



WP370 (v1.4) August 29, 2013

Reducing Switching Power with Intelligent Clock Gating

By: Frederic Rivoallon and Jagadeesh Balasubramanian

Xilinx delivers the first automated, fine-grain clock-gating solution that can reduce dynamic power by up to 30% in 7 series and Zynq®-7000 devices.


Xilinx intelligent clock-gating optimizations are automatically performed on the entire design, introduce no new tools to the flow, and generate no changes to the existing logic or to the clocks that alter the behavior of the design. And, in most cases, the timing is also preserved.

Intelligent Clock Gating Overview

Clock gating is a well understood power optimization technique employed in both ASIC and FPGA designs to eliminate unnecessary switching activity. This method usually requires the designers to add a small amount of logic to their RTL code to disable or deselect unnecessarily active sequential elements—registers, for example. Despite the obvious value of reduced dynamic power afforded by this method, the designer faces significant challenges when attempting to perform these optimizations manually:

- Truly reducing activity in the design requires intimate knowledge of the design itself and typically requires numerous changes to the RTL.
- Most ASIC and FPGA designs today comprise some combination of new, legacy, and third-party IP circuit designs, but typically only the new designs are candidates for clock-gating optimizations. Designers rarely if ever attempt these optimizations on legacy and IP design. They usually do not have sufficient depth of knowledge about the design and operation of the legacy RTL code, and it requires too much time to manually develop meaningful clock-gating optimizations.
- Applying clock-gating optimizations usually requires the addition of more tools and more steps to the design flow and can precipitate the creation of an intricate set of new clocks requiring complex timing analyses (as is often the case for ASIC optimization). Unless the gains in power efficiency are sufficient and essential to the success of the design, the additional complexity and time can be prohibitive and add risk.

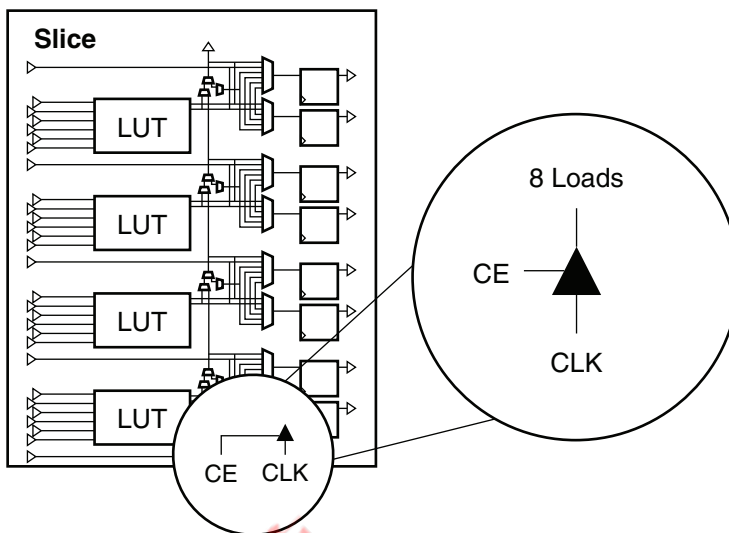
Xilinx has an automated capability linked to the place and route portion of the standard design flow that uses a set of innovative algorithms to perform an analysis on all portions of the design (including legacy and third-party IP blocks). Having analyzed the logic equations to detect sourcing registers that do not contribute to the result for each clock cycle, the software utilizes the abundant supply of clock enables (CEs) available in the logic to create fine-grain clock-gating or logic-gating signals that neutralize superfluous switching activity, as shown in [Figure 1](#).



WP370_01_062610

Figure 1: Intelligent Clock Gating Dramatically Reduces Switching Power Consumption

Each CE is ideally suited for power optimization because it connects to the basic logic cluster (the slice). The CE controls a small number of registers (only eight), providing the level of granularity that matches the minimum size of buses used by the vast majority of designs (see Figure 2).




WP370_02_071213

Figure 2: Clock Enables in the Slice

It is important to note that these optimizations do *not* alter the pre-existing logic or clock placement, nor do they create new clocks. The resulting design is logically equivalent to the original and the additional logic created is separate from previous logic, adding only 2% more LUTs (on average) to the original design. As a result, the optimization does not affect timing in the vast majority of cases because it does not add levels of logic to the original logic paths.

Additional Optimizations

Intelligent clock gating optimization also reduces power for dedicated block RAM in either simple or dual-port mode. These blocks provide several enables: an array enable, a write enable, and an output register clock enable. Most of the power savings comes from using the array enable, as shown in Figure 3.



WP370_04_062610

Figure 3: Intelligent Clock Gating Optimizations Can Leverage Block RAM Enables

For example, in a block RAM followed by a 2-to-1 multiplexer, the optimization implements an OR function in a LUT with the write enable (weR) and the select

(preselectR) and connects them to the ENARDEN of the block RAM. The OR function ensures that the block dissipates less power when no data is being written and when its output is not used (i.e., not selected in the multiplexer). Assuming a 50% toggle rate on the write enable of the block RAM, this optimization shows a 26% reduction in dynamic power. Here is an example of the Verilog code:

```

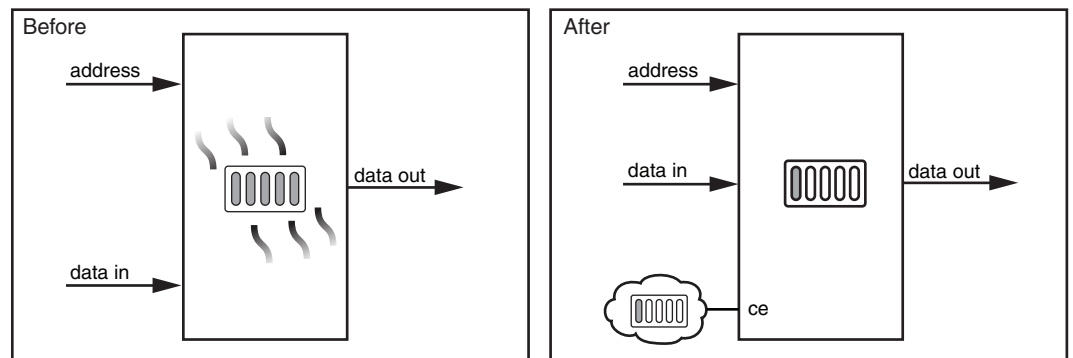
module ram_mux ( input      select, clk, we,
                input [7:0] bypass,
                input [10:0] addr,
                input [7:0] data_in,
                output reg [7:0] result_out );

    reg          preselectR, selectR, weR;
    reg [7:0]    data_out, mem [2047:0];

    always @(posedge clk)
    begin
        // RAM block 2048x8 (inferred)
        if (weR) mem[addr] = data_in;
        data_out <= mem[addr];
        // Registering inputs
        weR <= we;
        preselectR <= select;
        selectR <= preselectR;
        // Mux: RAM output and input data
        result_out <= selectR ? data_out : bypass;
    end
endmodule // ram_mux

```

In that same version of software, these optimizations can detect a clock enable implemented as logic and replace it with a dedicated CE, as shown in [Figure 4](#).



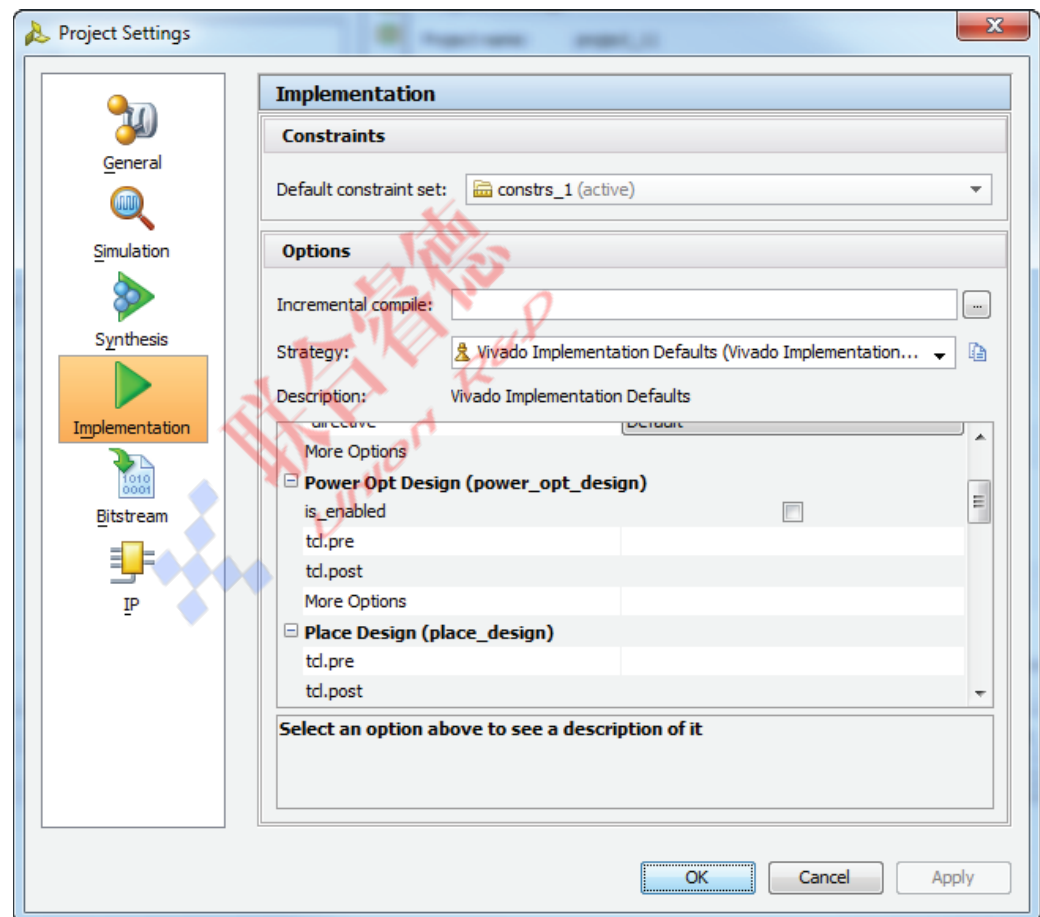
WP370_04_062610

Figure 4: Rewiring a Clock Enable as a Native CE to Reduce Power

Intelligent Clock Gating in the Design Implementation Flow

Vivado® Design Suite (v2013.1 and later) integrates the intelligent clock-gating optimization into the design implementation flow. All black boxes and IPs are also optimized. A subset of the block RAM power optimizations is performed by default in the `opt_design` stage after synthesis. These optimizations focus on minimizing the power consumption of block RAMs in the design while preserving performance.

A more extensive power optimization that includes additional block RAM power optimizations as well as registers and shift register (SRL) optimizations is available (optionally) during the `power_opt_design` step after the `opt_design` stage. It can be enabled through the `is_enabled` option under Power Opt Design in the Project Settings dialog. After the power optimizations, the user runs placement and routing as usual. See [Figure 5](#).



WP370_05_080213

Figure 5: Turning on the Power Optimization Option in Vivado Design Suite

By default, power optimizations work on the entire design. However, special constraints exist to perform a more targeted optimization on portions of the design such as modules, clock domain, or type of instances such as block RAMs.

Benchmark Results

The benefits of intelligent clock gating vary depending on the design; some designs show significant power savings while others do not. The effectiveness depends on whether the user has already optimized the design for power and if the design had logic structures that are amenable to clock gating. Customer design suites have shown dynamic power reduction of up to 30% and an average of 18%. See Figure 6 for a design-by-design breakdown on a suite of 50 designs.

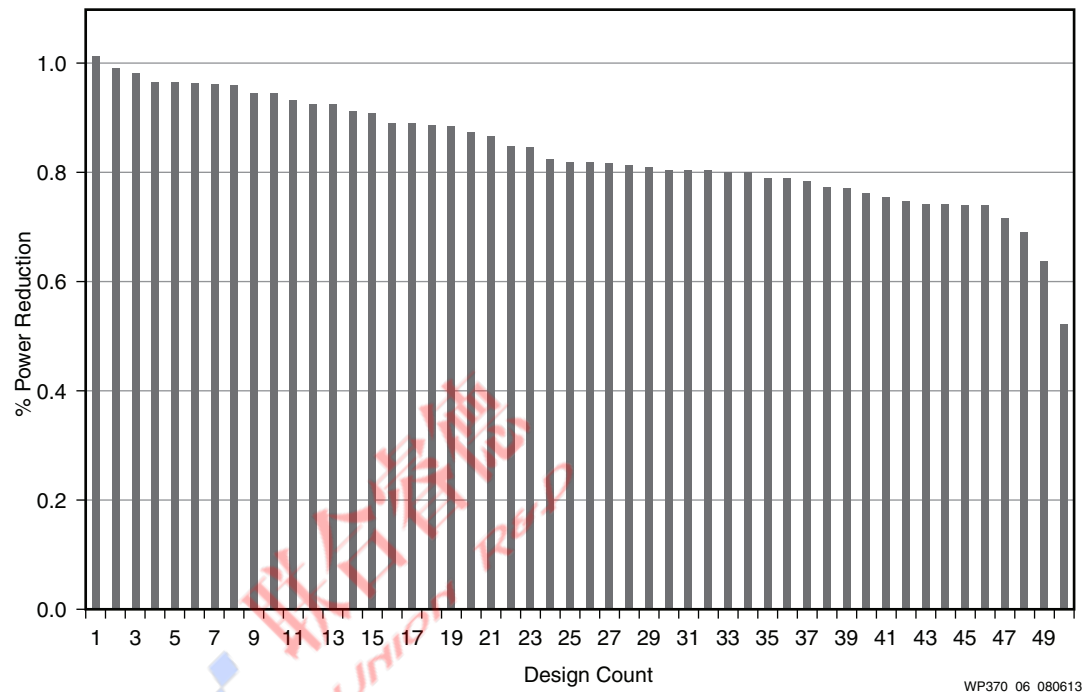


Figure 6: Dynamic Power Reduction

Summary

The intelligent clock-gating optimization feature provided in the Vivado Design Suite greatly simplifies the effort to reduce dynamic power in FPGA designs. The traditional approach to clock-gating optimization used in ASIC design presupposes an intimate knowledge of the design, thereby virtually precluding optimization of legacy and third-party IP blocks. New tools, new steps, and complex timing analyses are typically required to compensate for the inevitable new “gated clocks” and the changes in logic that are produced.

In contrast, Xilinx intelligent clock-gating optimizations are automatically performed on the entire design (or portions of it), introduce no new tools to the flow (compared to the default flow), and generate no changes to the existing logic or clocks that would alter the behavior or timing of the original design version.

For training on Vivado Design Suite power optimizations, view this video at: www.xilinx.com/csi/training/vivado/power-optimization-using-vivado.htm

For more information, go to the Xilinx power efficiency web page at: www.xilinx.com/products/technology/power/index.htm

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
05/03/10	1.0	Initial Xilinx release.
07/23/10	1.1	Added Additional Optimizations , Figure 3 , and Figure 4 .
10/05/10	1.2	Added Spartan-6 FPGA information. Updated Additional Optimizations . Added Figure 6 .
03/01/11	1.3	Added Kintex-7 and Virtex-7 FPGA information. Updated ISE version information. Updated Intelligent Clock Gating Overview .
08/29/13	1.4	Added Zynq-7000 AP SoCs. Updated Intelligent Clock Gating in the Design Implementation Flow , Benchmark Results , and Summary .

Notice of Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.