2006-01-01

Reducing Energy in FPGA Multipliers Through Glitch Reduction - Clock Power and Digit-Serial Addendum

Nathaniel Rollins

Michael J. Wirthlin
wirthlin@ee.byu.edu

Reducing Energy in FPGA Multipliers Through Glitch Reduction - Clock Power and Digit-Serial Multipliers Addendum

Nathaniel Rollins and Michael J. Wirthlin,
nhr2@ee.byu.edu, and *wirthlin@ee.byu.edu*
Department of Electrical and Computer Engineering
Brigham Young University, Provo, UT. 84602

Abstract

In a previous paper it was shown that reducing the amount of glitches in digital designs can significantly reduce the amount of dynamic power consumption. Pipelined multipliers and a bit-serial multiplier design were used to show this. The paper failed to mention how much of the dynamic power consumption was due to the clock distribution. Also the only digit-serial multiplier digit size investigated was a digit size of 1. This paper addresses the issue of dynamic clocking power and includes results of digit-serial multipliers with larger digit sizes.

1 Introduction

Field programmable gate arrays (FPGAs) are an attractive architecture choice for digital signal processing (DSP) applications. However, the relatively large amount of power that FPGAs consume make them less attractive than application-specific integrated circuits (ASICs) for wireless and handheld DSP applications. Consequently, energy consumption is becoming a critical design parameter for high-performance DSP designs.

In a previous paper [1] it was shown that a large amount of dynamic power consumption is often due to unproductive signal transitions called *glitches*. It was shown that pipelining can be used to significantly reduce the amount of glitching in an FPGA design and consequently reduce the amount of dynamic power consumption. Digit-serial techniques, are a form of pipelining, was also shown to reduce power consumption.

This previous study failed to show how clocking power can contribute a significant portion of the overall dynamic power consumption. It also only provided digit-serial results with a digit size of 1. This addition to the previous study addresses these two issues. First the energy metrics used in the first paper are used to examine how digit-serial multipliers of larger digit sizes can be used to reduce energy consumption. Next we show that clocking power can make up a significant amount of dynamic power consumption.

The programmable nature of FPGA interconnect results in an interconnect structure with significantly larger loading than custom circuits. The signal buffers, pass transistors and other programmable switching structures significantly increase the capacitive load of signal nets over dedicated metal wires. This loading burden increases both the delay of interconnect as well as the power. Several studies suggest that the primary contribution in global system timing is the global interconnect [2]. It has also been shown that most of the power dissipated by an FPGA design occurs in the interconnect[3].

Because of the relatively large capacitive loading of programmable interconnect, the switching activity of individual signal wires will have a significant contribution to the dynamic power of the circuit. Much of the dynamic switching power can be wasted in unproductive circuit "glitches". Signal glitching refers to the transitory switching activity within a circuit as logic values propagate through multiple levels of combinational logic. While glitching is not unique to FPGAs, the relatively high capacitive loading of programmable interconnect places a much higher power cost to signal glitching for FPGAs. Previous studies

have shown that power dissipation caused by glitching can makeup a significant amount of total dissipated power[4, 5].

An important technique for reducing FPGA power consumption is to reduce the amount of signal glitching within the circuit. Pipelining is one technique for reducing signal glitches. Previous studies have shown that pipelining can be used to reduce power by 90% [6]. A pipelined design has less logic between registers and therefore is less prone to glitching. Pipelining an FPGA design can come at little or no cost since flip-flops are included in every FPGA logic block and often go unused. Digit serial techniques, a form of pipelining, can also be used to reduce signal glitching in arithmetic circuits [7].

This paper will describe a methodology for estimating the power associated with glitches within FPGA datapath circuits. This paper will quantify the dynamic glitch power, energy per operation, energy delay, energy throughput, and energy density of each multiplier used in the study. This methodology will be applied to a wide variety of multiplier circuits including non-pipelined multipliers, pipelined multipliers (with various pipelining levels), and digit-serial multipliers. The methodology presented and results from the study can be used by designers and high-level synthesis tools to properly select arithmetic operators based on energy consumption.

This paper will begin by describing the importance of dynamic transient power in FPGA circuits. A methodology will be described for carefully estimating the static, dynamic, and glitch dynamic power of datapath circuits. This methodology will be applied to a combinational multiplier to identify the unproductive glitching power component. Next, the energy measures used to evaluate power will be introduced. These will include energy per operation, energy delay, and energy density. The paper will apply this methodology to a number of non-pipelined multiplier circuits and highlight the significant contribution of glitches in circuit power. The paper will continue by applying this methodology on pipelined circuits and digit-serial circuits. These circuit types will be compared by evaluating the "operation energy" of each implementation style.

2 Dynamic Transient Power

Dynamic power makes up a large portion of the total amount of power consumed by an FPGA design. FPGA interconnect is largely responsible for dynamic power consumption. It has been shown that the interconnect of an FPGA accounts for the majority of the area on an FPGA chip and also accounts for the majority of the power dissipated by FPGA designs[8]. The amount of power consumed by the interconnect and clock tree can account for up to 86% of total dissipated power[8]. Unnecessary and unproductive use of FPGA interconnect will therefore be very costly in terms of power.

Unproductive interconnect activity is usually caused by glitching. Glitching refers to spurious signal transitions on interconnect lines caused by unequal logic or interconnect delays. For example consider the signal activity of an N-bit ripple carry adder. When new inputs arrive at the adder, all N-bit sums are computed simultaneously but the carry bits must ripple from the least significant bit up to the most significant bit. The most significant bit of the adder could switch N times due to this rippling. Only the final transition can be called a productive transition and so any other transitions are called glitches. The carry-out of the 32nd bit of a 32-bit carry chain will have on average 2x more useless transitions (glitches) than useful transitions per cycle and the sum output will have on average 1.5x more useless transitions per cycle[9]. Fast carry chains in FPGAs may significantly reduce the glitching caused by such an example, however FPGA designs are laden with glitches caused unequal line lengths leading to combinational logic block (CLB) logic.

Estimating the glitches that occur in FPGA designs is important in order to understand how much power is consumed by glitching. A static simulator cannot estimate dynamic signal activity and thus is not sufficient for accurate glitching power analysis of an FPGA design. A previous study demonstrated how static simulations of FPGA circuits can under estimate the circuit signal activity. In that study, the dynamic power estimation was 24% less than the actual power consumption [4]. In designs with larger amounts of glitching (such as a multiplier) the accuracy of such a static power model will be even worse.

An accurate power estimation should take signal glitching into account. To produce an effective

power estimation, a back-annotated timing model is required in order to quantify net delays. A timing simulation which uses this back annotated timing model will then be able to simulate glitches so that the amount of glitching can be determined. A power analysis based on this timing simulation will produce a more accurate power consumption estimation.

ModelSim together with Xilinx's ISE tools can be used effectively to model dynamic transient signal activity and produce an accurate power consumption estimation. The Xilinx tools can be used to generate a back-annotated timing file which can be used by ModelSim to effectively model glitching through a timing simulation. The amount of glitching reported by ModelSim depends on the granularity chosen for the simulations. This study finds that a resolution of 100ps is sufficient. ModelSim simulations can be captured and recorded in a file which reports the switching activity of every net in a design. The switching activity of each net can be analyzed and tabulated into two categories: unproductive glitching transitions and productive signal transitions. Power estimation tools such as XPower[10] can then use this report to estimate the power consumption of the design.

The overall power consumption of a single circuit component can be broken into three categories: normalized static power, dynamic glitching power and the remaining dynamic power.

Static Power The static power of an individual circuit module is obtained by dividing the total static power of the device by the relative size of the circuit (i.e. # Circuit LUTs / Total LUTs). For circuits such as multipliers with large signal activity, this component is relatively small.

Dynamic Glitching Power The glitching power is obtained by counting the temporary signal glitches in the timing simulation. The percentage of signal glitches to total glitches is used to divide the dynamic power between glitching power and useful dynamic power.

Useful Dynamic Power The useful dynamic power is obtained by tabulating the "useful" transitions within the module. If the final value of a signal is different from the beginning of a clock cycle to the end, then a useful transition is the *last* transition that occurs, and all others are glitches. Otherwise, all transitions during the clock cycle were glitches.

This study will show that the dynamic glitching power and the remaining dynamic power make up the majority of the total power for arithmetic circuits such as multipliers. Reducing one or both of these parts is the key to reducing overall power consumption.

3 Multiplier Dynamic Glitch Power

This study uses multiplier designs to demonstrate the effects of glitching on total power consumption and on operation energy. A multiplier is a good design to demonstrate this due to it's large amount of net delays and varied net lengths which lead to a large number of glitches[8, 11]. The multiplier in Figure 1 shows that pipelining can be easily implemented by adding registers between multiplier stages. Additionally, a digit-serial multiplier based on the representation in Figure 1 can be easily created[7, 12]. A pipelined multiplier and a digit-serial multiplier are used in later sections to show how glitches can be reduced in order to lower the operation energy.

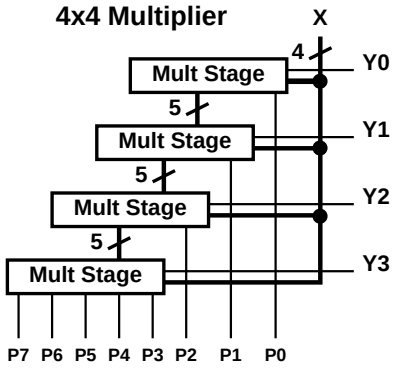

Figure 1: 4x4 multiplier.

For a non-pipelined multiplier the amount of glitching increases super-linearly as the size of the multiplier grows. To show this behaviour power estimates are performed on 4x4, 8x8, 16x16, and 32x32 multipliers. Random inputs are presented at the inputs for all simulations. The pie charts in Figure 2 show the breakdown of the total power into its constituent parts of static power, dynamic glitch power and remaining dynamic power. For all four multiplier sizes the amount of static power used by the design

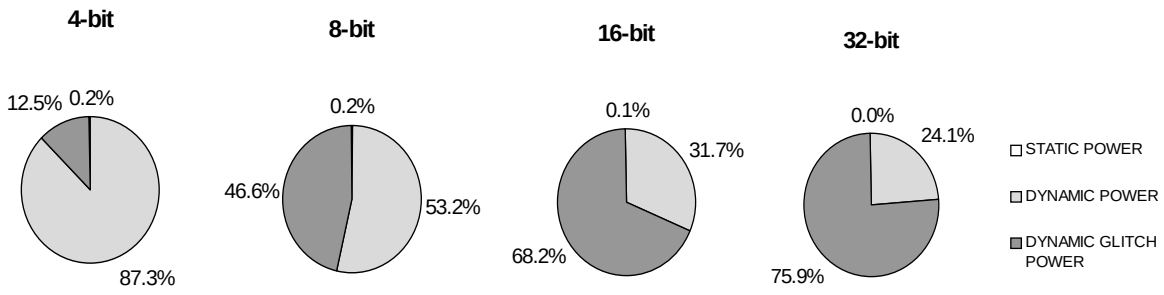

Figure 2: Breakdown of power constituents for a multiplier of various bitwidths.

makes up less than 1% of the total power. Note how the percentage of dynamic glitch power increases as the bitwidth of the multiplier grows. In the case of the 4x4 multiplier glitching accounts for about 12% of the total power. As the multiplier grows to a 32x32 multiplier, the total power is dominated by glitching which accounts for 76% of the total power.

4 Reducing Glitch Power Through Pipelining

Pipelining a design is an intuitive way to reduce glitching. A pipelined circuit has less glitching due to the reduced amount of logic between registers. With less logic between registers, the amount of interconnect between registers is also reduced. Pipelining also causes long routing interconnect to be broken up by registers resulting in a smaller range of logic and interconnect delays. Consequently less glitching occurs. In many cases pipelining can be implemented with little additional cost since often many of the flip flops within the design's CLBs go unused. When pipelining can be used to reduce glitching, the amount of power consumed by a design is also reduced.

Implementing pipelining on the multiplier design previously presented shows how fewer glitches reduces overall power consumption and operation energy. Pipeline stages are strategically inserted in the multipliers of different bitwidths (4x4, 8x8, 16x16, and 32x32). For each multiplier, pipelining is gradually introduced until the multiplier is completely pipelined.

Figure 3 shows how glitching is reduced as pipeline stages are inserted. The graph reports the number of glitches as a percentage of the total signal transitions for each multiplier. The glitching percentage drops with the amount of pipelining introduced. The almost linear quality of the graph indicates that the advantages gained from pipelining pay off right up until the multiplier is fully pipelined.

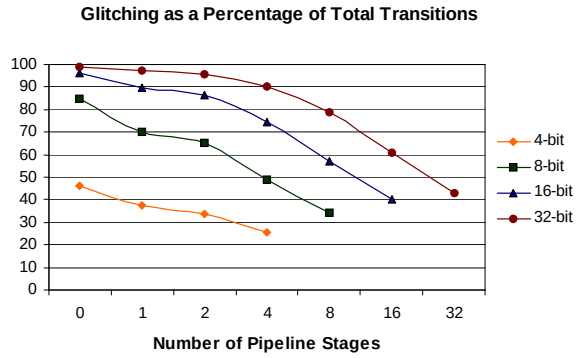

Figure 3: The amount of glitching as a percentage of total design transitions for multipliers of various widths and varying amounts of pipelining.

As the amount of glitching goes down the amount of power consumed due to glitching is also reduced. Figure 4 shows how the dynamic glitching power lowers as the amount of pipelining increases. This logarithmic graph agrees with the intuition that as the amount of glitching goes down (see Figure 3) the amount of power consumption due to glitching also goes down. The logarithmic quality of the graph indicates that as pipelining begins to be applied to the multiplier there is a large initial pay-off in reduc-

tion of power due to glitching. After a certain point there is less power savings to be had by increasing the amount of pipelining.

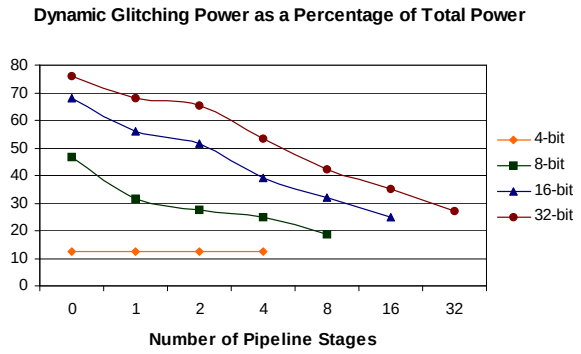

Figure 4: The amount of dynamic glitching power as a percentage of total power for multipliers of various widths and a varying amount of pipelining.

5 Reducing Glitch Power Through Digit-Serial Computation

Pipelining the stages of a multiplier has proven to be an effective way of reducing glitching (see Figure 3). The amount of pipelining available in the multiplier shown in Figure 1 is limited by number of multiplier stages. In other words an NxN multiplier can have a maximum of N pipeline stages. The graph in Figure 3 suggests that glitching could be further reduced if additional pipelining was available. Additional pipelining is available in a digit-serial multiplier where pipelining is applied at a smaller granularity[12]. A digit-serial multiplier is pipelined at the digit level. It can reduce the amount of glitching to less than 1% of total signal transitions (for operands of any width). When compared to the percentages shown in Figure 4 for the pipelined multiplier, the less than 1% glitching achieved by the digit-serial multiplier is very impressive. With almost zero glitches, the amount of power consumed by glitching in a digit-serial multiplier approaches zero. This means that at least 98% of the consumed power is due to useful dynamic power.

Figure 6 shows the power consumption of digit-serial multipliers of different radices (1, 2, 4, 8, 16, 32). When compared with the power consumption of the pipelined multipliers (Figure 5) a digit-serial multiplier with a digit size of 1 (a bit-serial multiplier) with 32-bit operands consumes the same amount of total power as a fully pipelined 4x4 multiplier or 6x less overall power than a fully pipelined 32x32 multiplier, and 160x less overall power than a non-pipelined 32x32 multiplier.

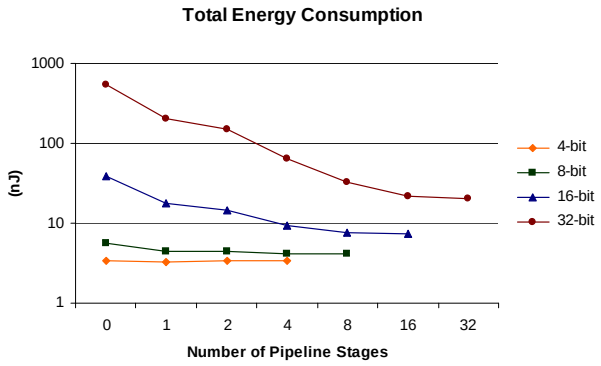

Figure 5: Total energy consumption (in nJ) of a multiplier of different widths and various amounts of pipelining.

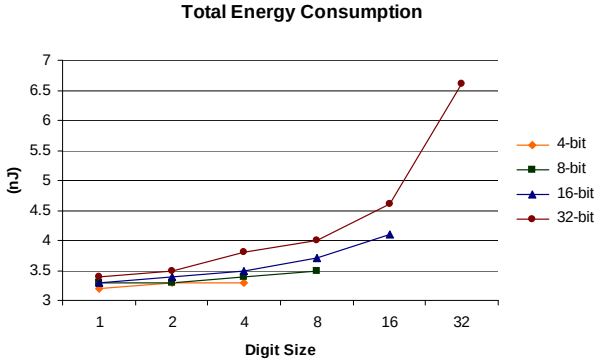

Figure 6: Total energy consumption (in nJ) of digit-serial multipliers of different widths and digit sizes.

The large power savings of the digit-serial multiplier comes at a cost. With such an extreme amount of pipelining the minimum clock period of the digit-serial multiplier is reduced allowing for a faster clock rate, but the pipelining not only increases the latency but also the throughput of the design. Whereas the throughput of an NxN multiplier based on the design

of Figure 1 is one product per cycle, the throughput of a digit-serial multiplier is one product per N/D cycles (where D is the digit size). Note that for traditional digit-serial multipliers one product is retreived once per $N*2/D$ cycles, but for this study we use an efficient digit-serial multiplier that produces a product in half as many cycles[7]. New operands are introduced to a digit-serial multiplier every N/D cycles. Since the throughput of a design directly affects operation energy, even though a digit-serial multiplier consumes less overall power it may have a larger operation energy than a pipelined multiplier.

6 Operation Energy

Most studies quantifying the power consumption of FPGA circuits report on the average power of the overall circuit. For high-throughput arithmetic circuits, however, the primary concern is the amount of *energy* required to perform a specific arithmetic operation. Understanding the energy consumption of individual datapath operators will allow the designer or synthesis tool to choose operators that meet a specific throughput constraint while minimizing energy consumption.

This study will use the glitch estimation technique presented in the previous section to estimate the amount of energy consumed by a number of multiplier operators. Four energy parameters will be used: energy per operation, energy delay, energy throughput, and energy density.

6.1 Energy per Operation

Energy per operation quantifies the amount of energy required to complete a single operation of a specific circuit operator. This measure, reported as nJ, is more useful than an average power measure when comparing arithmetic circuits with different implementation approaches. Specifically, this measure allows us to compare the energy efficiency of multi-cycle operators with single-cycle operators.

Circuit energy can be computed by integrating the circuit power consumption over time ($E = \int Pdt$). Assuming constant power consumption, the energy can be estimated by multiplying the average power of the circuit operator by the amount of time to complete a single operation. For single-cycle operators, the energy per operation is simply the average power of the circuit times the clock period, or $E_{cycle} = P \cdot t_{clk}$.

The pipelined multiplier is an example of a circuit with a single-cycle operation. The energy per operation for the pipelined multipliers is shown in Figure 7. Energy per operation is affected by glitching power. Just as the graph of power consumption in Figure 4 drops logarithmically, Figure 7 shows that energy per operation also drops logarithmically. The graph also shows that as the width of the multiplier increases, energy per operation grows exponentially.

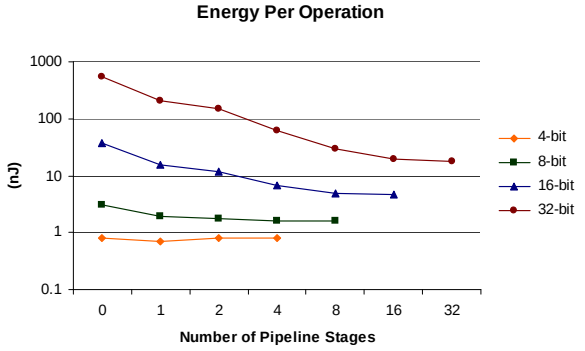

Figure 7: The energy per operation (in nJ) for a multiplier of different widths and various amounts of pipelining.

The energy per operation of multi-cycle circuit operators must take into consideration the number of clock cycles required to perform the operation. For a non-pipelined multi-cycle operator, the energy per operation is $E_{op} = P \cdot t_{clk} \cdot n$ where n is the number of clock cycles required to perform the operation.

Pipelined operators overlap the computation of several discrete operations in a single clock cycle. The energy consumed by a pipelined operator is shared among the various instances of the operator in the pipeline. Thus for pipelined operators, rather than considering the number of cycles required to perform the operation (n), the interval between successive initiations of the pipeline operation is considered (δ). The energy for a single computation is computed as $E_{op} = P \cdot t_{clk} \cdot \delta$. For $\delta = 1$ (which is the case for the pipelined multipliers), the operator is fully pipelined (i.e. a new operation can be initiated every clock cycle) and the energy per operation

is $E_{op} = P \cdot t_{clk}$.

The digit-serial multiplier is an example of a circuit with multi-cycle operation. The digit-serial multiplier does not have a $\delta = 1$ like the pipelined multipliers, but has $\delta = N/D$; where N is the bitwidth of the digit-serial multiplier operands, and D is the digit size. Thus, despite its much lower overall power consumption (Figure 5) in general the digit-serial multiplier has a large energy per operation.

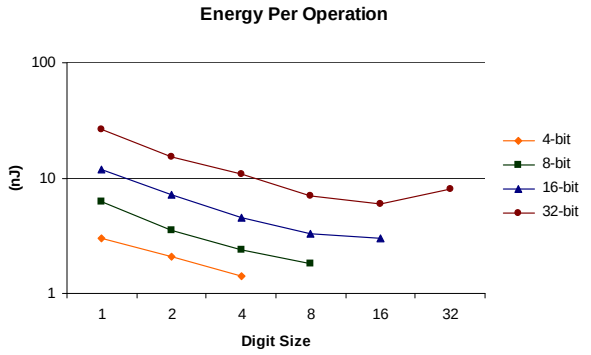

Figure 8: Energy per operation (in nJ) of a digit-multiplier with different digit sizes and operands of different widths.

The energy per operation of the digit-serial multipliers is shown in Figure 8. The plots show that the energy per operation of the digit-serial multiplier is always greater than the pipelined multiplier except for one case of the 4x4 pipelined multiplier (see Figure 7). The plots in Figure 8 show that increasing the digit size of the multiplier lowers the energy per operation up until a digit size of about 16 bits.

6.2 Energy Delay

Energy delay is a related measure that combines the energy efficiency and speed of an operator into a single parameter [13]. This measure is frequently used to balance the trade-off between reducing energy and increasing circuit speed. The energy delay of a circuit operator is computed as $E_{delay} = E_{cycle} \cdot \lambda$, where λ is the latency of a single computation of the operator. The latency of the operation is $\lambda = t_{min} \cdot n$ where n is the number of clock cycles required to complete a single operation and t_{min} is the minimum clock time period of the circuit. Recall that a pipelined multiplier will have n =number of pipeline stages, and a digit serial mulitiplier will have $n = N/D$ (where N is the operand widths and D is the digit size).

The minimum clock time period of a circuit decreases as the amount of pipelining increases. This smaller clock period provides greater throughput for designs with single-cycle operations such as the pipelined multipliers. However, as we increase the amount of pipelining the latency increases (n increases), potentially increasing the energy delay. Thus, as pipelining is increased the energy delay could increase due to an increased latency, but the energy delay could also be reduced due the the reduced amount of glitching.

Figure 9 shows how an increased amount of pipelining could either decrease or increase energy delay. In designs which have a large percentage of glitching (see Figure 2) such as the 32-bit multiplier, increasing the amount of pipelining decreases the energy delay. In designs with less glitching, energy delay rises as pipelining is increased.

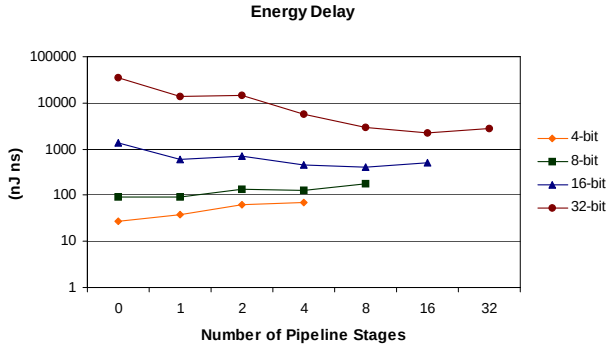

Figure 9: Energy delay (in nJ·ns) for a multiplier of different widths and various amounts of pipelining.

The trade-off of larger latency for less glitching is accented with the digit-serial multiplier. The digit-serial multiplier consumes almost no energy due to glitching, but relative to the pipelined multipliers requires a large number of cycles to complete an operation. The plots in Figure 10 shows the energy density of digit-serial multipliers. Due to the large amount of glitching in the pipelined multipliers, the digit-serial multipliers have a lower energy delay. For the digit-serial mulipliers the energy delay decreases with increasing digit size only up to a certain point (similar to energy per operation). The optimal digit size for energy delay appears to be 8 bits.

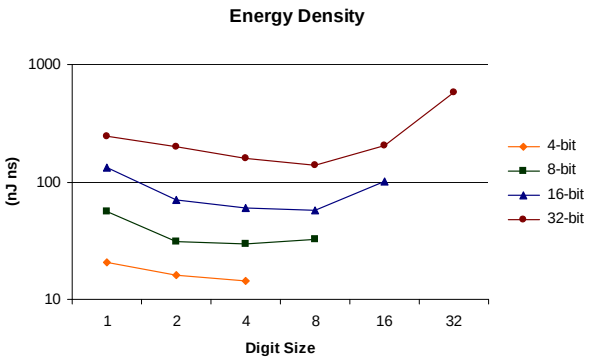

Figure 10: Energy delay (in nJ·ns) of digit-serial multipliers of different digit sizes and operands of different widths.

6.3 Energy Throughput

The energy throughput parameter is similar to energy delay in that it combines energy efficiency and circuit speed into a single parameter. Energy throughput differs from energy delay in that it does not take into account the number of clock cycles needed to complete a single operation (n), but rather it considers the number of clock cycles between successive operation initiations (δ). In most cases only one or the other is used. Energy throughput is calculated as $E_{thput} = E_{cycle} \cdot t_{min} \cdot \delta$, where t_{min} is the minimum clock time period of the circuit and δ is the interval between successive initiations of the pipeline operation. For pipelined multipliers $\delta = 1$ and for digit-serial multipliers $\delta = N/D$.

Sometimes it is more useful to consider throughput energy rather than energy delay. In some designs a pipelined operation is used successively in such a way that a new operation can be initiated on each clock cycle ($\delta = 1$). In these cases, energy throughput might be a more desirable parameter than energy delay. Energy delay is more useful for designs which require an operation to complete before a new one can be initiated ($\delta = pipedepth$).

Where increasing pipeline depth raises energy delay (due to an increased latency), it lowers energy throughput. Pipelining benefits energy throughput in two ways: it reduces glitches which lowers overall energy consumption, and it reduces the minimum clock period - providing a greater throughput. Figure 11 shows how pipelining greatly benefits energy throughput. Energy throughput (in nJ·ns) is displayed on a logarithmic scale therefore over a range of different pipeline depths it appears linear. The graph shows how critical the first stages of pipelining are. Implementing only a single pipeline stage in the 32x32 multiplier reduces the energy throughput by an order of magnitude.

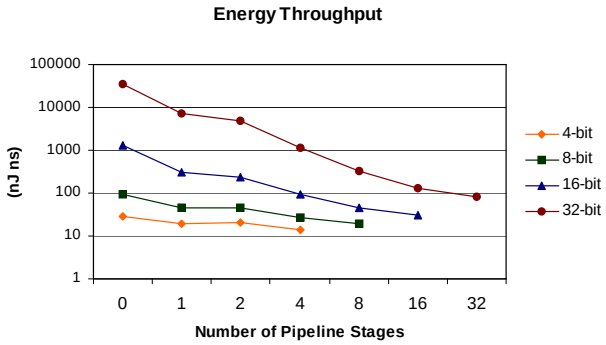

Figure 11: Energy throughput (in nJ·ns) for a multiplier of different widths and various amounts of pipelining.

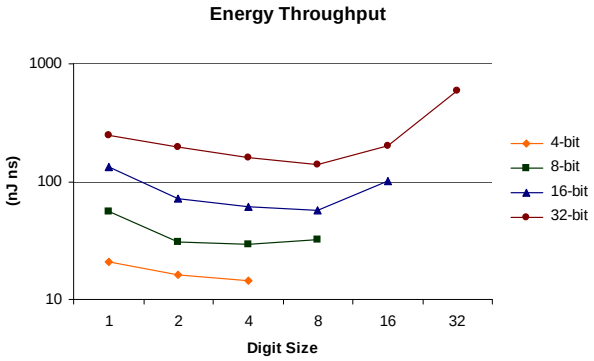

Figure 12: Energy throughput (in nJ·ns) of digit-serial multipliers of different digit sizes and operands of different widths.

Energy throughput for the digit-serial multiplier is the same as energy density since $n = \delta = N/D$ (where N is the size of the digit-serial multiplier operands and D is the digit size). In other words, the data introduction rate is the same as throughput. Despite its low power consumption (see Figure 6) the digit-serial multiplier tends to have a large energy throughput and energy delay. The digit-serial

multiplier benefits from having almost no glitching and a smaller clock period than the pipelined multipliers. However, it suffers from having a large δ value. In almost every case, this large δ value overshadows the advantages of a low minimum clock period and minimized glitching. Figure 12 shows the energy throughput for the digit serial multipliers. As with energy delay, the optimal digit size appears to be 8 bits. Comparing this figure with Figure 11 reveals that as the pipeline depth of the multipliers increases, they overtake the digit-serial multipliers in efficiency.

6.4 Energy Density

The final energy parameter used in this study is energy density, $E_{density}$. This parameter normalizes the amount of energy required to perform a single operation to the number of logic resources used by the circuit. This parameter can be calculated as $E_{density} = E_{cycle}/A$, where A is the "area" of the circuit. It is important to note that the amount of dynamic energy consumed by a circuit module is not necessarily linearly related to the size of the circuit. As described earlier, the primary contributer to dynamic power is signal glitches. Since circuit size affects energy density, large circuits that reduce glitching through pipelining may have a lower energy density than smaller pipelined circuits.

Figure 13 shows the energy density for the pipelined multipliers. From the graphs in this figure it may appear that energy density is not directly related to circuit size, but multiplier size does impact energy density through average net activity rate. The 4x4 multiplier has the smallest area but the largest energy density. The 32-bit multiplier has very high energy density when it is non-pipelined, however as pipelining is introduced the energy density drops until it reports the lowest energy density.

The energy trends shown in Figure 13 are the result of glitching and average net activity rate. As expected, glitching is reduced as pipelining is introduced, lowering the overall energy consumption. The 32-bit multiplier demonstrates this trend. However, the size of the multiplier also affects the energy density. The nets of a smaller multiplier have a higher average activity rate than a larger multiplier, resulting in larger energy densities for smaller multipliers.

The graph in Figure 14 shows the energy density

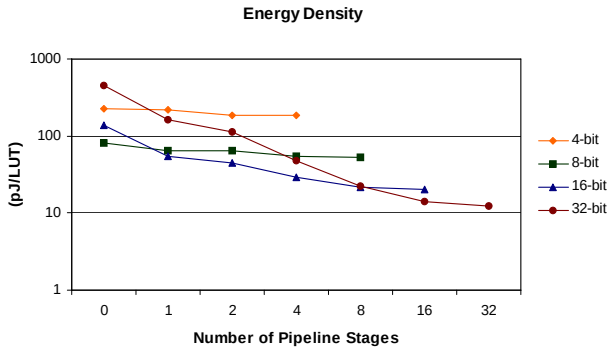

Figure 13: Energy density (in pJ/LUT) for a multiplier of different widths and various amounts of pipelining.

of the digit-serial multipliers. The plots in this figure show that despite being a circuit with a multi-cycle operation, the digit-serial energy density is often lower than the energy density of the pipelined multipliers. The digit-serial multipliers generally have a smaller area than the pipelined multipliers. Whereas the area of the pipelined multipliers has no direct correlation with its energy density, the energy density of the digit-serial multipliers is inversely proportional to area. The digit-serial multipliers with 4-bit operands have the smallest area and the highest energy density. Conversely, the digit-serial multipliers with 32-bit operands have the largest area and the smallest energy density.

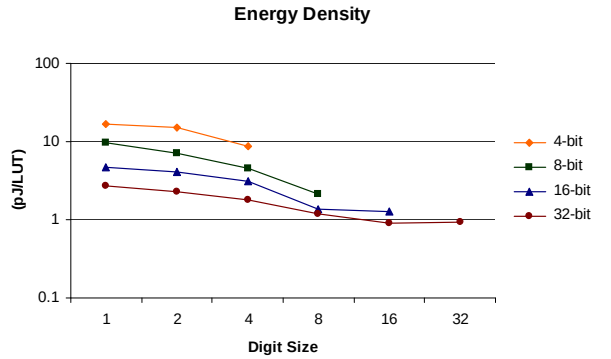

Figure 14: Energy density (in pJ/LUT) of digit-serial multipliers of different digit sizes and operands of different widths.

7 Clock Energy

Often the additional registers required for pipelining can come at little or no cost in FPGA designs, since the configurable logic blocks (CLBs) used in the design have registers included that often go unused. When an ASIC design is pipelined, the power savings due to glitch reduction has to be weighted against the power increase due to the increase in register count. Unlike an ASIC design however, in an FPGA design there is little or no increase in clock power as pipeline depths or digit sizes are increased. Figures 15 and 16 show the amount of clock energy required for the pipelined multipliers and digit-serial multipliers respectively. For the digit-serial multipliers, the clock energy consumed is constant, regardless of the operand bit width or digit size. The amount of power consumed by the clock is a small fraction of the total power consumption.

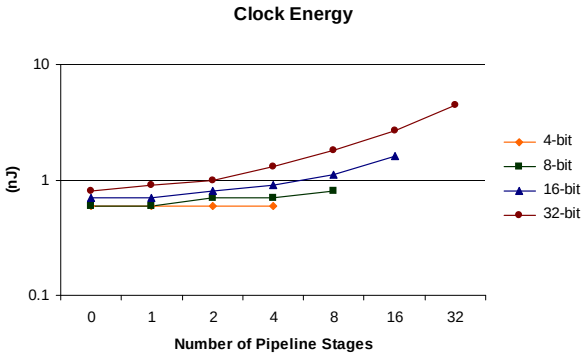

Figure 15: Clock energy consumption in pipeplined multipliers

8 Conclusion

This paper presents a methodology for estimating the glitches of FPGA circuits and uses this methodology to determine the amount of power wasted in glitching. This paper shows that the majority of dynamic power in non-pipelined multipliers is consumed by glitches. The glitch power can be reduced by pipelining the circuit. Results in the paper show an exponential decrease in glitch power by increasing the pipelining depth.

Several energy related measures were introduced

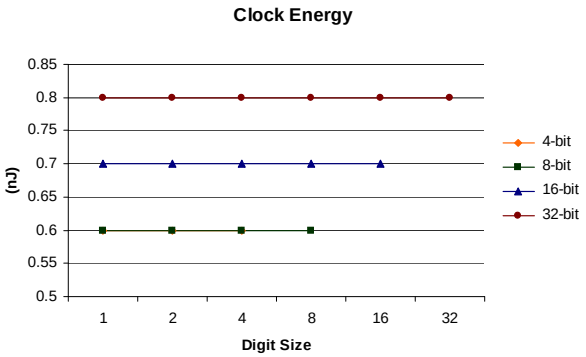

Figure 16: Clock energy consumption in digit-serial multipliers

and used to provide better tools for comparing the energy requirements of different implementation approaches. Single cycle and multi-cycle circuit implementations can be compared by estimating the energy per operation rather than the average power consumption. This paper demonstrates that while multi-cycle digit-serial operators have significantly lower average power, the energy per operation is higher than deeply pipelined parallel approaches.

Energy delay, energy throughput and energy density were also used to compare the multiplier implementation approaches. Energy delay is used to balance the trade-off between energy reduction and operator latency. Increasing the pipelining depth exponentially reduces glitch power, but also increases latency. The amount of pipelining which produces the best trade-off depends on the bitwidth of the multiplier. In general, as the bitwidth increases, the amount of pipelining producing the lowest energy delay also increases.

The metrics and estimates generated in this paper will be used as part of a high-level datapath synthesis tool. This tool will use energy estimates of the various multiplier implementation approaches to select the proper multiplication approach. High-power low-pipelined circuits will be selected for datapath circuits with tight latency constraints while low-power, highly pipelined circuits will be selected for latency tolerant datapath circuits.

References

[1] Nathaniel Rollins and Michael Wirthlin. Reducing energy in FPGA multipliers through glitch reduction. *BYU Dspace - https://dspace.byu.edu/handle/1877/61*, August 2005.

[2] Verghese George, Hui Zhang, and Jan Rabaey. The design of a low energy FPGA. In *International Symposium on Low Power Electronics and Design 1999. Proceedings*, pages 188–193, August 1999.

[3] Juergen Becker, Michael Huebner, and Michael Ullmann. Power estimation and power measurement of Xilinx Virtex FPGAs: Trade-offs and limitations. In *Proceedings of the 16th Symposium on Integrated Circuits and Systems Design (SBCCI'03)*. IEEE Computer Society Press, 2003.

[4] Nathan Rollins. SLAAC1V power user's manual. Technical report, Department of Electrical and Computer Engineering, Brigham Young University, 2004.

[5] Anand Raghunathan, Sujit Dey, and Niraj K. Jha. Register transfer level power optimization with emphasis on glitch analysis and reduction. *IEEE Transactions on Computer-aided Design of Intergerated Circuits and Systems*, 18(8):1114–1131, August 1999.

[6] Steven J.E. Wilton, Su-Shin Ang, and Wayne Luk. The impact of pipelining on energy per operation in field-programmable gate arrays. In *Field-Programmable Logic and Applications. Proceedings of the 13th International Workshop, FPL 2004*, Lecture Notes in Computer Science, LNCS 3203, pages 719–728. Springer-Verlag, August 2004.

[7] Yun-Nan Chang, Janardhan H. Satyanarayana, and Keshab K. Parhi. Systematic design of high-speed and low-power digit-serial multipliers. 45(12):1585–1596, December 1998.

[8] Eric Kusse and Jan Rabaey. Low-energy embedded FPGA structures. In *International Symposium on Low Power Electronics and Design 1998*, pages 155–160, August 1998.

[9] Jeroen Leijten, Jef van Meerbergen, and Jochen Jess. Analysis and reduction of glitches in synchronous networks. In *European Design and Test Conference, 1995.ED&TC Proceedings*, pages 398–403, March 1995.

[10] Xilinx, Inc. *XPower Manual*.

[11] Suleman Sirri Demirsy, Andrew G. Dempster, and Izzet Kale. Power analysis of multiplier blocks. In *International Symposium on Circuits and Systems*, volume 1, pages I–297–I–300, May 2002.

[12] R. F. Lyon. Two's complement pipeline multipliers. *IEEE Transactions on Communications*, pages 418–425, April 1976.

[13] Ricardo Gonzalez and Mark Horowitz. Energy dissapation in general purpose microprocessors. *IEEE Journal of Solid-State Circuits*, 31(9):1277–1284, September 1996.