

Utilizing Clock-Gating Efficiency to Reduce Power

Mitch Dale, Calypto Design Systems

1/15/2008 8:48 AM EST

Clock gating can be applied at multiple levels of abstraction, but RTL is the most effective point in the design process. Measuring clock-gating efficiency is an accurate guide to power optimization because it takes into account switching activity. RTL designers can use clock-gating efficiency to pinpoint hotspots and concentrate their optimization efforts. As consumers continue to demand more functionality in smaller, more energy efficient devices, power optimization rules a hardware designer's life. It typically takes multiple iterations over weeks of optimization to achieve power goals and budgets. While power should be optimized at all stages of the design flow, many times it is only addressed after initial register transfer level (RTL) synthesis runs uncover inefficiencies.

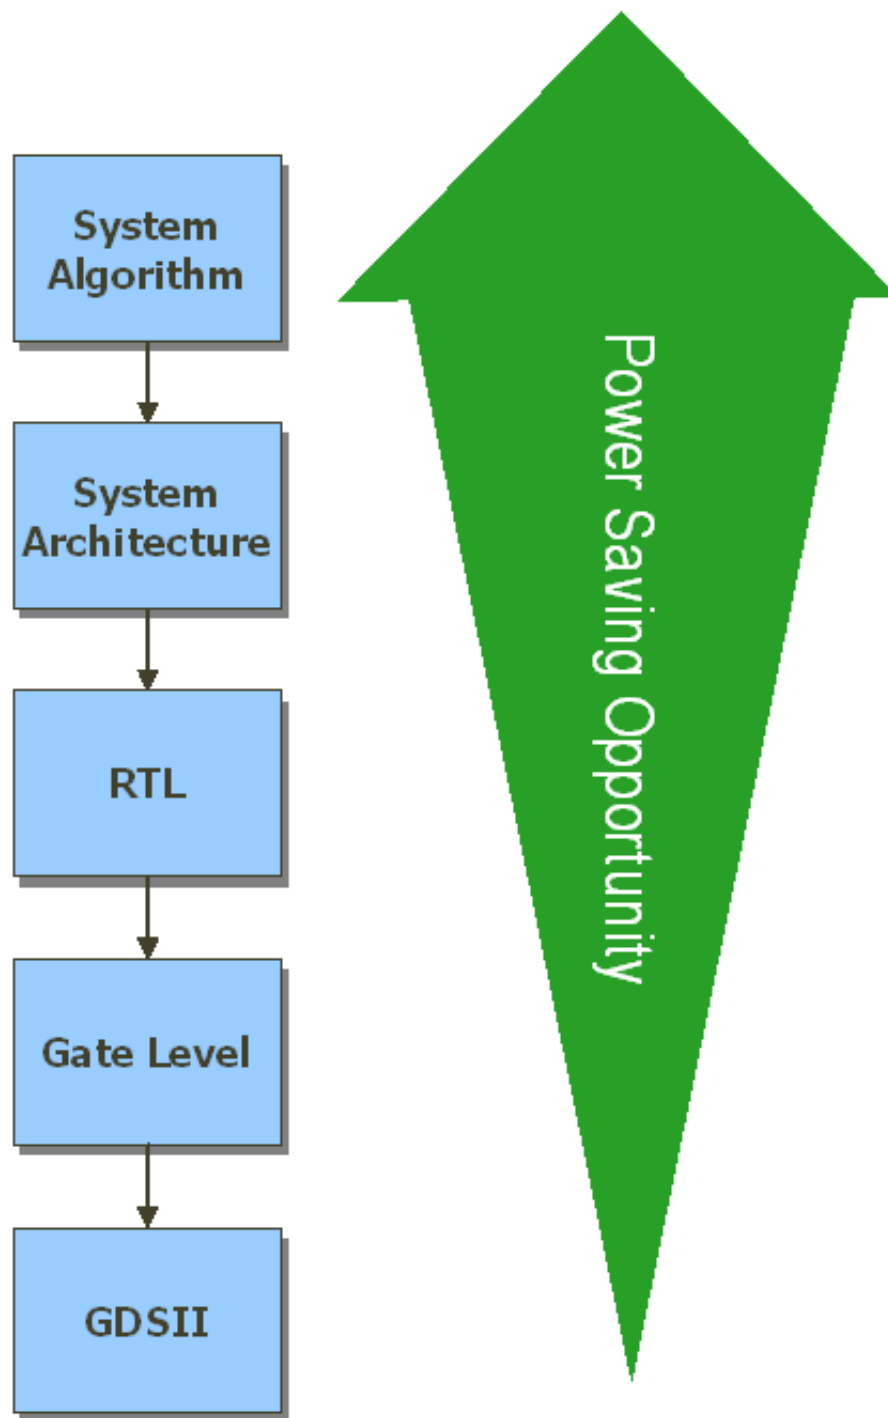
RTL clock gating is the most commonly used optimization technique for improving energy efficiency, but leads to the question of how well a design is clock gated. The traditional method of looking at the percentage of registers clock gated is not indicative of the energy efficiency because it does not take into account switching activity. The average clock-gating efficiency is a much better indicator of energy consumption because it measures of both the number of registers gated and the duration they are turned off.

Dynamic Power Optimization at Multiple Design Stages

A design's energy consumption is a function of the switching activity, which in turn is totally dependent on the system application and hardware implementation. Designers typically have little control over the application; video must be compressed at a given rate, packets routed within a given latency or instructions executed at a certain frequency. In contrast, there are multiple design techniques and tricks the designer has at his or her disposal when implementing the hardware. Clock gating is an accepted design technique for optimizing power, and can be applied at the system level, RTL and gate-level. The granularity of clock gating and the impact it has on overall energy consumption depends on the design stage.

Although each design stage offers an opportunity to save power, higher levels of abstraction have greater impact on reducing power and lowering costs. Starting at the system-level, a CPU may have multiple sleep modes, each disabling specific blocks in the design. Moving down in abstraction, RTL clock gating shuts off unused computations while leaving other logic active for data processing. For example, turning off unused computations such as shift, multiple or add within an Arithmetic Logic Unit (ALU) based on the operator.

Gate-level clock gating can provide a finer grain of control by not clocking a register if it is not changing states. As in the case of an output hold, the clock can be gated with the hold condition so that the register is not unnecessarily toggling. *Figure 1* shows potential opportunities for saving power at different stages of design.



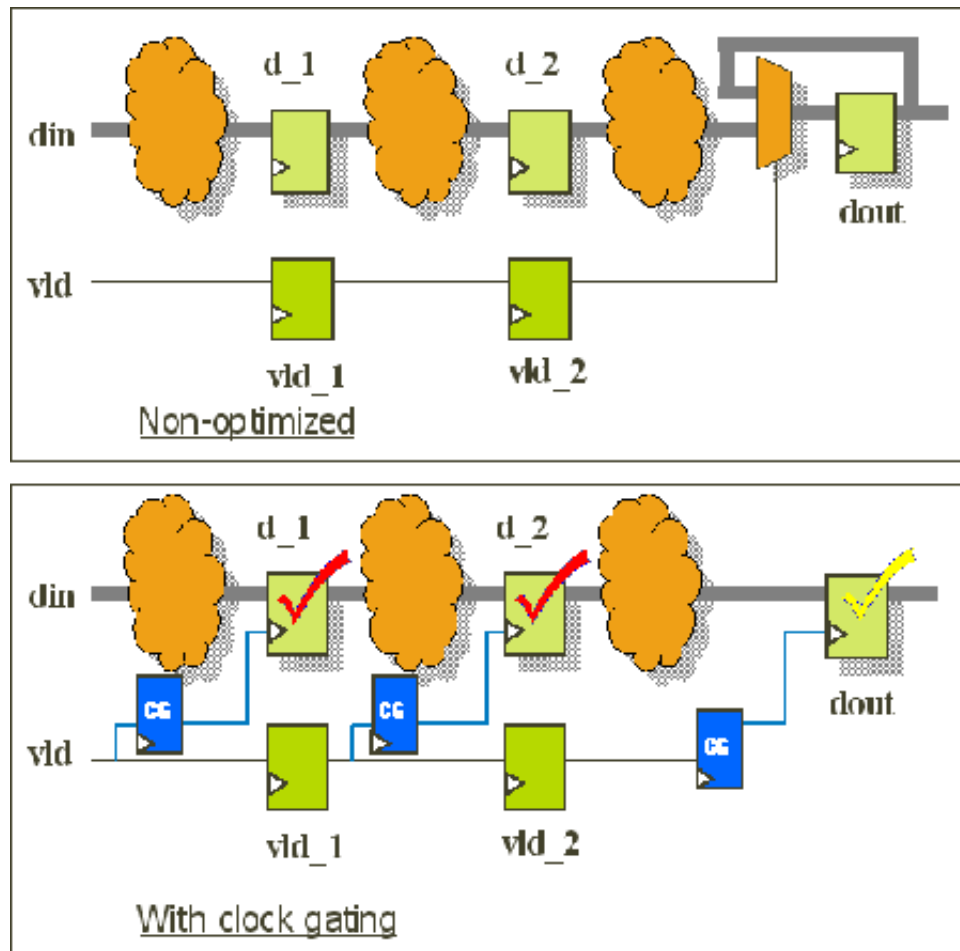
1. Opportunities abound for saving power at different stages of design.

Decomposing Clock Gating

Clocks consume power because they continuously toggle the registers and associated downstream logic. This is referred to as toggle rate and is a major contributor to dynamic or switching power. To reduce dynamic power consumption, clock gating turns off clocks while still maintaining the original design functionality.

Clock gating is a two-step process. The first step is identifying enable conditions, simple combinational logic, such as an output hold on a register or more involved sequential logic that spans multiple clocks. Combinational enables can be identified by today's RTL synthesis tools as long as the RTL code fits a recognized coding style. Sequential enables are more difficult to recognize and typically are done manually by hardware designers. The second step in clock gating involves inserting clock-gating cells into the clock path using the enable logic.

Commercially available synthesis tools accomplish the second task automatically. *Figure 2* shows that combinational and sequential clock gating.



2. Red checks identify sequential analysis and yellow check combinational analysis blocks.

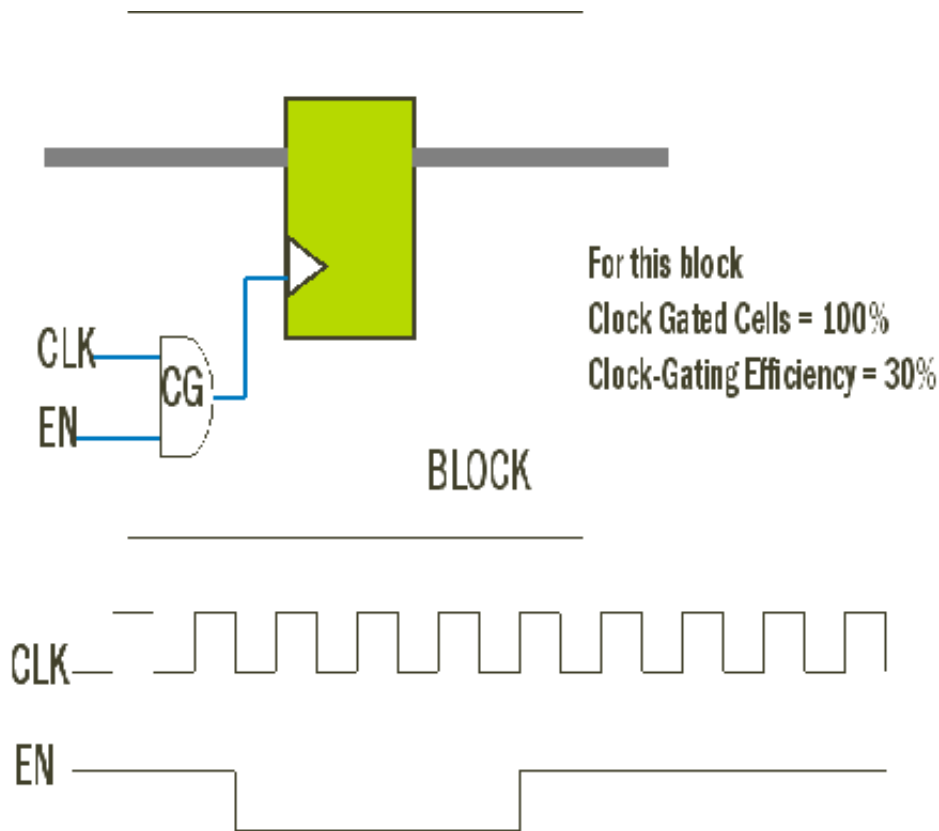
Sequential clock gating has a greater impact on energy efficiency than combinational clock gating because it turns off registers for longer periods of time. In fact, sequential clock gating has been shown to reduce power by up to 60% on design blocks. Sequential clock gating requires sequential analysis based on activity over multiple clock cycles to decide which registers can be gated and under what enable conditions.

What is Clock-Gating Efficiency?

RTL is the best point in the design process to optimize power. At this point in the design flow, there is flexibility in the implementation to make significant improvements in energy efficiency. There is accurate information available from synthesis to reflect the total impact on power, timing and area, as well. What's needed is a good RTL metric to evaluate how well a design is clock gated and to help identify candidate clock-gating optimizations within the design.

A typical metric used to measure the effectiveness of clock gating is the percentage of registers in the design that are clock gated. While this gives designers an indication of the number of clock-gated registers in the design, it has poor correlation to actual power savings. That's because dynamic power consumption depends on the toggle rate. Clock-gating efficiency, on the other hand, considers the toggle rate, making it a more telling indicator of actual dynamic power consumption.

Clock-gating efficiency is defined as the percentage of time a register is gated for a given stimulus or switching activity. The average clock-gating efficiency can be computed as the average of all clock-gating efficiencies in the design. Clearly, clock-gating efficiency depends on representative switching activity. A design may have multiple modes and multiple operating conditions. For example, designers will typically have a switching activity file based on idle, typical and peak modes. Because switching activity is only as representative as the testbench itself, the selection of a representative testbench is critical to good power estimation.



3. Clock gating efficiency takes switching activity into account.

Figure 3 shows a simple clock gated register. Since there is only one register, the block is 100% clock gated. However, when looking at the enable signal over time, the clock is inactive for only three of the 10 cycles making clock-gating efficiency 30%. By calculating the average of clock-gating efficiency for all registers over a given set of stimulus, a design team would have a better idea of how well a design is clock gated.

A designer's goal is to improve the average clock-gating efficiency as much as possible, and it is not practical to achieve 100%. The optimal clock-gating efficiency is design and application dependent. Moreover, clock gating is not an always-good power optimization as the added enable logic and clock-gating cells have associated power, timing and power costs. The simple rule thumb is to look for clock-gating conditions that disable wide registers over multiple cycles.

Using Clock-Gating Efficiency to Guide Power Optimization

Adding clock gating may not always be accompanied by reduced power because dynamic power is also a function of clock frequency, voltage and capacitance. Even though it is not an absolute indication of power, it is a good metric for hardware designers to gain visibility into energy efficiency at RTL without time-consuming power analysis or synthesis.

Clock-gating efficiency is a solid metric to guide hardware designers toward an energy efficient implementation. Registers with low clock-gating efficiency are good candidates for clock gating. Considerations when using this metric include:

- Clock gating adds logic that consumes power. Low efficiency together with knowledge of the design will point to areas where greater power savings maybe possible. A good example of this is low-efficiency datapath registers.
- The intrinsic limit for improving clock-gating efficiency depends on the functionality of a block. For example, a high-speed Ethernet IPv6 interface block may have little opportunity for improvement. An efficiency of 10-15% may be near optimal under typical traffic scenarios.
- Power savings from clock gating is increased with the size of the logic following the clock-gated register. The greater the fan-out, the better.
- Ranking design registers by clock-gating efficiency. Optimize registers where the duration of shutting off the clock is multiple cycles.
- Pay attention to the impact on timing. Avoid using timing-critical signals in enable conditions.

Conclusion

Hardware designers commonly use clock gating to reduce toggle rates on registers, lowering dynamic power consumption. Clock gating can be applied at multiple levels of abstraction, but RTL is the most effective point in the design process. Measuring clock-gating efficiency is an accurate guide to power optimization because it takes into account switching activity. RTL designers can use clock-gating efficiency to pinpoint hotspots and concentrate their optimization efforts.

About the Author:

Mitch Dale is Director of Product Marketing with Calypto Design Systems. He can be reached at mdale@calypto.com.