,i}^{in} + \sum_{j=1}^{O_k} l_{k,j}^{out}}{I_{1k}+I_{2k}+O_k} \end{aligned} \tag{4}$$

where $l_{k,i}^{in}$, $l_{k,j}^{out}$ are the Manhattan distances from the module pin center of the k^{th} module to its input pin i and its output pin j, respectively. I_{1k} and I_{2k} are the number of the k^{th} module input pins used for the connection, and O_k is the number of its output pins used for the connection. Thus, the parameter L models the power increase at the shortest distances (when compared to the linear fit) where the position of all connected module pins can no longer be approximated by the position of the module pin center.

It is important to emphasize the importance of using only the module pin center in order to represent all the connections of the module. We compute the distance between the modules by applying the RST algorithm on module pin centers. The problem of finding RST is NP-complete and is often computationally very expensive. However, we do not apply this algorithm on a pin-to-pin basis. Instead, the Steiner tree connects the module pin centers, so the algorithm does not depend on the word-length of the module's operands, but only on the number of connected modules. As this number is relatively small compared to the number of routed nets, the computation time to obtain the RST is highly reduced.

Consequently, the power obtained in (2) represents the average power per connection, which is later multiplied by the sum of the switching activities, sw_i, of all the connections, in order to obtain the total interconnect power estimate of a given data-path bus:

$$P_{total}^{int} = P_{int} \cdot \sum sw_i \qquad (5)$$

5. EXPERIMENTAL RESULTS

We split the model evaluation into two sets of experiments. In the first set, we evaluate the accuracy of the effective capacitances obtained for different types of wires by measuring the power on-board. In the second set, we evaluate the power model presented in section 4 by using these effective capacitances and the wire-length provided by MARWEL. We use several DSP circuits to compare the power estimates to the physical measurements for various input signal statistics and module positions. Additionally, we compare the accuracy of the presented model to the accuracy of the low-level tool XPower.

5.1 Effective capacitance evaluation

The experiments were performed on four different size multipliers and five different size adders. The characterization set used for the multivariable regression considered only the power values corresponding to the input signals with autocorrelation coefficient equal to 0, as they provided the largest consumption and thus, the best accuracy. Furthermore, for each module and each autocorrelation coefficient, we computed two values, δP, that corresponds to the power difference for the module positions 1 and 2, in the left-hand side of equation (1), and P_{cap} that corresponds to the right-hand side of the same equation, computed from the obtained effective capacitance values, listed in Table 1.

Fig. 5 shows the relative errors when the computed P_{cap} is compared to the measured δP. It can be observed that, the resulting discrepancy is always smaller than 12.5% and probably occurs due to the local wire parasitics as explained in [7].

Table 1: Effective capacitances for different wire types.

Wire type	Capacitance per CLB [fF]
Long	178.133
Hex	86.578
Double	71.47
Direct	≈ 0

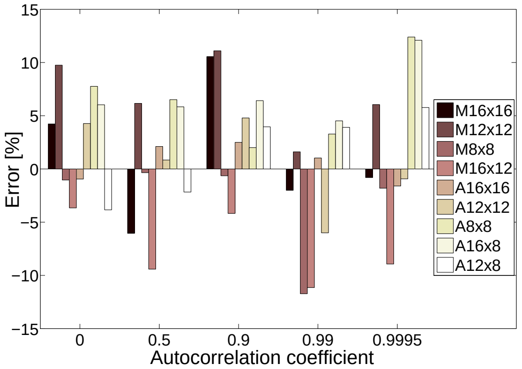

Figure 5: Error of the interconnect power computed with the effective capacitance values.

5.2 Interconnect model evaluation

As previously mentioned, the interconnect power model needs two parameters: the distance between the modules (i.e. the length of the RST), and the ordering of the pins on the component boundaries (in order to compute parameter L). We have built a library containing parameter L for the arithmetic modules. For the components not belonging to this library, L had to be computed additionally. The length of the RST was obtained by using Geosteiner [9]. The coordinates of the module pin centers that are needed for the RST computation, were obtained from the floorplans of the placed designs. As the model uses only these coordinates and does not require any other placement information, it can be easily integrated into power optimization techniques that perform high-level synthesis combined with floorplanning. In these cases, the accuracy of the model will depend on the accuracy of the floorplan estimate.

The experiments are divided into three subsets. The first one considers the connections between all combinations of two different modules; an adder and a multiplier. This is the characterization set used to obtain the coefficients k_i of the proposed model. In order to account for the interconnect capacitance noise [2], five different placements were generated for each distance (varied in a wide range) between the two modules. Next, for each distance, a mean power value was computed. Finally, the coefficients k_i were obtained by using multivariable regression over the mean power values for various distances.

Fig. 6a shows the relative errors for each different placement in the characterization set versus the distance between the two modules. The estimates obtained after placement, are compared to the power values computed by using MARWEL after place-and-route and the effective capacitance values. It can be seen that, in most cases, the error lies in the range [-20%,+20%], with an absolute maximum error of 40%. Furthermore, most of the largest errors are obtained

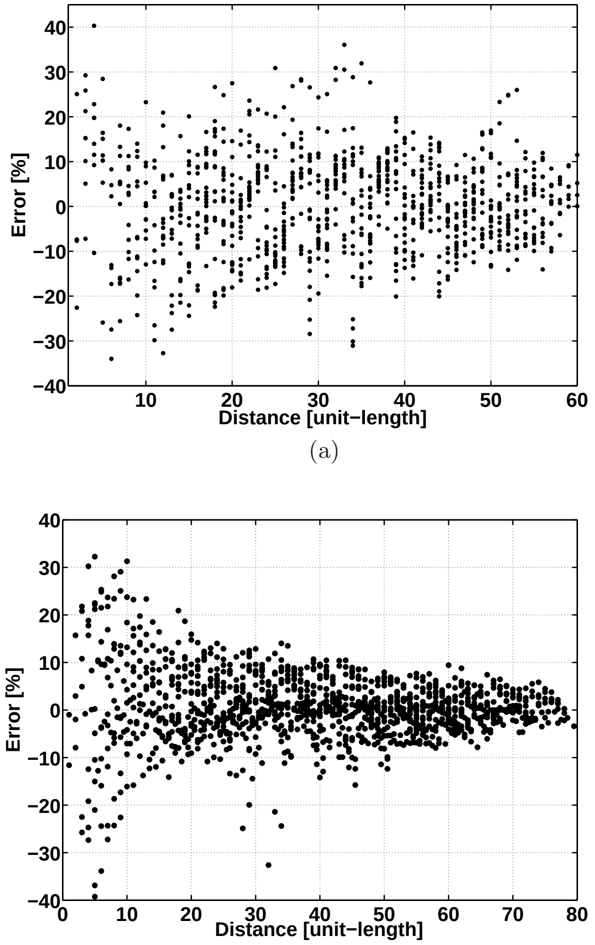

(a)

(b)

Figure 6: Interconnect estimation errors for the connections a) between two modules , b) between module and IO pins

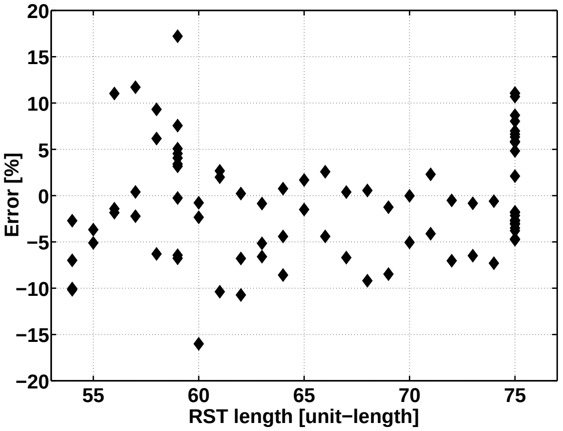

Figure 7: Interconnect estimation errors for the connections between five modules

for the smallest distances, so their impact on the error of the absolute total power estimate of a design will be very small, as the shortest interconnects represent a small portion of the total power. Another important observation is that the coefficients k_1, k_2 and k_3 obtained for the different combinations of the two modules were practically identical. This means that the same three coefficients can be applied to any combination of modules.

The same procedure has been repeated for the connections between both types of modules and the I/O pins. Still, different values for the coefficients k_i were obtained (in further text referred to as k_{io}). We believe that this effect occurs because the router uses much tighter bounds when routing the connections from or to I/O pins, compared to the routing inside the chip core. The error performance of the model in this case is presented in Fig. 6b. It can be seen that the power variation between different placements is much smaller than the power variance in Fig. 6a, leading to the same conclusion about the different routing bounds.

It was noted that the values of the first two coefficients, k_1^{mod} and k_2^{mod}, are clearly higher than the values of k_1^{io} and k_2^{io}. However, the third coefficient is smaller. This is due to the shape of the module and the influence of the direct wires used for the minimal distance routing. When two modules are separated by the unit-distance, they are completely aligned, and thus, all the lines can be routed as direct lines. The capacitance of the direct wires is close to zero, which leads to the small value for the coefficient k_3^{mod}. However, in the case of the interconnections between a module and the I/O pins, many lines are going around the module, leading to a higher number of double and hex lines, which results in a higher value for the coefficient k_3^{io}.

As a conclusion, we have two sets of coefficients used for the interconnection estimation. One is to be applied when the connections lie inside the core of the chip, and the other, when connections to or from the I/O pins are considered. Once obtained, the coefficient sets k_{mod} and k_{io}, are used for interconnect power model verification in the following experimental sets.

Fig. 7 shows the relative errors of the interconnect power model when applied to the second set of experiments, which consists of a design with five connected modules. The positions of four modules are fixed, while the position of the remaining module is varied throughout the chip. The coefficients k_{mod}, obtained from the experiments considering two modules, are used here in order to obtain power estimates. Power values are obtained by using MARWEL and effective capacitance values. It can be seen that the model provides very good estimates in the range [-20%,+20%]. In both, Fig. 6 and Fig. 7, a random distribution of the error residuals can be seen. This distribution suggests that a linear model presented by (2) is a good fit for observed data.

The third set of experiments was designed to test DSP circuits. In this analysis, we have included the switching activity obtained from bit-level DFG simulations. For each net, the switching activity was computed separately as the average number of transitions during one clock cycle. On the one hand, it was summed over all bits in a signal word in order to obtain the total switching activity of the particular data-path bus between the modules according to equation 5. On the other hand, the value of the switching activity for each interconnect was multiplied by the corresponding number of wires of each type used for its routing obtained from MARWEL, and also by their effective capacitance values in order to obtain the power of each interconnect. These power values were summed over all global interconnects in the design as to form the measured interconnect power.

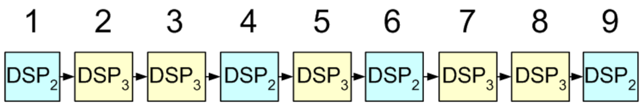

Figure 8: SYSTEM block schematic

Table 2: The number of different arithmetic module types in three configurations of $SYSTEM$ design

Benchmark	Operator-Impl.	Number
$SYS1$	$Mult_{LUT}$	9
	$Mult_{EMB}$	22
	$Adder$	22
$SYS2$	$Mult_{LUT}$	13
	$Mult_{EMB}$	18
	$Adder$	22
$SYS3$	$Mult_{LUT}$	5
	$Mult_{EMB}$	26
	$Adder$	22

The evaluation set consists of two groups: small DSP designs and large designs that represent industrial applications. In the first group, three DSP designs that implement different arithmetic expressions can be identified. They have the following functions:

$$\begin{aligned} DSP_1 &= (x_1x_2+1)x_3x_4 + (256x_1+x_2) \\ DSP_2 &= ((x_1+x_2)(x_3+x_4) + x_1x_2)x_2(x_3+x_4) \\ DSP_3 &= (x_2x_3)x_2 + (x_1+x_3)x_2 \end{aligned} \tag{6}$$

The second group consists of three different configurations of a large DSP design $SYSTEM$, and $CORDIC$, a design taken from [15] representing industrial application. The large design consists of four modules of type DSP_2 and five modules of type DSP_3 connected in a determined order as presented in Fig. 8. Three different configurations of the $SYSTEM$ design are obtained by varying the number of multipliers implemented in LUTs and embedded multipliers. Their characteristics are presented in Table 2.

Table 3 shows the results for each benchmark when data with different values of autocorrelation coefficient are applied to its inputs. Table also includes the number of slices and embedded multipliers used by each design. The results for DSP_1 are obtained for four different placements. The first placement is achieved without using any area constraints. For the second one, the relative positions of the modules are kept as in the first placement, but all the modules are placed far from the I/O pins. Third, a bounding box with the size of a quarter of the FPGA surface is applied as an area constraint, and it is placed on the opposite side of the pins. In the fourth position, an area constraint for only one of the multipliers is created and it is placed far from the I/O pins. Evaluating power in four different positions also enabled us to confirm that the interconnect power values computed using MARWEL and effective capacitances could serve as a fair substitute for direct power measurements. For each design position, we substracted the computed interconnect power from the measured dynamic power. The results, which represent the logic power, should be the same in all four positions. Indeed, the maximum relative difference between these logic power values was 2.05%.

Table 3: Relative errors for the proposed model (HLM) and XPower (XP), for different autocorrelation coefficients.

Bench.	Slic.	Emb.	ρ	Er(HLM)[%]	Er(XP)[%]
$DSP1$	290	0	0	-19.9	174.45
			0.9	-19.4	167.56
			0.99	-20.3	164.21
			0.9995	-22.54	157.06
			0	-5.97	38.48
			0.9	-5.55	35.47
			0.99	-4.75	36.25
			0.9995	-3.62	39.48
			0	-1.7	38.95
			0.9	-1.07	33.84
			0.99	-0.37	32.71
			0.9995	2.86	33.80
			0	6.51	87.27
			0.9	8.06	87.16
			0.99	10.59	89.33
			0.9995	15.3	99.97
$DSP2$	192	2	0	-8.09	258.45
			0.9	-13.63	233.50
			0.99	-18.73	216.23
			0.9995	-5.6	245.27
$DSP3$	212	2	0	6.32	328.79
			0.9	4.5	316.48
			0.99	-1.93	281.70
			0.9995	-9	246.91
$CORD$	591	0	NA	-9.22	NA
$SYS1$	1972	22	0	-16.86	NA
			0.9	-16.26	NA
			0.99	-16.01	NA
			0.9995	-16.34	NA
$SYS2$	1692	18	0	-16.17	NA
			0.9	-14.83	NA
			0.99	-13.7	NA
			0.9995	-12.88	NA
$SYS3$	1444	26	0	-19.21	NA
			0.9	-18.12	NA
			0.99	-17.36	NA
			0.9995	-17.21	NA

Beside the estimates obtained by the proposed model, Table 3 also includes the error of XPower with respect to the reference power values, as described next. We have used the advanced power reports provided by XPower (from ISE 10.1). First, we ran the Modelsim gate-level timing simulation of the placed-and-routed design, and as a result obtained a .vcd file. This file contains detailed information on the toggling rates and frequencies of all the signals in the design, and it was used as the input simulation file for XPower. We were unable to use the new tool, XPower Analyzer, as the power values are displayed in milliwatts for now, and as such, it can only be used for large designs where this precision does not have a significant impact on the accuracy.

The information about the power of each individual element is listed in the XPower advanced report. We generated a script that parses the XPower report and extracts the information on power of the global nets in the circuit. Then, the power values for all global nets were added to obtain the total interconnect power value. The benchmark taken from [15] was not evaluated by XPower as the input data are unavailable to us. In this case, for the sake of comparison with the power model presented here, we have assumed 0.5 switching activity on all nets in the design. We did not evaluate different configurations of the large design with XPower neither, since the generated .vcd file was 5 GB large for the recommended resolution of 1ps used by Modelsim, and the XPower tool was not able to parse this file correctly.

The results are presented in Table 3. For the HLM, we observed that the highest underestimates were reported for the designs where the modules were placed tightly next to each other (such as DSP_1^1 and all configurations of the large design), and thus, generated congestion in the routing lines. Still, the highest detected error was -22.54%, with almost all of the errors lying in the range [-20%,+20%], thus, proving the applicability of the model in an optimization process.

On the other hand, the XPower tool shows large overestimate errors. We believe that this is due to the fact that the static power reported by XPower is a constant for the Virtex II Pro device, and that the tool is calibrated to estimate the power of large designs. The power values for interconnects are higher than their real values in order to compensate for the increase in static power due to the higher temperature generated by the activity of the large designs. In order to have some kind of indicator about the accuracy of XPower for larger designs, we chose one configuration of the large design, then parted the input vector data set into two halves, simulated and applied XPower to each half, and finally computed the mean of the power results. It was observed that XPower accuracy did not increase although the test design is approximately 10 times larger than the small test designs. However, it has not been possible to confirm that static power compensation is the source of XPower errors, as the power values obtained by using MARWEL and the effective capacitance values represent only the dynamic interconnect power, and the increase in static power due to the larger design activity is not reflected in their values. Thus, the comparison of the measured power value for this large DSP design to XPower estimates is equivalent to the same comparison for the small designs. The accuracy of XPower stays the same as expected since the static power is excluded from the analysis.

We believe that XPower is aimed at coarse architecture optimization (order of watts), while HLM is also aimed at detailed architecture refinement (order of milliwatts).

6. CONCLUSION

We have presented a high-level approach to estimate the power consumption of interconnections in data-path oriented FPGA designs. The proposed methodology has been verified through physical on-board measurements. The results show that the accuracy of the models in most cases, lies within 20% of the power measurements. The model performance has been explored over a wide range of input parameters, signal components and module positions on a chip. The accuracy of the model has also been verified through on-board measurements of some DSP benchmarks. The results suggest the applicability of the estimation model in high-level power optimization techniques combined with floorplanning.

The work presented here considers only modules with registered inputs and outputs, as is the case in pipelined designs. However, in non-pipelined designs, the amount of glitching can represent a high percentage of the total power. Our future work is oriented toward extending the models to include glitching effects and routing congestion.

7. ACKNOWLEDGMENTS

This work was supported in part by the Spanish Ministry of Education and Science under project TEC2006-13067-C03-03.

8. REFERENCES

[1] Altera. www.altera.com.

[2] J. H. Anderson and F. N. Najm. Interconnect capacitance estimation for fpgas. In *Proc. on ASP-DAC*, pages 713–718, Jan. 2004.

[3] S. Balachandran and D. Bhatia. A-priori wirelength and interconnect estimation based on circuit characteristics. In *Proc. on SLIP*, pages 77–84, 2003.

[4] D. Chen, J. Cong, and Y. Fan. Low-power high-level synthesis for fpga architectures. In *Proc. on ISLPED*, pages 134–139, Aug. 2003.

[5] P. Christie and D. Stroobandt. The interpretation and application of rent's rule. *IEEE Trans. on VLSI*, 8(6):639–648, Dec. 2000.

[6] Y. S. D. Elléouet and N. Julien. An fpga power aware design flow. In *Proc. on PATMOS*, pages 415–424, Sept. 2006.

[7] V. Degalahal and T. Tuan. Methodology for high level estimation of fpga power consumption. In *Proc. ASP-DAC*, pages 657–660, Jan. 2005.

[8] M. French, L. Wang, T. Anderson, and M. Wirthlin. Post synthesis level power modelling of fpgas. In *Proc. on FCCM*, pages 281–282, Apr. 2005.

[9] GeoSteiner. http://www.diku.dk/geosteiner/.

[10] GTL. http://www.infosun.fim.uni-passau.de/gtl/.

[11] R. Jevtic and C. Carreras. Power estimation of embedded multiplier blocks in fpgas. *IEEE Trans. on VLSI*, to be published, 2009.

[12] P. Kannan, S. Balachandran, and D. Bhatia. fgrep - fast generic routing demand estimation for placed fpga circuits. In *Proc. on FPL*, pages 37–47, Aug. 2001.

[13] T. Mak, P. Sedcole, P. Y. K. Cheung, and W. Luk. Interconnection lengths and delays estimation for communication links in fpgas. In *Proc. on SLIP*, pages 1–9, Apr. 2008.

[14] V. Manohararajah, G. R. Chiu, D. P. Singh, and S. D. Brown. Difficulty of predicting interconnect delay in a timing driven fpga cad flow. In *Proc. on SLIP*, pages 3–8, Mar. 2006.

[15] Opencores. http://www.opencores.org/.

[16] L. Shang, A. S. Kaviani, and K. Bathala. Dynamic power consumption in virtex-ii fpga family. In *Proc. FPGA*, pages 157–164, Feb 2002.

[17] N. Sherwani. *Algorithms for Physical VLSI Design Automation*. Cluwer Academic Publisher, Boston/Dordrecht/London, 1999.

[18] A. Singh and M. Marek-Sadowska. Efficient circuit clustering for area and power reduction in fpgas. *ACM Trans. on DAES*, 7(4):643–663, Oct. 2002.

[19] Xilinx. www.xilinx.com.

[20] P. Zarkesh-Ha, J. A. Davis, and J. D. Meindl. Prediction of net-length distribution for global interconnects in a heterogeneous system-on-a-chip. *IEEE Trans. on VLSI*, 8(6):649–659, Dec. 2000.