

Introduction

HDL coding styles can have a significant effect on the quality of results that you achieve for programmable logic designs. Synthesis tools optimize HDL code for both logic utilization and performance. However, sometimes the best optimizations require human understanding of the design, and synthesis tools have no information about the purpose or intent of the design. You are often in the best position to improve your quality of results.

This chapter addresses HDL coding style recommendations to ensure optimal synthesis results when targeting Altera® devices, including the following sections:

- “Quartus II Language Templates” on page 6–1
- “Using Altera Megafunctions” on page 6–2
- “Instantiating Altera Megafunctions in HDL Code” on page 6–3
- “Inferring Multiplier and DSP Functions from HDL Code” on page 6–6
- “Inferring Memory Functions from HDL Code” on page 6–12
- “Coding Guidelines for Registers and Latches” on page 6–36
- “General Coding Guidelines” on page 6–46
- “Designing with Low-Level Primitives” on page 6–71



For additional guidelines about structuring your design, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*. For additional hand-crafted techniques, you can use to optimize design blocks for the adaptive logic modules (ALMs) in many Altera devices, including a collection of circuit building blocks and related discussions, refer to the *Advanced Synthesis Cookbook: A Design Guide for Stratix II, Stratix III, and Stratix IV Devices*.

Altera’s website also provides design examples for other types of functions and to target specific applications. Refer to the [Design Examples](#) page and the [Reference Designs](#) page.

For style recommendations, options, or HDL attributes specific to your synthesis tool (including Quartus® II integrated synthesis and other EDA tools), refer to the tool vendor’s documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

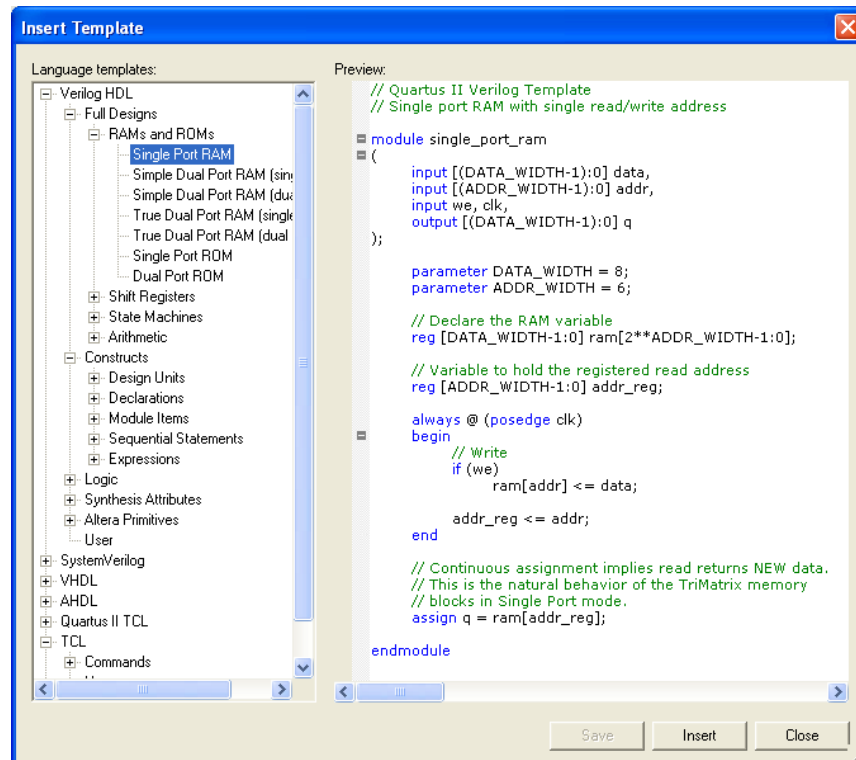
Quartus II Language Templates

The Quartus II software provides Verilog HDL, VHDL, AHDL, Tcl script, and megafunction language templates that can help you with your design.

Many of the Verilog HDL and VHDL examples in this document correspond with examples in the templates. You can easily insert examples from this document into your HDL source code using the **Insert Template** dialog box in the Quartus II user interface, shown in [Figure 6-1](#).

To open the **Insert Template** dialog box when you have a file open in the Quartus II Text Editor, on the Edit menu, click **Insert Template**. Alternatively, you can right-click in the Text Editor window and choose **Insert Template**.

Figure 6-1. Insert Template Dialog Box



Using Altera Megafunctions

Altera provides parameterizable megafunctions that are optimized for Altera device architectures. Using megafunctions instead of coding your own logic saves valuable design time. Additionally, the Altera-provided megafunctions may offer more efficient logic synthesis and device implementation. You can scale the megafunction's size and set various options by setting parameters. Megafunctions include the library of parameterized modules (LPM) and Altera device-specific megafunctions.

To use megafunctions in your HDL code, you can instantiate them as described in ["Instantiating Altera Megafunctions in HDL Code"](#) on page 6-3.

Sometimes it is preferable to make your code independent of device family or vendor. In this case, you might not want to instantiate megafunctions directly. For some types of logic functions, such as memories and DSP functions, you can infer a megafunction instead of instantiating it. Synthesis tools, including Quartus II integrated synthesis, recognize certain types of HDL code and automatically infer the appropriate megafunction. The synthesis tool uses the Altera megafunction code when compiling your design—even when you do not specifically instantiate the megafunction. Synthesis tools infer megafunctions to take advantage of logic that is optimized for Altera devices or to target dedicated architectural blocks.

In cases where you prefer to use generic HDL code instead of instantiating a megafunction, follow the guidelines and coding examples in “[Inferring Multiplier and DSP Functions from HDL Code](#)” on page 6-6 and “[Inferring Memory Functions from HDL Code](#)” on page 6-12 to ensure your HDL code infers the appropriate Altera megafunction.



You must use megafunctions to access some Altera device-specific architecture features. You can infer or instantiate megafunctions to target some features such as memory and DSP blocks. You must instantiate megafunctions to target certain device and high-speed features such as LVDS drivers, phase-locked loops (PLLs), transceivers, and double-data rate input/output (DDIO) circuitry.

For some designs, generic HDL code can provide better results than instantiating a megafunction. Refer to the following general guidelines and examples that describe when to use standard HDL code and when to use megafunctions:

- For simple addition or subtraction functions, use the + or – symbol instead of an LPM function. Instantiating an LPM function for simple arithmetic operations can result in a less efficient result because the function is hard coded and the synthesis algorithms cannot take advantage of basic logic optimizations.
- For simple multiplexers and decoders, use array notation (such as `out = data [sel]`) instead of an LPM function. Array notation works very well and has simple syntax. You can use the `lpm_mux` function to take advantage of architectural features such as cascade chains in APEX™ series devices, but use the LPM function only if you understand the device architecture in detail and want to force a specific implementation.
- Avoid division operations where possible. Division is an inherently slow operation. Many designers use multiplication creatively to produce division results.

Instantiating Altera Megafunctions in HDL Code

The following sections describe how to use megafunctions by instantiating them in your HDL code with the following methods:

- “[Instantiating Megafunctions Using the MegaWizard Plug-In Manager](#)”—You can use the MegaWizard™ Plug-In Manager to parameterize the function and create a wrapper file.
- “[Creating a Netlist File for Other Synthesis Tools](#)”—You can optionally create a netlist file instead of a wrapper file.

- “Instantiating Megafunctions Using the Port and Parameter Definition”—You can instantiate the function directly in your HDL code.

Instantiating Megafunctions Using the MegaWizard Plug-In Manager

Use the MegaWizard Plug-In Manager as described in this section to create megafunctions in the Quartus II GUI that you can instantiate in your HDL code. The MegaWizard Plug-In Manager provides a GUI to customize and parameterize megafunctions, and ensures that you set all megafunction parameters properly. When you finish setting parameters, you can specify which files you want generated. Depending on which language you choose, the MegaWizard Plug-In Manager instantiates the megafunction with the correct parameters and generates a megafunction variation file (wrapper file) in Verilog HDL (.v), VHDL (.vhd), or AHDL (.tdf), along with other supporting files.

The MegaWizard Plug-In Manager provides options to create the following files:

- A sample instantiation template for the language of the variation file (`_inst.v | vhd | tdf`).
- Component Declaration File (.cmp) that can be used in VHDL Design Files
- ADHL Include File (.inc) that can be used in Text Design Files (.tdf)
- Quartus II Block Symbol File (.bsf) for schematic designs
- Verilog HDL module declaration file that can be used when instantiating the megafunction as a black box in a third-party synthesis tool (`_bb.v`).
- If you enable the option to generate a synthesis area and timing estimation netlist, the MegaWizard Plug-In Manager generates an additional synthesis netlist file (`_syn.v`). Refer to “Creating a Netlist File for Other Synthesis Tools” on page 6-5 for details.

Table 6-1 lists and describes the files generated by the MegaWizard Plug-In Manager.

Table 6-1. MegaWizard Plug-In Manager Generated Files (Part 1 of 2)

File	Description
<code><output file>.v (1)</code>	Verilog HDL Variation Wrapper File—Megafunction wrapper file for instantiation in a Verilog HDL design.
<code><output file>.vhd (1)</code>	VHDL Variation Wrapper File—Megafunction wrapper file for instantiation in a VHDL design.
<code><output file>.tdf (1)</code>	AHDL Variation Wrapper File—Megafunction wrapper file for instantiation in an AHDL design.
<code><output file>.inc</code>	ADHL Include File—Used in AHDL designs.
<code><output file>.cmp</code>	Component Declaration File—Used in VHDL designs.
<code><output file>.bsf</code>	Block Symbol File—Used in Quartus II Block Design Files (.bdf).
<code><output file>_inst.v</code>	Verilog HDL Instantiation Template—Sample Verilog HDL instantiation of the module in the megafunction wrapper file.
<code><output file>_inst.vhd</code>	VHDL Instantiation Template—Sample VHDL instantiation of the entity in the megafunction wrapper file.
<code><output file>_inst.tdf</code>	Text Design File Instantiation Template—Sample AHDL instantiation of the subdesign in the megafunction wrapper file.

Table 6-1. MegaWizard Plug-In Manager Generated Files (Part 2 of 2)

File	Description
<code><output file>_bb.v</code>	Black box Verilog HDL Module Declaration—Hollow-body module declaration that can be used in Verilog HDL designs to specify port directions when creating black boxes in third-party synthesis tools.
<code><output file>_syn.v (2)</code>	Synthesis area and timing estimation netlist—Megafunction netlist may be used by third-party synthesis tools to improve area and timing estimations.

Notes to Table 6-1:

- (1) The MegaWizard Plug-In Manager generates either the `.v`, `.vhd`, or `.edf` file, depending on the language you select for the output file on the megafunction-selection page of the wizard.
- (2) The MegaWizard Plug-In Manager generates this file only if you turn on the **Generate netlist** option under **Timing and resource estimation** on the **EDA** page of the wizard.


Creating a Netlist File for Other Synthesis Tools

When you use certain megafunctions with third-party EDA synthesis tools (that is, tools other than Quartus II integrated synthesis), you can optionally create a netlist for area and timing estimation instead of a wrapper file.

The netlist file is a representation of the customized logic used in the Quartus II software. The file provides the connectivity of architectural elements in the megafunction but may not represent true functionality. This information enables certain third-party synthesis tools to better report area and timing estimates. In addition, synthesis tools can use the timing information to focus timing-driven optimizations and improve the quality of results.


To generate the netlist, turn on **Generate a synthesis area and timing estimation netlist** on the EDA page of the MegaWizard Plug-In Manager. The netlist file is called `<output file>_syn.v`. If you use this netlist for synthesis, you must include the megafunction wrapper file `<output file>.v | vhd` in your Quartus II project for placement and routing.


Your synthesis tool may call the Quartus II software in the background to generate this netlist, so you might not be required to perform the extra step of turning on this option.

 For information about support for area and timing estimation netlists in your synthesis tool, refer to the tool vendor's documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Instantiating Megafunctions Using the Port and Parameter Definition

You can instantiate the megafunction directly in your Verilog HDL, VHDL, or AHDL code by calling the megafunction and setting its parameters as you would any other module, component, or subdesign.



 Refer to the specific megafunction in the Quartus II Help for a list of the megafunction ports and parameters. The Quartus II Help also provides a sample VHDL component declaration and AHDL function prototype for each megafunction.

-  Altera strongly recommends that you use the MegaWizard Plug-In Manager for complex megafunctions such as PLLs, transceivers, and LVDS drivers. For details about using the MegaWizard Plug-In Manager, refer to “Instantiating Megafunctions Using the MegaWizard Plug-In Manager” on page 6-4.

Inferring Multiplier and DSP Functions from HDL Code

The following sections describe how to infer multiplier and DSP functions from generic HDL code, and, if applicable, how to target the dedicated DSP block architecture in Altera devices:

- “Multipliers—Inferring the LPM_MULT Megafunction from HDL Code”
- “Multiply-Accumulators and Multiply-Adders—Inferring ALTMULT_ACCUM and ALTMULT_ADD Megafunctions from HDL Code” on page 6-8

-  For synthesis tool features and options, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.
-  For more design examples involving advanced multiply functions and complex DSP functions, refer to the [DSP Design Examples](#) page on Altera’s website.

Multipliers—Inferring the LPM_MULT Megafunction from HDL Code

To infer multiplier functions, synthesis tools look for multipliers and convert them to LPM_MULT or ALTMULT_ADD megafunctions, or may map them directly to device atoms. For devices with DSP blocks, the software can implement the function in a DSP block instead of logic, depending on device utilization. The Quartus II Fitter can also place input and output registers in DSP blocks (that is, perform register packing) to improve performance and area utilization.

-  For additional information about the DSP block and the supported functions, refer to the appropriate Altera device family handbook and Altera’s [DSP Solutions Center](#) website.

[Example 6-1](#) and [Example 6-2](#) show Verilog HDL code examples, and [Example 6-3](#) and [Example 6-4](#) show VHDL code examples, for unsigned and signed multipliers that synthesis tools can infer as an LPM_MULT or ALTMULT_ADD megafunction. Each example fits into one DSP block 9-bit element. In addition, when register packing occurs, no extra logic cells for registers are required.

-  The signed declaration in Verilog HDL is a feature of the Verilog 2001 Standard.

Example 6-1. Verilog HDL Unsigned Multiplier

```
module unsigned_mult (out, a, b);
    output [15:0] out;
    input  [7:0] a;
    input  [7:0] b;
    assign out = a * b;
endmodule
```

Example 6-2. Verilog HDL Signed Multiplier with Input and Output Registers (Pipelining = 2)

```
module signed_mult (out, clk, a, b);
    output [15:0] out;
    input clk;
    input signed [7:0] a;
    input signed [7:0] b;

    reg signed [7:0] a_reg;
    reg signed [7:0] b_reg;
    reg signed [15:0] out;
    wire signed [15:0] mult_out;

    assign mult_out = a_reg * b_reg;

    always @ (posedge clk)
    begin
        a_reg <= a;
        b_reg <= b;
        out <= mult_out;
    end
endmodule
```

Example 6-3. VHDL Unsigned Multiplier with Input and Output Registers (Pipelining = 2)

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsigned_mult IS
    PORT (
        a: IN UNSIGNED (7 DOWNTO 0);
        b: IN UNSIGNED (7 DOWNTO 0);
        clk: IN STD_LOGIC;
        aclr: IN STD_LOGIC;
        result: OUT UNSIGNED (15 DOWNTO 0)
    );
END unsigned_mult;

ARCHITECTURE rtl OF unsigned_mult IS
    SIGNAL a_reg, b_reg: UNSIGNED (7 DOWNTO 0);
BEGIN
    PROCESS (clk, aclr)
    BEGIN
        IF (aclr = '1') THEN
            a_reg <= (OTHERS => '0');
            b_reg <= (OTHERS => '0');
            result <= (OTHERS => '0');
        ELSIF (clk'event AND clk = '1') THEN
            a_reg <= a;
            b_reg <= b;
            result <= a_reg * b_reg;
        END IF;
    END PROCESS;
END rtl;
```

Example 6-4. VHDL Signed Multiplier

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY signed_mult IS
  PORT (
    a: IN SIGNED (7 DOWNTO 0);
    b: IN SIGNED (7 DOWNTO 0);
    result: OUT SIGNED (15 DOWNTO 0)
  );
END signed_mult;

BEGIN
  result <= a * b;
END rtl;

```

Multiply-Accumulators and Multiply-Adders—Inferring ALTMULT_ACCUM and ALTMULT_ADD Megafunctions from HDL Code

Synthesis tools detect multiply-accumulators or multiply-adders and convert them to ALTMULT_ACCUM or ALTMULT_ADD megafunctions, respectively, or may map them directly to device atoms. The Quartus II software then places these functions in DSP blocks during placement and routing.



Synthesis tools infer multiply-accumulator and multiply-adder functions only if the Altera device family has dedicated DSP blocks that support these functions.

A simple multiply-accumulator consists of a multiplier feeding an addition operator. The addition operator feeds a set of registers that then feeds the second input to the addition operator. A simple multiply-adder consists of two to four multipliers feeding one or two levels of addition, subtraction, or addition/subtraction operators. Addition is always the second-level operator, if it is used. In addition to the multiply-accumulator and multiply-adder, the Quartus II Fitter also places input and output registers into the DSP blocks to pack registers and improve performance and area utilization.

Some device families offer additional advanced multiply-add and accumulate functions, such as complex multiplication, input shift register, or larger multiplications.



For details about advanced DSP block features, refer to the appropriate device handbook. For more design examples involving DSP functions and inferring advanced features in the multiply-add and multiply-accumulate circuitry, refer to the [DSP Design Examples](#) page on Altera's website.

The Verilog HDL and VHDL code samples shown in [Example 6-5](#) through [Example 6-8](#) infer multiply-accumulators and multiply-adders with input, output, and pipeline registers as well as an optional asynchronous clear signal. Using the three sets of registers provides the best performance through the function, with a latency of 3. You can remove the registers in your design to reduce the latency.

Example 6-5. Verilog HDL Unsigned Multiply-Accumulator

```
module unsig_altmult_accum (dataout, dataa, datab, clk, aclr, clken);
    input [7:0] dataa;
    input [7:0] datab;
    input clk;
    input aclr;
    input clken;
    output [31:0] dataout;

    reg [31:0] dataout;
    reg [7:0] dataa_reg;
    reg [7:0] datab_reg;
    reg [15:0] multa_reg;
    wire [15:0] multa;
    wire [31:0] adder_out;
    assign multa = dataa_reg * datab_reg;
    assign adder_out = multa_reg + dataout;

    always @ (posedge clk or posedge aclr)
    begin
        if (aclr)
            begin
                dataa_reg <= 8'b0;
                datab_reg <= 8'b0;
                multa_reg <= 16'b0;
                dataout <= 32'b0;
            end
        else if (clken)
            begin
                dataa_reg <= dataa;
                datab_reg <= datab;
                multa_reg <= multa;
                dataout <= adder_out;
            end
        end
    end
endmodule
```

Example 6-6. Verilog HDL Signed Multiply-Adder

```
module sig_altmult_add (dataa, datab, datac, datad, clock, aclr,
result);
    input signed [15:0] dataa, datab, datac, datad;
    input clock, aclr;
    output reg signed [32:0] result;

    reg signed [15:0] dataa_reg, datab_reg, datac_reg, datad_reg;
    reg signed [31:0] mult0_result, mult1_result;

    always @ (posedge clock or posedge aclr) begin
        if (aclr) begin
            dataa_reg <= 16'b0;
            datab_reg <= 16'b0;
            datac_reg <= 16'b0;
            datad_reg <= 16'b0;
            mult0_result <= 32'b0;
            mult1_result <= 32'b0;
            result <= 33'b0;
        end
        else begin
            dataa_reg <= dataa;
            datab_reg <= datab;
            datac_reg <= datac;
            datad_reg <= datad;
            mult0_result <= dataa_reg * datab_reg;
            mult1_result <= datac_reg * datad_reg;
            result <= mult0_result + mult1_result;
        end
    end
end
endmodule
```

Example 6-7. VHDL Signed Multiply-Accumulator

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY sig_altmult_accum IS
  PORT (
    a: IN SIGNED(7 DOWNTO 0);
    b: IN SIGNED (7 DOWNTO 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    accum_out: OUT SIGNED (15 DOWNTO 0)
  ) ;
END sig_altmult_accum;

ARCHITECTURE rtl OF sig_altmult_accum IS
  SIGNAL a_reg, b_reg: SIGNED (7 DOWNTO 0);
  SIGNAL pdt_reg: SIGNED (15 DOWNTO 0);
  SIGNAL adder_out: SIGNED (15 DOWNTO 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') then
      a_reg <= (others => '0');
      b_reg <= (others => '0');
      pdt_reg <= (others => '0');
      adder_out <= (others => '0');
    ELSIF (clk'event and clk = '1') THEN
      a_reg <= (a);
      b_reg <= (b);
      pdt_reg <= a_reg * b_reg;
      adder_out <= adder_out + pdt_reg;
    END IF;
  END process;
  accum_out <= adder_out;
END rtl;
```

Example 6-8. VHDL Unsigned Multiply-Adder

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsignedmult_add IS
  PORT (
    a: IN UNSIGNED (7 DOWNTO 0);
    b: IN UNSIGNED (7 DOWNTO 0);
    c: IN UNSIGNED (7 DOWNTO 0);
    d: IN UNSIGNED (7 DOWNTO 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    result: OUT UNSIGNED (15 DOWNTO 0)
  );
END unsignedmult_add;

ARCHITECTURE rtl OF unsignedmult_add IS
  SIGNAL a_reg, b_reg, c_reg, d_reg: UNSIGNED (7 DOWNTO 0);
  SIGNAL pdt_reg, pdt2_reg: UNSIGNED (15 DOWNTO 0);
  SIGNAL result_reg: UNSIGNED (15 DOWNTO 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') THEN
      a_reg <= (OTHERS => '0');
      b_reg <= (OTHERS => '0');
      c_reg <= (OTHERS => '0');
      d_reg <= (OTHERS => '0');
      pdt_reg <= (OTHERS => '0');
      pdt2_reg <= (OTHERS => '0');

    ELSIF (clk'event AND clk = '1') THEN
      a_reg <= a;
      b_reg <= b;
      c_reg <= c;
      d_reg <= d;
      pdt_reg <= a_reg * b_reg;
      pdt2_reg <= c_reg * d_reg;
      result_reg <= pdt_reg + pdt2_reg;
    END IF;
  END PROCESS;
  result <= result_reg;
END rtl;

```

Inferring Memory Functions from HDL Code

The following sections describe how to infer memory functions from generic HDL code and, if applicable, to target the dedicated memory architecture in Altera devices:

- “RAM Functions—Inferring ALTSYNCRAM and ALTDPRAM Megafunctions from HDL Code” on page 6-13
- “ROM Functions—Inferring ALTSYNCRAM and LPM_ROM Megafunctions from HDL Code” on page 6-29
- “Shift Registers—Inferring the ALTSHIFT_TAPS Megafunction from HDL Code” on page 6-32

- For synthesis tool features and options, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Altera's dedicated memory architecture offers a number of advanced features that can be easily targeted using the MegaWizard Plug-In Manager, as described in "Instantiating Altera Megafunctions in HDL Code" on page 6-3. The coding recommendations in the following sections provide portable examples of generic HDL code that infer the appropriate megafunction. However, if you want to use some of the advanced memory features in Altera devices, consider using the megafunction directly so that you can control the ports and parameters more easily.

RAM Functions—Inferring ALTSYNCRAM and ALTDPRAM Megafunctions from HDL Code

To infer RAM functions, synthesis tools detect sets of registers and logic that can be replaced with the ALTSYNCRAM or ALTDPRAM megafunctions for device families that have dedicated RAM blocks, or may map them directly to device memory atoms. Tools typically consider all signals and variables that have a two-dimensional array type and then create a RAM block, if applicable, based on the way the signals, variables, or both are assigned, referenced, or both in the HDL source description. This section provides examples demonstrating the coding styles that are inferred to create a memory block.

Standard synthesis tools recognize single-port and simple dual-port (one read port and one write port) RAM blocks. Some tools (such as the Quartus II software) also recognize true dual-port RAM blocks that map to the memory blocks in certain Altera devices. Tools usually do not infer small RAM blocks because small RAM blocks typically can be implemented more efficiently using the registers in regular logic.

If you are using Quartus II integrated synthesis, you can direct the software to infer ROM blocks for all sizes with the **Allow Any RAM Size for Recognition** option under **More Settings** on the **Analysis & Synthesis Settings** page of the **Settings** dialog box.

- If your design contains a RAM block that your synthesis tool does not recognize and infer, the design might require a large amount of system memory that can potentially cause compilation problems.

Some synthesis tools provide options to control the implementation of inferred RAM blocks for Altera devices with TriMatrix memory blocks. For example, Quartus II integrated synthesis provides the `ramstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block. Quartus II integrated synthesis does not map inferred memory into Stratix® III MLABs unless the HDL code specifies the appropriate `ramstyle` attribute, although the Fitter may map some memories to MLABs.

- For details about using the `ramstyle` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

When you use a formal verification flow, Altera recommends that you create RAM blocks in separate entities or modules that contain only the RAM logic. In certain formal verification flows, for example, when using Quartus II integrated synthesis, the entity or module containing the inferred RAM is put into a black box automatically because formal verification tools do not support RAM blocks. The Quartus II software issues a warning message when this occurs. If the entity or module contains any additional logic outside the RAM block, this logic also must be treated as a black box for formal verification and therefore cannot be verified.

The following subsections present several guidelines for inferring RAM functions that match the dedicated memory architecture in Altera devices, and then provide recommended HDL code for different types of memory logic.

Use Synchronous Memory Blocks

Altera recommends using synchronous memory blocks for Altera designs. The TriMatrix memory blocks in Altera's newest devices are synchronous, so RAM designs that are targeted towards architectures that contain these dedicated memory blocks must be synchronous to be mapped directly into the device architecture. For these devices, asynchronous memory logic is implemented in regular logic cells.

Synchronous memories are supported in all Altera device families. A memory block is considered synchronous if it uses one of the following read behaviors:

- Memory read occurs in a Verilog always block with a clock signal or a VHDL clocked process.
- Memory read occurs outside a clocked block, but there is a synchronous read address (that is, the address used in the read statement is registered). This type of logic is not always inferred as a memory block, depending on the target device architecture.



The synchronous memory structures in Altera devices differ from the structures in other vendors' devices. Match your design to the target device architecture to achieve the best results.

Later subsections provide coding recommendations for various memory types. All of these examples are synchronous to ensure that they can be directly mapped into the dedicated memory architecture available in Altera FPGAs.



For additional information about the dedicated memory blocks in your specific device, refer to the appropriate Altera device family data sheet on the Altera website at www.altera.com.

Avoid Unsupported Reset and Control Conditions

To ensure that your HDL code can be implemented in the target device architecture, avoid unsupported reset conditions or other control logic that does not exist in the device architecture.

The RAM contents of Altera memory blocks cannot be cleared with a reset signal during device operation. If your HDL code describes a RAM with a reset signal for the RAM contents, the logic is implemented in regular logic cells instead of a memory block. As a general rule, avoid putting RAM read or write operations in an always block or process block with a reset signal. If you want to specify memory contents, initialize the memory as described in [“Specifying Initial Memory Contents at Power-Up” on page 6-27](#) or write the data to the RAM during device operation.

[Example 6-9](#) shows an example of undesirable code where there is a reset signal that clears part of the RAM contents. Avoid this coding style because it is not supported in Altera memories.

Example 6-9. Verilog RAM with Reset Signal that Clears RAM Contents: Not Supported in Device Architecture

```
module clear_ram
(
    input clock,
    input reset,
    input we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out
);

    reg [7:0] mem [0:31];
    integer i;

    always @ (posedge clock or posedge reset)
    begin
        if (reset == 1'b1)
            mem[address] <= 0;
        else if (we == 1'b1)
            mem[address] <= data_in;

        data_out <= mem[address];
    end
endmodule
```

[Example 6-10](#) shows an example of undesirable code where the reset signal affects the RAM, although the effect may not be intended. Avoid this coding style because it is not supported in Altera memories.

Example 6-10. Verilog RAM with Reset Signal that Affects RAM: Not Supported in Device Architecture

```

module bad_reset
(
    input clock,
    input reset,
    input we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out,
    input d,
    output reg q
);

    reg [7:0] mem [0:31];
    integer i;

    always @ (posedge clock or posedge reset)
    begin
        if (reset == 1'b1)
            q <= 0;
        else
            begin
                if (we == 1'b1)
                    mem[address] <= data_in;

                data_out <= mem[address];
                q <= d;
            end
        end
    endmodule

```

In addition to reset signals, other control logic can prevent memory logic from being inferred as a memory block. For example, you cannot use a clock enable on the read address registers in Stratix devices, because doing so affects the output latch of the RAM, and therefore the synthesized result in the device RAM architecture would not match the HDL description. In Stratix II, Cyclone® II, Arria® GX, and other newer devices, however, you can use the address stall feature as a read address clock enable, so there is no such limitation. Check the documentation on your device architecture to ensure that your code matches the hardware available in the device.

Check Read-During-Write Behavior

It is important to check the read-during-write behavior of the memory block described in your HDL design as compared to the behavior in your target device architecture. Your HDL source code specifies the memory behavior when you read and write from the same memory address in the same clock cycle. The code specifies that the read returns either the old data at the address, or the new data being written to the address. This is referred to as the read-during-write behavior of the memory block. Altera memory blocks have different read-during-write behavior depending on the target device family, memory mode, and block type.

Synthesis tools map an HDL design into the target device architecture, with the goal of maintaining the functionality described in your source code. Therefore, if your source code specifies unsupported read-during-write behavior for the device RAM blocks, the software must implement the logic outside the RAM hardware in regular logic cells.

One common problem occurs when there is a continuous read in the HDL code, as shown in the following samples. You should avoid using these coding styles:

```
//Verilog HDL concurrent signal assignment
assign q = ram[raddr_reg];

-- VHDL concurrent signal assignment
q <= ram(raddr_reg);
```

When a write operation occurs, this type of HDL implies that the read should immediately reflect the new data at the address, independent of the read clock. However, that is not the behavior of TriMatrix memory blocks. In the device architecture, the new data is not available until the next edge of the read clock. Therefore, if the synthesis tool mapped the logic directly to a TriMatrix memory block, the device functionality and gate-level simulation results would not match the HDL description or function simulation results. If the write clock and read clock are the same, the synthesis tool can infer memory blocks and add extra bypass logic so that the device behavior does match the HDL behavior. If the write and read clocks are different, the synthesis tool cannot reliably add bypass logic, so the logic is implemented in regular logic cells instead of dedicated RAM blocks. The examples in the following sections discuss some of these differences for read-during-write conditions.

In addition, the MLAB feature in Stratix III logic array blocks (LABs) does not easily support old data or new data behavior for a read during write in the dedicated device architecture. Implementing the extra logic to support this behavior significantly reduces timing performance through the memory.



For best performance in MLAB memories, your design should not depend on the read data during a write operation.

In many synthesis tools, you can specify that the read-during-write behavior is not important to your design; for example, if you never read from the same address to which you write in the same clock cycle. For Quartus II integrated synthesis, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than use the behavior specified by your HDL code. Using this type of attribute prevents the synthesis tool from using extra logic to implement the memory block, and in some cases, can allow memory inference when it would otherwise be impossible.



For more information about attribute syntax, the `no_rw_check` attribute value, or specific options for your synthesis tool, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

The following subsections provide coding recommendations for various memory types. Each example describes the read-during-write behavior and addresses the support for the memory type in Altera devices.

Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior

The code examples in this section show Verilog HDL and VHDL code that infers simple dual-port, single-clock synchronous RAM. Single-port RAM blocks use a similar coding style.

The read-during-write behavior in these examples is to read the old data at the memory address. Refer to “[Check Read-During-Write Behavior](#)” on page 6-16 for details. Altera recommends that you use this coding style for most RAM blocks as long as your design does not require that a simultaneous read and write to the same RAM location read the new value that is currently being written to that RAM location. For best performance in MLAB memories, use the appropriate attribute so that your design does not depend on the read data during a write operation.

If you require that the read-during-write results in new data, refer to “[Single-Clock Synchronous RAM with New Data Read-During-Write Behavior](#)” on page 6-19.

The simple dual-port RAM code samples shown in [Example 6-11](#) and [Example 6-12](#) map directly into Altera TriMatrix memory.

Single-port versions of memory blocks (that is, using the same read address and write address signals) can allow better RAM utilization than dual-port memory blocks, depending on the device family.

Example 6-11. Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior

```
module single_clk_ram(  
    output reg [7:0] q,  
    input [7:0] d,  
    input [6:0] write_address, read_address,  
    input we, clk  
);  
    reg [7:0] mem [127:0];  
  
    always @ (posedge clk) begin  
        if (we)  
            mem[write_address] <= d;  
        q <= mem[read_address]; // q doesn't get d in this clock cycle  
    end  
endmodule
```

Example 6-12. VHDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_ram IS
  PORT (
    clock: IN STD_LOGIC;
    data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
    write_address: IN INTEGER RANGE 0 to 31;
    read_address: IN INTEGER RANGE 0 to 31;
    we: IN STD_LOGIC;
    q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
  );
END single_clock_ram;

ARCHITECTURE rtl OF single_clock_ram IS
  TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
  SIGNAL ram_block: MEM;
BEGIN
  PROCESS (clock)
  BEGIN
    IF (clock'event AND clock = '1') THEN
      IF (we = '1') THEN
        ram_block(write_address) <= data;
      END IF;
      q <= ram_block(read_address);
      -- VHDL semantics imply that q doesn't get data
      -- in this clock cycle
    END IF;
  END PROCESS;
END rtl;

```

Single-Clock Synchronous RAM with New Data Read-During-Write Behavior

The examples in this section describe RAM blocks in which a simultaneous read and write to the same location reads the new value that is currently being written to that RAM location.

To implement this behavior in the target device, synthesis software adds bypass logic around the RAM block. This bypass logic increases the area utilization of the design and decreases the performance if the RAM block is part of the design's critical path. Refer to [“Check Read-During-Write Behavior” on page 6-16](#) for details. If this behavior is not required for your design, use the examples from [“Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior” on page 6-17](#).

The simple dual-port RAM shown in [Example 6-13](#) and [Example 6-14](#) require the software to create bypass logic around the RAM block.

Single-port versions of the Verilog memory block (that is, using the same read address and write address signals) do not require any logic cells to create bypass logic in the Arria GX, Stratix, and Cyclone series of devices, because the device memory supports new data read-during-write behavior when in single-port mode (same clock, same read, and write address).

Example 6-13. Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```

module single_clock_wr_ram(
    output reg [7:0] q,
    input [7:0] d,
    input [6:0] write_address, read_address,
    input we, clk
);
    reg [7:0] mem [127:0];

    always @ (posedge clk) begin
        if (we)
            mem[write_address] = d;
        q = mem[read_address]; // q does get d in this clock cycle if
we is high
    end
endmodule

```



Example 6-13 is similar to **Example 6-11**, but **Example 6-13** uses a blocking assignment for the write so that the data is assigned immediately.

An alternative way to create a single-clock RAM is to use an assign statement to read the address of mem to create the output q, as shown in the following coding style. By itself, the code describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, then a read-during-write would result in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary. For this reason, avoid using this alternate type of coding style:

```

reg [7:0] mem [127:0];
reg [6:0] read_address_reg;

always @ (posedge clk) begin
    if (we)
        mem[write_address] <= d;

    read_address_reg <= read_address;
end

assign q = mem[read_address_reg];

```

The VHDL sample in [Example 6-14](#) uses a concurrent signal assignment to read from the RAM. By itself, this example describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, a read-during-write results in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary.

Example 6-14. VHDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_rw_ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
    );
END single_clock_rw_ram;

ARCHITECTURE rtl OF single_clock_rw_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;
    SIGNAL read_address_reg: INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
    q <= ram_block(read_address_reg);
END rtl;
```

[Example 6-14](#) does not infer a RAM block for the APEX series of devices, ACEX®, or the FLEX® series of devices by default because the read-during-write behavior depends on surrounding logic. For Quartus II integrated synthesis, if you do not require the read-through-write capability, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than use the behavior specified by your HDL code.

Simple Dual-Port, Dual-Clock Synchronous RAM

In dual clock designs, synthesis tools cannot accurately infer the read-during-write behavior because it depends on the timing of the two clocks within the target device. Therefore, the read-during-write behavior of the synthesized design is undefined and may differ from your original HDL code. Refer to [“Check Read-During-Write Behavior” on page 6-16](#) for details.

When Quartus II integrated synthesis infers this type of RAM, it issues a warning because of the undefined read-during-write behavior. If this functionality is acceptable in your design, you can avoid the warning by adding the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM.

The code samples shown in [Example 6-15](#) and [Example 6-16](#) show Verilog HDL and VHDL code that infers dual-clock synchronous RAM. The exact behavior depends on the relationship between the clocks.

Example 6-15. Verilog HDL Simple Dual-Port, Dual-Clock Synchronous RAM

```
module dual_clock_ram(  
    output reg [7:0] q,  
    input [7:0] d,  
    input [6:0] write_address, read_address,  
    input we, clk1, clk2  
);  
    reg [6:0] read_address_reg;  
    reg [7:0] mem [127:0];  
  
    always @ (posedge clk1)  
    begin  
        if (we)  
            mem[write_address] <= d;  
    end  
  
    always @ (posedge clk2) begin  
        q <= mem[read_address_reg];  
        read_address_reg <= read_address;  
    end  
endmodule
```

Example 6-16. VHDL Simple Dual-Port, Dual-Clock Synchronous RAM

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY dual_clock_ram IS
    PORT (
        clock1, clock2: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (3 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (3 DOWNTO 0)
    );
END dual_clock_ram;
ARCHITECTURE rtl OF dual_clock_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(3 DOWNTO 0);
    SIGNAL ram_block: MEM;
    SIGNAL read_address_reg : INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock1)
    BEGIN
        IF (clock1'event AND clock1 = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
        END IF;
    END PROCESS;
    PROCESS (clock2)
    BEGIN
        IF (clock2'event AND clock2 = '1') THEN
            q <= ram_block(read_address_reg);
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
END rtl;

```

True Dual-Port Synchronous RAM

The code examples in this section show Verilog HDL and VHDL code that infers true dual-port synchronous RAM. Different synthesis tools may differ in their support for these types of memories. This section describes the inference rules for Quartus II integrated synthesis. This type of RAM inference is supported only for the Arria GX, Stratix, and Cyclone series of devices.

Altera TriMatrix memory blocks have two independent address ports, allowing for operations on two unique addresses simultaneously. A read operation and a write operation can share the same port if they share the same address. The Quartus II software infers true dual-port RAMs in Verilog HDL and VHDL with any combination of independent read or write operations in the same clock cycle, with at most two unique port addresses, performing two reads and one write, two writes and one read, or two writes and two reads in one clock cycle with one or two unique addresses.

In the TriMatrix RAM block architecture, there is no priority between the two ports. Therefore, if you write to the same location on both ports at the same time, the result is indeterminate in the device architecture. You must ensure your HDL code does not imply priority for writes to the memory block, if you want the design to be implemented in a dedicated hardware memory block. For example, if both ports are defined in the same process block, the code is synthesized and simulated sequentially so there would be a priority between the two ports. If your code does imply a priority, the logic cannot be implemented in the device RAM blocks and is implemented in regular logic cells.

You must also consider the read-during-write behavior of the RAM block, to ensure that it can be mapped directly to the device RAM architecture. Refer to “[Check Read-During-Write Behavior](#)” on page 6-16 for details.

When a read and write operation occur on the same port for the same address, the read operation may behave as follows:

- **Read new data**—This mode matches the behavior of TriMatrix memory blocks.
- **Read old data**—This mode is supported only by Stratix IV, Arria II GX, Stratix III, and Cyclone III TriMatrix memory blocks. This behavior is not possible in TriMatrix memory blocks of other families.

When a read and write operation occur on different ports for the same address (also known as mixed port), the read operation may behave as follows:

- **Read new data**—Quartus II integrated synthesis supports this mode by creating bypass logic around the TriMatrix memory block.
- **Read old data**—This behavior is supported by TriMatrix memory blocks.

The Verilog HDL single-clock code sample shown in [Example 6-17](#) maps directly into Altera TriMatrix memory. When a read and write operation occur on the same port for the same address, the new data being written to the memory is read. When a read and write operation occur on different ports for the same address, the old data in the memory is read. Simultaneous writes to the same location on both ports results in indeterminate behavior.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

Example 6-17. Verilog HDL True Dual-Port RAM with Single Clock

```

module true_dual_port_ram_single_clock
(
    input [(DATA_WIDTH-1):0] data_a, data_b,
    input [(ADDR_WIDTH-1):0] addr_a, addr_b,
    input we_a, we_b, clk,
    output reg [(DATA_WIDTH-1):0] q_a, q_b
);

    parameter DATA_WIDTH = 8;
    parameter ADDR_WIDTH = 6;

    // Declare the RAM variable
    reg [DATA_WIDTH-1:0] ram[2**ADDR_WIDTH-1:0];

    always @ (posedge clk)
    begin // Port A
        if (we_a)
            begin
                ram[addr_a] <= data_a;
                q_a <= data_a;
            end
        else
            q_a <= ram[addr_a];
    end
    always @ (posedge clk)
    begin // Port b
        if (we_b)
            begin
                ram[addr_b] <= data_b;
                q_b <= data_b;
            end
        else
            q_b <= ram[addr_b];
        end
    end

endmodule

```

If you use the following Verilog HDL read statements instead of the `if-else` statements in [Example 6-17](#), the HDL code specifies that the read results in old data when a read and write operation occur at the same time for the same address on the same port or mixed ports. (This behavior is supported only in the TriMatrix memories of Stratix IV, Arria II GX, Stratix III and Cyclone III devices, and is not inferred as memory for other device families):

```

always @ (posedge clk)
begin // Port A
    if (we_a)
        ram[addr_a] <= data_a;

    q_a <= ram[addr_a];
end

always @ (posedge clk)
begin // Port B
    if (we_b)
        ram[addr_b] <= data_b;

    q_b <= ram[addr_b];
end

```

The VHDL single-clock code sample shown in [Example 6-18](#) maps directly into Altera TriMatrix memory. When a read and write operation occurs on the same port for the same address, the new data being written to the memory is read. When a read and write operation occurs on different ports for the same address, the old data in the memory is read. Because simultaneous writes to the same location on both ports results in indeterminate behavior, Altera recommends that you avoid this condition.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

Example 6-18. VHDL True Dual-Port RAM with Single Clock (Part 1 of 2)

```
library ieee;
use ieee.std_logic_1164.all;

entity true_dual_port_ram_single_clock is

    generic
    (
        DATA_WIDTH : natural := 8;
        ADDR_WIDTH  : natural := 6
    );

    port
    (
        clk: in std_logic;
        addr_a: in natural range 0 to 2**ADDR_WIDTH - 1;
        addr_b: in natural range 0 to 2**ADDR_WIDTH - 1;
        data_a: in std_logic_vector((DATA_WIDTH-1) downto 0);
        data_b: in std_logic_vector((DATA_WIDTH-1) downto 0);
        we_a: in std_logic := '1';
        we_b: in std_logic := '1';
        q_a: out std_logic_vector((DATA_WIDTH -1) downto 0);
        q_b: out std_logic_vector((DATA_WIDTH -1) downto 0)
    );

end true_dual_port_ram_single_clock;

architecture rtl of true_dual_port_ram_single_clock is

    -- Build a 2-D array type for the RAM
    subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);
    type memory_t is addr_a(raddr'high downto 0) of word_t;

    -- Declare the RAM signal.
    signal ram : memory_t;
```

Example 6-18. VHDL True Dual-Port RAM with Single Clock (Part 2of 2)

```

begin

    process(clk)
    begin
        if(rising_edge(clk)) then -- Port A
            if(we_a = '1') then
                ram(addr_a) <= data_a;

                -- Read-during-write on the same port returns NEW data
                q_a <= data_a;
            else
                -- Read-during-write on the mixed port returns OLD data
                q_a <= ram(addr_a);
            end if;
        end if;

    end process;

    process(clk)
    begin
        if(rising_edge(clk)) then -- Port B
            if(we_b = '1') then
                ram(addr_b) <= data_b;

                -- Read-during-write on the same port returns NEW data
                q_b <= data_b;
            else
                -- Read-during-write on the mixed port returns OLD data
                q_b <= ram(addr_b);
            end if;
        end if;
    end process;

end rtl;

```

Specifying Initial Memory Contents at Power-Up


Your synthesis tool may offer various ways to specify the initial contents of an inferred memory.



Certain device memory types do not support initialized memory, such as the M-RAM blocks in Stratix and Stratix II devices.

There are slight power-up and initialization differences between dedicated RAM blocks and the Stratix III MLAB memory due to the continuous read of the MLAB. Altera dedicated RAM block outputs always power-up to zero and are set to the initial value on the first read. For example, if address 0 is pre-initialized to FF, the RAM block powers up with the output at 0. A subsequent read after power up from address 0 outputs the pre-initialized value of FF. Therefore, if a RAM is powered up and an enable (read enable or clock enable) is held low, then the power-up output of 0 is maintained until the first valid read cycle. The Stratix III MLAB is implemented using registers that power-up to 0, but are initialized to their initial value immediately at power-up or reset. Therefore, the initial value regardless of the enable status is seen. Quartus II integrated synthesis does not map inferred memory to MLABs unless the HDL code specifies the appropriate `ramstyle` attribute.

Quartus II integrated synthesis supports the `ram_init_file` synthesis attribute that allows you to specify a Memory Initialization File (`.mif`) for an inferred RAM block.

 For information about the `ram_init_file` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the tool vendor's documentation.

In Verilog HDL, you can use an initial block to initialize the contents of an inferred memory. Quartus II integrated synthesis automatically converts the initial block into a `.mif` for the inferred RAM. [Example 6-19](#) shows Verilog HDL code that infers a simple dual-port RAM block and corresponding `.mif` file.


Example 6-19. Verilog HDL RAM with Initialized Contents

```
module ram_with_init(
    output reg [7:0] q,
    input [7:0] d,
    input [4:0] write_address, read_address,
    input we, clk
);
    reg [7:0] mem [0:31];
    integer i;

    initial begin
        for (i = 0; i < 32; i = i + 1)
            mem[i] = i[7:0];
        end

    always @ (posedge clk) begin
        if (we)
            mem[write_address] <= d;
        q <= mem[read_address];
        end
    endmodule
```

Quartus II integrated synthesis and other synthesis tools also support the `$readmemb` and `$readmemh` commands so that RAM and ROM initialization work identically in synthesis and simulation. [Example 6-20](#) shows an initial block that initializes an inferred RAM block using the `$readmemb` command.

 Refer to the Verilog Language Reference Manual (LRM) 1364-2001 Section 17.2.8 for details about the format of the `ram.txt` file.

Example 6-20. Verilog HDL RAM Initialized with the `readmemb` Command

```
reg [7:0] ram[0:15];
initial
begin
    $readmemb("ram.txt", ram);
end
```

In VHDL, you can initialize the contents of an inferred memory by specifying a default value for the corresponding signal. Quartus II integrated synthesis automatically converts the default value into a `.mif` file for the inferred RAM. [Example 6-21](#) shows VHDL code that infers a simple dual-port RAM block and corresponding `.mif` file.

Example 6-21. VHDL RAM with Initialized Contents

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
use ieee.numeric_std.all;

ENTITY ram_with_init IS
    PORT(
        clock: IN STD_LOGIC;
        data: IN UNSIGNED (7 DOWNTO 0);
        write_address: IN integer RANGE 0 to 31;
        read_address: IN integer RANGE 0 to 31;
        we: IN std_logic;
        q: OUT UNSIGNED (7 DOWNTO 0));
END;

ARCHITECTURE rtl OF ram_with_init IS

    TYPE MEM IS ARRAY(31 DOWNTO 0) OF unsigned(7 DOWNTO 0);
    FUNCTION initialize_ram
        return MEM is
        variable result : MEM;
    BEGIN
        FOR i IN 31 DOWNTO 0 LOOP
            result(i) := to_unsigned(natural(i), natural'(8));
        END LOOP;
        RETURN result;
    END initialize_ram;

    SIGNAL ram_block : MEM := initialize_ram;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            q <= ram_block(read_address);
        END IF;
    END PROCESS;
END rtl;
    
```

ROM Functions—Inferring ALTSYNCRAM and LPM_ROM Megafunctions from HDL Code


To infer ROM functions, synthesis tools detect sets of registers and logic that can be replaced with the ALTSYNCRAM or LPM_ROM megafunctions, depending on the target device family, only for device families that have dedicated memory blocks.

ROMs are inferred when a CASE statement exists in which a value is set to a constant for every choice in the case statement. Because small ROMs typically achieve the best performance when they are implemented using the registers in regular logic, each ROM function must meet a minimum size requirement to be inferred and placed into memory.




If you use Quartus II integrated synthesis, you can direct the software to infer ROM blocks for all sizes with the **Allow Any ROM Size for Recognition** option under **More Settings** on the **Analysis & Synthesis Settings** page of the **Settings** dialog box.

Some synthesis tools provide options to control the implementation of inferred ROM blocks for Altera devices with TriMatrix memory blocks. For example, Quartus II integrated synthesis provides the `romstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block.

 For details about using the `romstyle` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

When you are using a formal verification flow, Altera recommends that you create ROM blocks in separate entities or modules that contain only the ROM logic because you may need to treat the entity and module as a black box during formal verification.

 Because formal verification tools do not support ROM megafunctions, Quartus II integrated synthesis does not infer ROM megafunctions when a formal verification tool is selected.

The Verilog HDL and VHDL code samples shown in [Example 6-22](#) through [Example 6-25](#) infer synchronous ROM blocks. Depending on the device family's dedicated RAM architecture, the ROM logic may have to be synchronous; consult the device family handbook for details.

For device architectures with synchronous RAM blocks, such as the Stratix series devices and newer device families, either the address or the output has to be registered for ROM code to be inferred. When output registers are used, the registers are implemented using the input registers of the RAM block, but the functionality of the ROM is not changed. If you register the address, the power-up state of the inferred ROM can be different from the HDL design. In this scenario, the synthesis software issues a warning. The Quartus II Help explains the condition under which the functionality changes when you use Quartus II integrated synthesis.

These ROM code samples map directly to the Altera TriMatrix memory architecture.

Example 6-22. Verilog HDL Synchronous ROM

```
module sync_rom (clock, address, data_out);
    input clock;
    input [7:0] address;
    output [5:0] data_out;

    reg [5:0] data_out;

    always @ (posedge clock)
    begin
        case (address)
            8'b00000000: data_out = 6'b101111;
            8'b00000001: data_out = 6'b110110;
            ...
            8'b11111110: data_out = 6'b000001;
            8'b11111111: data_out = 6'b101010;
        endcase
    end
endmodule
```

Example 6-23. VHDL Synchronous ROM

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY sync_rom IS
  PORT (
    clock: IN STD_LOGIC;
    address: IN STD_LOGIC_VECTOR(7 downto 0);
    data_out: OUT STD_LOGIC_VECTOR(5 downto 0)
  );
END sync_rom;

ARCHITECTURE rtl OF sync_rom IS
BEGIN
  PROCESS (clock)
  BEGIN
    IF rising_edge (clock) THEN
      CASE address IS
        WHEN "00000000" => data_out <= "101111";
        WHEN "00000001" => data_out <= "110110";
        ...
        WHEN "11111110" => data_out <= "000001";
        WHEN "11111111" => data_out <= "101010";
        WHEN OTHERS => data_out <= "101111";
      END CASE;
    END IF;
  END PROCESS;
END rtl;
```

Example 6-24. Verilog HDL Dual-Port Synchronous ROM Using readmemb

```
module dual_port_rom (
  input [(addr_width-1):0] addr_a, addr_b,
  input clk,
  output reg [(data_width-1):0] q_a, q_b
);
  parameter data_width = 8;
  parameter addr_width = 8;

  reg [data_width-1:0] rom[2**addr_width-1:0];

  initial // Read the memory contents in the file
  dual_port_rom_init.txt.
  begin
    $readmemb("dual_port_rom_init.txt", rom);
  end

  always @ (posedge clk)
  begin
    q_a <= rom[addr_a];
    q_b <= rom[addr_b];
  end
endmodule
```

Example 6-25. VHDL Dual-Port Synchronous ROM Using Initialization Function

```

library ieee;
use ieee.std_logic_1164.all;
use ieee.numeric_std.all;

entity dual_port_rom is
  generic (
    DATA_WIDTH : natural := 8;
    ADDR_WIDTH  : natural := 8
  );
  port (
    clk      : in std_logic;
    addr_a: in natural range 0 to 2**ADDR_WIDTH - 1;
    addr_b: in natural range 0 to 2**ADDR_WIDTH - 1;
    q_a     : out std_logic_vector((DATA_WIDTH - 1) downto 0);
    q_b     : out std_logic_vector((DATA_WIDTH - 1) downto 0)
  );
end entity;

architecture rtl of dual_port_rom is
  -- Build a 2-D array type for the ROM
  subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);
  type memory_t is array(addr_a'high downto 0) of word_t;

  function init_rom
    return memory_t is
    variable tmp : memory_t := (others => (others => '0'));
  begin
    for addr_pos in 0 to 2**ADDR_WIDTH - 1 loop
      -- Initialize each address with the address itself
      tmp(addr_pos) := std_logic_vector(to_unsigned(addr_pos,
DATA_WIDTH));
    end loop;
    return tmp;
  end init_rom;

  -- Declare the ROM signal and specify a default initialization value.
  signal rom : memory_t := init_rom;
begin
  process(clk)
  begin
    if (rising_edge(clk)) then
      q_a <= rom(addr_a);
      q_b <= rom(addr_b);
    end if;
  end process;
end rtl;

```

Shift Registers—Inferring the ALTSHIFT_TAPS Megafunction from HDL Code

To infer shift registers, synthesis tools detect a group of shift registers of the same length and convert them to an ALTSHIFT_TAPS megafunction. To be detected, all the shift registers must have the following characteristics:

- Use the same clock and clock enable
- Do not have any other secondary signals
- Have equally spaced taps that are at least three registers apart

When you use a formal verification flow, Altera recommends that you create shift register blocks in separate entities or modules containing only the shift register logic, because you might have to treat the entity or module as a black box during formal verification.



Because formal verification tools do not support shift register megafunctions, Quartus II integrated synthesis does not infer the `ALTSHIFT_TAPS` megafunction when a formal verification tool is selected. You can select EDA tools for use with your Quartus II project on the **EDA Tool Settings** page of the **Settings** dialog box.



For more information about the `ALTSHIFT_TAPS` megafunction, refer to the [ALTSHIFT_TAPS Megafunction User Guide](#).

Synthesis software recognizes shift registers only for device families that have dedicated RAM blocks and the software uses certain guidelines to determine the best implementation. The following guidelines are followed in Quartus II integrated synthesis and also are generally followed by other EDA tools:

- For FLEX 10K® and ACEX 1K devices, the software does not infer `ALTSHIFT_TAPS` megafunctions because FLEX 10K and ACEX 1K devices have a relatively small amount of dedicated memory.
- For APEX 20K and APEX II devices, the software infers the `ALTSHIFT_TAPS` megafunction only if the shift register has more than a total of 128 bits. Smaller shift registers typically do not benefit from implementation in dedicated memory.
- For the Arria GX, Stratix, and Cyclone series devices, the software determines whether to infer the `ALTSHIFT_TAPS` megafunction based on the width of the registered bus (W), the length between each tap (L), and the number of taps (N).
 - If the registered bus width is one ($W = 1$), the software infers `ALTSHIFT_TAPS` if the number of taps times the length between each tap is greater than or equal to 64 ($N \times L \geq 64$).
 - If the registered bus width is greater than one ($W > 1$), the software infers `ALTSHIFT_TAPS` if the registered bus width times the number of taps times the length between each tap is greater than or equal to 32 ($W \times N \times L \geq 32$).

If the length between each tap (L) is not a power of two, the software uses more logic to decode the read and write counters. This situation occurs because for different sizes of shift registers, external decode logic that uses logic elements (LEs) or Adaptive Logic Modules (ALMs) is required to implement the function. This decode logic eliminates the performance and utilization advantages of implementing shift registers in memory.

The registers that the software maps to the `ALTSHIFT_TAPS` megafunction and places in RAM are not available in a Verilog HDL or VHDL output file for simulation tools because their node names do not exist after synthesis.

Simple Shift Register

The code samples shown in [Example 6-26](#) and [Example 6-27](#) show a simple, single-bit wide, 64-bit long shift register. The synthesis software implements the register ($W = 1$ and $M = 64$) in an `ALTSHIFT_TAPS` megafunction for supported devices. If the length of the register is less than 64 bits, the software implements the shift register in logic.

Example 6-26. Verilog HDL Single-Bit Wide, 64-Bit Long Shift Register

```

module shift_1x64 (clk, shift, sr_in, sr_out);
    input clk, shift;
    input sr_in;
    output sr_out;

    reg [63:0] sr;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
            begin
                sr[63:1] <= sr[62:0];
                sr[0] <= sr_in;
            end
        end
        assign sr_out = sr[63];
    endmodule

```

Example 6-27. VHDL Single-Bit Wide, 64-Bit Long Shift Register

```

LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_1x64 IS
PORT (
    clk: IN STD_LOGIC;
    shift: IN STD_LOGIC;
    sr_in: IN STD_LOGIC;
    sr_out: OUT STD_LOGIC
);
END shift_1x64;

ARCHITECTURE arch OF shift_1x64 IS
TYPE sr_length IS ARRAY (63 DOWNTO 0) OF STD_LOGIC;
SIGNAL sr: sr_length;
BEGIN
PROCESS (clk)
BEGIN
IF (clk'EVENT and clk = '1') THEN
IF (shift = '1') THEN
sr(63 DOWNTO 1) <= sr(62 DOWNTO 0);
sr(0) <= sr_in;
END IF;
END IF;
END PROCESS;
sr_out <= sr(63);
END arch;

```

Shift Register with Evenly Spaced Taps

The code samples shown in [Example 6-28](#) and [Example 6-29](#) show a Verilog HDL and VHDL 8-bit wide, 64-bit long shift register ($W > 1$ and $M = 64$) with evenly spaced taps at 15, 31, and 47. The synthesis software implements this function in a single ALTSHIFT_TAPS megafunction and maps it to RAM in supported devices.

Example 6-28. Verilog HDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```
module shift_8x64_taps (clk, shift, sr_in, sr_out, sr_tap_one,
sr_tap_two, sr_tap_three );
    input clk, shift;
    input [7:0] sr_in;
    output [7:0] sr_tap_one, sr_tap_two, sr_tap_three, sr_out;

    reg [7:0] sr [63:0];
    integer n;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
        begin
            for (n = 63; n>0; n = n-1)
            begin
                sr[n] <= sr[n-1];
            end
            sr[0] <= sr_in;
        end

        end

        assign sr_tap_one = sr[15];
        assign sr_tap_two = sr[31];
        assign sr_tap_three = sr[47];
        assign sr_out = sr[63];
    endmodule
```

Example 6-29. VHDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```

LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_8x64_taps IS
    PORT (
        clk: IN STD_LOGIC;
        shift: IN STD_LOGIC;
        sr_in: IN STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_one: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_two : OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_three: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_out: OUT STD_LOGIC_VECTOR(7 DOWNTO 0)
    );
END shift_8x64_taps;

ARCHITECTURE arch OF shift_8x64_taps IS
    SUBTYPE sr_width IS STD_LOGIC_VECTOR(7 DOWNTO 0);
    TYPE sr_length IS ARRAY (63 DOWNTO 0) OF sr_width;
    SIGNAL sr: sr_length;
BEGIN
    PROCESS (clk)
    BEGIN
        IF (clk'EVENT and clk = '1') THEN
            IF (shift = '1') THEN
                sr(63 DOWNTO 1) <= sr(62 DOWNTO 0);
                sr(0) <= sr_in;
            END IF;
        END IF;
    END PROCESS;
    sr_tap_one <= sr(15);
    sr_tap_two <= sr(31);
    sr_tap_three <= sr(47);
    sr_out <= sr(63);
END arch;

```

Coding Guidelines for Registers and Latches

This section provides device-specific coding recommendations for Altera registers and latches. Understanding the architecture of the target Altera device helps ensure that your code produces the expected results and achieves the optimal quality of results.

This section provides guidelines in the following areas:

- [“Register Power-Up Values in Altera Devices”](#)
- [“Secondary Register Control Signals Such as Clear and Clock Enable” on page 6-38](#)
- [“Latches” on page 6-42](#)

Register Power-Up Values in Altera Devices

Registers in the device core always power up to a low (0) logic level on all Altera devices. However, there are ways to implement logic such that registers behave as if they were powering up to a high (1) logic level.

If you use a preset signal on a device that does not support presets in the register architecture, then your synthesis tool may convert the preset signal to a clear signal, which requires synthesis to perform an optimization referred to as NOT gate push-back. NOT gate push-back adds an inverter to the input and the output of the register so that the reset and power-up conditions will appear to be high but the device operates as expected. In this case, your synthesis tool may issue a message informing you about the power-up condition. The register itself powers up low, but the register output is inverted, so the signal that arrives at all destinations is high.


Due to these effects, if you specify a non-zero reset value, you may cause your synthesis tool to use the asynchronous clear (`ac1r`) signals available on the registers to implement the high bits with NOT gate push-back. In that case, the registers look as though they power up to the specified reset value. You see this behavior, for example, if your design targets FLEX 10KE or ACEX devices.

When a load signal is available in the device, your synthesis tools can implement a reset of 1 or 0 value by using an asynchronous load of 1 or 0. When the synthesis tool uses an asynchronous load signal, it is not performing NOT gate push-back, so the registers power up to a 0 logic level.

 For additional details, refer to the appropriate device family handbook or the appropriate handbook of the Altera website.


Designers typically use an explicit reset signal for the design, which forces all registers into their appropriate values after reset but not necessarily at power-up. You can create your design such that the asynchronous reset allows the board to operate in a safe condition and then you can bring up the design with the reset active. This is a good practice so you do not depend on the power-up conditions of the device.


You can make the your design more stable and avoid potential glitches by synchronizing external or combinational logic of the device architecture before you drive the asynchronous control ports of registers.

 For additional information about good synchronous design practices, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.

If you want to force a particular power-up condition for your design, use the synthesis options available in your synthesis tool. With Quartus II integrated synthesis, you can apply the **Power-Up Level** logic option. You can also apply the option with an `altera_attribute` assignment in your source code. Using this option forces synthesis to perform NOT gate push-back because synthesis tools cannot actually change the power-up states of core registers.

You can apply the Quartus II integrated synthesis **Power-Up Level** logic option to a specific register or to a design entity, module or subdesign. If you do so, every register in that block receives the value. Registers power up to 0 by default; therefore you can use this assignment to force all registers to power up to 1 using NOT gate push-back.

 Be aware that using NOT gate push-back as a global assignment could slightly degrade the quality of results due to the number of inverters that are required. In some situations, issues are caused by enable or secondary control logic inference. It may also be more difficult to migrate such a design to an ASIC or a HardCopy® device. You can simulate the power-up behavior in a functional simulation if you use initialization.

 The **Power-Up Level** option and the `altera_attribute` assignment are described in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

Some synthesis tools can also read the default or initial values for registered signals and implement this behavior in the device. For example, Quartus II integrated synthesis converts default values for registered signals into **Power-Up Level** settings. That way, the synthesized behavior matches the power-up state of the HDL code during a functional simulation.

For example, the code samples in [Example 6-30](#) and [Example 6-31](#) both infer a register for `q` and set its power-up level to high (while the reset value is 0).

Example 6-30. Verilog Register with Reset and High Power-Up Value

```
reg q = 1'b1;

always @ (posedge clk or posedge aclr)
begin
  if (aclr)
    q <= 1'b0;
  else
    q <= d;
end
```

Example 6-31. VHDL Register with Reset and High Power-Up Level

```
SIGNAL q : STD_LOGIC := '1'; -- q has a default value of '1'

PROCESS (clk, reset)
BEGIN
  IF (reset = '1') THEN
    q <= '0';
  ELSIF (rising_edge(clk)) THEN
    q <= d;
  END IF;
END PROCESS;
```

Secondary Register Control Signals Such as Clear and Clock Enable

FPGA device architectures contain registers, also known as “flipflops”. The registers in Altera FPGAs provide a number of secondary control signals (such as clear and enable signals) that you can use to implement control logic for each register without using extra logic cells. Device families vary in their support for secondary signals, so consult the device family data sheet to verify which signals are available in your target device.

To make the most efficient use of the signals in the device, your HDL code should match the device architecture as closely as possible. The control signals have a certain priority due to the nature of the architecture, so your HDL code should follow that priority where possible.

Your synthesis tool can emulate any control signals using regular logic, so getting functionally correct results is always possible. However, if your design requirements are flexible in terms of which control signals are used and in what priority, match your design to the target device architecture to achieve the most efficient results. If the priority of the signals in your design is not the same as that of the target architecture, then extra logic may be required to implement the control signals. This extra logic uses additional device resources and can cause additional delays for the control signals.

In addition, there are certain cases where using logic other than the dedicated control logic in the device architecture can have a larger impact. For example, the clock enable signal has priority over the synchronous reset or clear signal in the device architecture. The clock enable turns off the clock line in the LAB, and the clear signal is synchronous. So in the device architecture, the synchronous clear takes effect only when a clock edge occurs.

If you code a register with a synchronous clear signal that has priority over the clock enable signal, the software must emulate the clock enable functionality using data inputs to the registers. Because the signal does not use the clock enable port of a register, you cannot apply a Clock Enable Multicycle constraint. In this case, following the priority of signals available in the device is clearly the best choice for the priority of these control signals, and using a different priority causes unexpected results with an assignment to the clock enable signal.




The priority order for secondary control signals in Altera devices differs from the order for other vendors' devices. If your design requirements are flexible regarding priority, verify that the secondary control signals meet design performance requirements when migrating designs between FPGA vendors and try to match your target device architecture to achieve the best results.

The signal order is the same for all Altera device families, although as noted previously, not all device families provide every signal. The following priority order is observed:

1. Asynchronous Clear, `aclr`—highest priority
2. Preset, `pre`
3. Asynchronous Load, `aload`
4. Enable, `ena`
5. Synchronous Clear, `sclr`
6. Synchronous Load, `sload`
7. Data In, `data`—lowest priority

The following examples provide Verilog HDL and VHDL code that creates a register with the `aclr`, `aload`, and `ena` control signals.

 The Verilog HDL example (Example 6-32) does not have `adata` on the sensitivity list, but the VHDL example (Example 6-33) does. This is a limitation of the Verilog HDL language—there is no way to describe an asynchronous load signal (in which `q` toggles if `adata` toggles while `aload` is high). All synthesis tools should infer an `aload` signal from this construct despite this limitation. When they perform such inference, you may see information or warning messages from the synthesis tool.

Example 6-32. Verilog HDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals

```

module dff_control(clk, aclr, aload, ena, data, adata, q);
    input clk, aclr, aload, ena, data, adata;
    output q;

    reg q;

    always @ (posedge clk or posedge aclr or posedge aload)
    begin
        if (aclr)
            q <= 1'b0;
        else if (aload)
            q <= adata;
        else if (ena)
            q <= data;
    end
endmodule

```

Example 6-33. VHDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY dff_control IS
    PORT (
        clk: IN STD_LOGIC;
        aclr: IN STD_LOGIC;
        aload: IN STD_LOGIC;
        adata: IN STD_LOGIC;
        ena: IN STD_LOGIC;
        data: IN STD_LOGIC;
        q: OUT STD_LOGIC
    );
END dff_control;

ARCHITECTURE rtl OF dff_control IS
BEGIN
    PROCESS (clk, aclr, aload, adata)
    BEGIN
        IF (aclr = '1') THEN
            q <= '0';
        ELSIF (aload = '1') THEN
            q <= adata;
        ELSE
            IF (clk = '1' AND clk'event) THEN
                IF (ena = '1') THEN
                    q <= data;
                END IF;
            END IF;
        END IF;
    END PROCESS;
END rtl;

```

The preset signal is not available in many device families; therefore, it is not included in the examples.

Creating many registers with different `sload` and `sclr` signals can make packing the registers into LABs difficult for the Quartus II Fitter because the `sclr` and `sload` signals are LAB-wide signals. In addition, using the LAB-wide `sload` signal prevents the Fitter from packing registers using the quick feedback path in the device architecture, which means that some registers cannot be packed with other logic.

Synthesis tools typically restrict use of `sload` and `sclr` signals to cases in which there are enough registers with common signals to allow good LAB packing. Using the look-up table (LUT) to implement the signals is always more flexible if it is available. Because different device families offer different numbers of control signals, inference of these signals is also device-specific. For example, Stratix II devices have more flexibility than Stratix devices with respect to secondary control signals, so synthesis tools might infer more `sload` and `sclr` signals for Stratix II devices.

If you use these additional control signals, use them in the priority order that matches the device architecture. To achieve the most efficient results, ensure the `sclr` signal has a higher priority than the `sload` signal in the same way that `acclr` has higher priority than `aload` in the previous examples. Remember that the register signals are not inferred unless the design meets the conditions described previously. However, if your HDL described the desired behavior, the software always implements logic with the correct functionality.

In Verilog HDL, the following code for `sload` and `sclr` could replace the `if (ena) q <= data;` statements in the Verilog HDL example shown in [Example 6-32](#) (after adding the control signals to the module declaration).

Example 6-34. Verilog HDL `sload` and `sclr` Control Signals

```
if (ena) begin
    if (sclr)
        q <= 1'b0;
    else if (sload)
        q <= sdata;
    else
        q <= data;
end
```

In VHDL, the following code for `sload` and `sclr` could replace the `IF (ena = '1') THEN q <= data; END IF;` statements in the VHDL example shown in [Example 6-33](#) on [page 6-40](#) (after adding the control signals to the entity declaration).

Example 6-35. VHDL `sload` and `sclr` Control Signals

```
IF (ena = '1') THEN
    IF (sclr = '1') THEN
        q <= '0';
    ELSIF (sload = '1') THEN
        q <= sdata;
    ELSE
        q <= data;
    END IF;
END IF;
```

Latches

A latch is a small combinational loop that holds the value of a signal until a new value is assigned.



Altera recommends that you design without the use of latches whenever possible.



For additional information about the issues involved in designing with latches and combinational loops, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.

Latches can be inferred from HDL code when you did not intend to use a latch, as described in “[Unintentional Latch Generation](#)”. If you do intend to infer a latch, it is important to infer it correctly to guarantee correct device operation as detailed in “[Inferring Latches Correctly](#)” on page 6-43.

Unintentional Latch Generation

When you are designing combinational logic, certain coding styles can create an unintentional latch. For example, when `CASE` or `IF` statements do not cover all possible input conditions, latches may be required to hold the output if a new output value is not assigned. Check your synthesis tool messages for references to inferred latches. If your code unintentionally creates a latch, make code changes to remove the latch.



Latches have limited support in formal verification tools. Therefore, ensure that you do not infer latches unintentionally. For example, an incomplete `CASE` statement may create a latch when you are using formal verification in your design flow.

The `full_case` attribute can be used in Verilog HDL designs to treat unspecified cases as don't care values (X). However, using the `full_case` attribute can cause simulation mismatches because this attribute is a synthesis-only attribute, so simulation tools still treat the unspecified cases as latches.



Refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook* for more information about using attributes in your synthesis tool. The *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook* provides an example explaining possible simulation mismatches.

Omitting the final `else` or `when others` clause in an `if` or `case` statement can also generate a latch. Don't care (X) assignments on the default conditions are useful in preventing latch generation. For the best logic optimization, assign the default case or final `else` value to don't care (X) instead of a logic value.

The VHDL sample code shown in [Example 6-36](#) prevents unintentional latches. Without the final `else` clause, this code creates unintentional latches to cover the remaining combinations of the `sel` inputs. When you are targeting a Stratix device with this code, omitting the final `else` condition can cause the synthesis software to use up to six LEs, instead of the three it uses with the `else` statement. Additionally, assigning the final `else` clause to 1 instead of X can result in slightly more LEs because the synthesis software cannot perform as much optimization when you specify a constant value compared to a don't care value.

Example 6-36. VHDL Code Preventing Unintentional Latch Creation

```

LIBRARY ieee;
USE IEEE.std_logic_1164.all;

ENTITY nolatch IS
    PORT (a,b,c: IN STD_LOGIC;
          sel: IN STD_LOGIC_VECTOR (4 DOWNTO 0);
          oput: OUT STD_LOGIC);
END nolatch;

ARCHITECTURE rtl OF nolatch IS
BEGIN
    PROCESS (a,b,c,sel) BEGIN
        if sel = "00000" THEN
            oput <= a;
        ELSIF sel = "00001" THEN
            oput <= b;
        ELSIF sel = "00010" THEN
            oput <= c;
        ELSE
            --- Prevents latch inference
            oput <= 'X'; --/
        END if;
    END PROCESS;
END rtl;

```

Inferring Latches Correctly

Synthesis tools can infer a latch that does not exhibit the glitch and timing hazard problems typically associated with combinational loops.



Any use of latches generates warnings and is flagged if the design is migrated to a HardCopy ASIC. In addition, timing analysis does not completely model latch timing in some cases. Do not use latches unless you are very certain that your design requires it, and you fully understand the impact of using the latches.

When using Quartus II integrated synthesis, latches that are inferred by the software are reported in the **User-Specified and Inferred Latches** section of the Compilation Report. This report indicates whether the latch is considered safe and free of timing hazards.

If a latch or combinational loop in your design is not listed in the **User-Specified and Inferred Latches** section, it means that it was not inferred as a safe latch by the software and is not considered glitch-free.

All combinational loops listed in the **Analysis & Synthesis Logic Cells Representing Combinational Loops** table in the Compilation Report are at risk of timing hazards. These entries indicate possible problems with your design that you should investigate. However, it is possible to have a correct design that includes combinational loops. For example, it is possible that the combinational loop cannot be sensitized. This can occur in cases where there is an electrical path in the hardware, but either the designer knows that the circuit never encounters data that causes that path to be activated, or the surrounding logic is set up in a mutually exclusive manner that prevents that path from ever being sensitized, independent of the data input.

For macrocell-based devices such as MAX[®] 7000AE and MAX 3000A, all data (D-type) latches and set-reset (S-R) latches listed in the **Analysis & Synthesis User-Specified and Inferred Latches** table have an implementation free of timing hazards such as glitches. The implementation includes a cover term to ensure there is no glitching, and includes a single macrocell in the feedback loop.

For 4-input LUT-based devices such as Stratix devices, the Cyclone series, and MAX II devices, all latches in the **User-Specified and Inferred Latches** table with a single LUT in the feedback loop are free of timing hazards when a single input changes. Because of the hardware behavior of the LUT, the output does not glitch when a single input toggles between two values that are supposed to produce the same output value. For example, a D-type input toggling when the enable input is inactive, or a set input toggling when a reset input with higher priority is active. This hardware behavior of the LUT means that no cover term is required for a loop around a single LUT. The Quartus II software uses a single LUT in the feedback loop whenever possible. A latch that has data, enable, set, and reset inputs in addition to the output fed back to the input cannot be implemented in a single 4-input LUT. If the Quartus II software cannot implement the latch with a single-LUT loop because there are too many inputs, the **User-Specified and Inferred Latches** table indicates that the latch is not free of timing hazards.

For 6-input LUT-based devices, the software can implement all latch inputs with a single adaptive look-up table (ALUT) in the combinational loop. Therefore, all latches in the **User-Specified and Inferred Latches** table are free of timing hazards when a single input changes.

If a latch is listed as a safe latch, other Quartus II optimizations, such as physical synthesis netlist optimizations in the Fitter, maintain the hazard-free performance.

To ensure hazard-free behavior, only one control input may change at a time. Changing two inputs simultaneously, such as deasserting set and reset at the same time, or changing data and enable at the same time, can produce incorrect behavior in any latch.

Quartus II integrated synthesis infers latches from `always` blocks in Verilog HDL and `process` statements in VHDL, but not from continuous assignments in Verilog HDL or concurrent signal assignments in VHDL. These rules are the same as for register inference. The software infers registers or flipflops only from `always` blocks and `process` statements.

The Verilog HDL code sample shown in [Example 6-37](#) infers a S-R latch correctly in the Quartus II software.

Example 6-37. Verilog HDL Set-Reset Latch

```
module simple_latch (
    input SetTerm,
    input ResetTerm,
    output reg LatchOut
);

    always @ (SetTerm or ResetTerm) begin
        if (SetTerm)
            LatchOut = 1'b1
        else if (ResetTerm)
            LatchOut = 1'b0
        end
endmodule
```

The VHDL code sample shown in [Example 6-38](#) infers a D-type latch correctly in the Quartus II software.

Example 6-38. VHDL Data Type Latch

```
LIBRARY IEEE;
USE IEEE.std_logic_1164.all;

ENTITY simple_latch IS
PORT (
enable, data      : IN STD_LOGIC;
q                 : OUT STD_LOGIC
);
END simple_latch;

ARCHITECTURE rtl OF simple_latch IS
BEGIN

    latch : PROCESS (enable, data)
    BEGIN
    IF (enable = '1') THEN
    q <= data;
    END IF;
    END PROCESS latch;
END rtl;
```

The following example shows a Verilog HDL continuous assignment that does not infer a latch in the Quartus II software. (The behavior is similar to a latch, but it may not function correctly as a latch and its timing is not analyzed as a latch):

```
assign latch_out = (~en & latch_out) | (en & data);
```

Quartus II integrated synthesis also creates safe latches when possible for instantiations of the LPM_LATCH megafunction. You can use this megafunction to create a latch with any combination of data, enable, set, and reset inputs. The same limitations apply for creating safe latches as for inferring latches from HDL code.

Inferring the Altera LPM_LATCH function in another synthesis tool ensures that the implementation is also recognized as a latch in the Quartus II software. If a third-party synthesis tool implements a latch using the LPM_LATCH megafunction, then the Quartus II integrated synthesis lists the latch in the **User-Specified and Inferred Latches** table in the same way as it lists latches created in HDL source code. The coding style necessary to produce an LPM_LATCH implementation may depend on your synthesis tool. Some third-party synthesis tools list the number of LPM_LATCH functions that are inferred.

For LUT-based families, the Fitter uses global routing for control signals including signals that Analysis and Synthesis identifies as latch enables. In some cases the global insertion delay may decrease the timing performance. If necessary, you can turn off the **Quartus II Global Signal** logic option to manually prevent the use of global signals. Global latch enables are listed in the **Global & Other Fast Signals** table in the Compilation Report.

General Coding Guidelines

This section helps you understand how synthesis tools map various types of HDL code into the target Altera device. Following Altera recommended coding styles, and in some cases designing logic structures to match the appropriate device architecture, can provide significant improvements in the design's quality of results.

This section provides coding guidelines for the following logic structures:

- **“Tri-State Signals”**. This section explains how to create tri-state signals for bidirectional I/O pins.
- **“Clock Multiplexing” on page 6-47**. This section provides recommendations for multiplexing clock signals.
- **“Adder Trees” on page 6-51**. This section explains the different coding styles that lead to optimal results for devices with 4-input look-up tables and 6-input ALUTs.
- **“State Machines” on page 6-53**. This section helps ensure the best results when you use state machines.
- **“Multiplexers” on page 6-60**. This section explains how multiplexers can be synthesized for 4-input LUT devices, addresses common problems, and provides guidelines to achieve optimal resource utilization.
- **“Cyclic Redundancy Check Functions” on page 6-68**. This section provides guidelines for getting good results when designing Cyclic Redundancy Check (CRC) functions.
- **“Comparators” on page 6-69**. This section explains different comparator implementations and provides suggestions for controlling the implementation.
- **“Counters” on page 6-71**. This section provides guidelines to ensure your counter design targets the device architecture optimally.

Tri-State Signals

When you target Altera devices, you should use tri-state signals only when they are attached to top-level bidirectional or output pins. Avoid lower level bidirectional pins, and avoid using the Z logic value unless it is driving an output or bidirectional pin.

Synthesis tools implement designs with internal tri-state signals correctly in Altera devices using multiplexer logic, but Altera does not recommend this coding practice.



In hierarchical block-based or incremental design flows, a hierarchical boundary cannot contain any bidirectional ports, unless the lower level bidirectional port is connected directly through the hierarchy to a top-level output pin without connecting to any other design logic. If you use boundary tri-states in a lower level block, synthesis software must push the tri-states through the hierarchy to the top-level to make use of the tri-state drivers on output pins of Altera devices. Because pushing tri-states requires optimizing through hierarchies, lower level tri-states are restricted with block-based design methodologies.

The code examples shown in [Example 6-39](#) and [Example 6-40](#) show Verilog HDL and VHDL code that creates tri-state bidirectional signals.

Example 6-39. Verilog HDL Tri-State Signal

```
module tristate (myinput, myenable, mybidir);
    input myinput, myenable;
    inout mybidir;
    assign mybidir = (myenable ? myinput : 1'bZ);
endmodule
```

Example 6-40. VHDL Tri-State Signal

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_arith.all;

ENTITY tristate IS
PORT (
    mybidir : INOUT STD_LOGIC;
    myinput : IN STD_LOGIC;
    myenable : IN STD_LOGIC
);
END tristate;

ARCHITECTURE rtl OF tristate IS
BEGIN
    mybidir <= 'Z' WHEN (myenable = '0') ELSE myinput;
END rtl;
```

Clock Multiplexing

Clock multiplexing is sometimes used to operate the same logic function with different clock sources. This type of logic can introduce glitches that create functional problems, and the delay inherent in the combinational logic can lead to timing problems. Clock multiplexers trigger warnings from a wide range of design rule check and timing analysis tools.


Altera recommends using dedicated hardware to perform clock multiplexing when it is available, instead of using multiplexing logic. For example, you can use the Clock Switchover feature or the Clock Control Block available in certain Altera devices. These dedicated hardware blocks avoid glitches, ensure that you use global low-skew routing lines, and avoid any possible hold time problems on the device due to logic delay on the clock line. Many Altera devices also support dynamic PLL reconfiguration, which is the safest and most robust method of changing clock rates during device operation.

Refer to the appropriate device data sheet or handbook for device-specific information about clocking structures. Also refer to the *ALTCLKCTRL Megafunction User Guide*, the *ALTPLL Megafunction User Guide*, and the *Phase-Locked Loops Reconfiguration (ALTPLL_RECONFIG) Megafunction User Guide*.

If you implement a clock multiplexer in logic cells because the design has too many clocks to use the clock control block, or if dynamic reconfiguration is too complex for your design, it is important to consider simultaneous toggling inputs and ensure glitch-free transitions.

Figure 6-2 shows a simple representation of a clock multiplexer (mux) in a device with 6-input LUTs.


Figure 6-2. Simple Clock Multiplexer in a 6-Input LUT



The data sheet for your target device describes how LUT outputs may glitch during a simultaneous toggle of input signals, independent of the LUT function. Although in practice the 4:1 MUX function does not generate detectable glitches during simultaneous data input toggles, it is possible to construct cell implementations that do exhibit significant glitches, so this simple clock mux structure is not recommended. An additional problem with this implementation is that the output behaves erratically during a change in the `clk_select` signals. This behavior could create timing violations on all registers fed by the system clock and result in possible metastability.

A more sophisticated clock select structure can eliminate the simultaneous toggle and switching problems, as shown in Figure 6-3.

Figure 6-3. Glitch-Free Clock Multiplexer Structure



This structure can be generalized for any number of clock channels. [Example 6-41](#) contains a parameterized version in Verilog HDL. The design enforces that no clock activates until all others have been inactive for at least a few cycles, and that activation occurs while the clock is low. The design applies a `synthesis_keep` directive to the AND gates on the right side of the figure, which ensures there are no simultaneous toggles on the input of the `clk_out` OR gate.

It is important to note that switching from clock A to clock B requires that clock A continue to operate for at least a few cycles. If the old clock stops immediately, the design sticks. The select signals are implemented as a “one-hot” control in this example, but you can use other encoding if you prefer. The input side logic is asynchronous and is not critical. This design can tolerate extreme glitching during the switch process.

Example 6-41. Verilog HDL Clock Multiplexing Design to Avoid Glitches

```

module clock_mux (clk,clk_select,clk_out);

parameter num_clocks = 4;

input [num_clocks-1:0] clk;
input [num_clocks-1:0] clk_select; // one hot
output clk_out;

genvar i;

reg [num_clocks-1:0] ena_r0;
reg [num_clocks-1:0] ena_r1;
reg [num_clocks-1:0] ena_r2;
wire [num_clocks-1:0] qualified_sel;

// A look-up-table (LUT) can glitch when multiple inputs
// change simultaneously. Use the keep attribute to
// insert a hard logic cell buffer and prevent
// the unrelated clocks from appearing on the same LUT.

wire [num_clocks-1:0] gated_clks /* synthesis keep */;

initial begin
ena_r0 = 0;
ena_r1 = 0;
ena_r2 = 0;
end

generate
for (i=0; i<num_clocks; i=i+1)
begin : lp0
wire [num_clocks-1:0] tmp_mask;
assign tmp_mask = {num_clocks{1'b1}} ^ (1 << i);

assign qualified_sel[i] = clk_select[i] &
(~|(ena_r2 & tmp_mask));

always @(posedge clk[i]) begin
ena_r0[i] <= qualified_sel[i];
ena_r1[i] <= ena_r0[i];
end

always @(negedge clk[i]) begin
ena_r2[i] <= ena_r1[i];
end

assign gated_clks[i] = clk[i] & ena_r2[i];
end
endgenerate

// These will not exhibit simultaneous toggle by construction
assign clk_out = |gated_clks;

endmodule

```

Adder Trees

Structuring adder trees appropriately to match your targeted Altera device architecture can result in significant performance and density improvements. A good example of an application using a large adder tree is a finite impulse response (FIR) correlator. Using a pipelined binary or ternary adder tree appropriately can greatly improve the quality of your results.

This section explains why coding recommendations are different for Altera 4-input LUT devices and 6-input LUT devices.

Architectures with 4-Input LUTs in Logic Elements

Architectures such as Stratix devices and the Cyclone series, APEX series, and FLEX series of devices contain 4-input LUTs as the standard combinational structure in the LE.

If your design can tolerate pipelining, the fastest way to add three numbers A, B, and C in devices that use 4-input lookup tables is to add A + B, register the output, and then add the registered output to C. Adding A + B takes one level of logic (one bit is added in one LE), so this runs at full clock speed. This can be extended to as many numbers as desired.

[Example 6-42](#) shows five numbers A, B, C, D, and E are added. Adding five numbers in devices that use 4-input lookup tables requires four adders and three levels of registers for a total of 64 LEs (for 16-bit numbers).

Example 6-42. Verilog HDL Pipelined Binary Tree

```
module binary_adder_tree (a, b, c, d, e, clk, out);
    parameter width = 16;
    input [width-1:0] a, b, c, d, e;
    input clk;
    output [width-1:0] out;

    wire [width-1:0] sum1, sum2, sum3, sum4;
    reg [width-1:0] sumreg1, sumreg2, sumreg3, sumreg4;
    // Registers

    always @ (posedge CLK)
        begin
            sumreg1 <= sum1;
            sumreg2 <= sum2;
            sumreg3 <= sum3;
            sumreg4 <= sum4;
        end

    // 2-bit additions
    assign sum1 = A + B;
    assign sum2 = C + D;
    assign sum3 = sumreg1 + sumreg2;
    assign sum4 = sumreg3 + E;
    assign out = sumreg4;
endmodule
```

Architectures with 6-Input LUTs in Adaptive Logic Modules

High-performance Altera device families use a 6-input LUT in their basic logic structure, so these devices benefit from a different coding style from the previous example presented for 4-input LUTs. Specifically, in these devices, ALMs can simultaneously add three bits. Therefore, the tree in [Example 6-42](#) must be two levels deep and contain just two add-by-three inputs instead of four add-by-two inputs.

Although the code in the previous example compiles successfully for 6-input LUT devices, the code is inefficient and does not take advantage of the 6-input adaptive ALUT. By restructuring the tree as a ternary tree, the design becomes much more efficient, significantly improving density utilization. Therefore, when you are targeting with ALUTs and ALMs, large pipelined binary adder trees designed for 4-input LUT architectures should be rewritten to take advantage of the advanced device architecture.

[Example 6-43](#) uses just 32 ALUTs in a Stratix II device—more than a 4:1 advantage over the number of LUTs in the prior example implemented in a Stratix device.



You cannot pack a LAB full when using this type of coding style because of the number of LAB inputs. However, in a typical design, the Quartus II Fitter can pack other logic into each LAB to take advantage of the unused ALMs.

Example 6-43. Verilog HDL Pipelined Ternary Tree

```
module ternary_adder_tree (a, b, c, d, e, clk, out);
    parameter width = 16;
    input [width-1:0] a, b, c, d, e;
    input clk;
    output [width-1:0] out;

    wire [width-1:0] sum1, sum2;
    reg [width-1:0] sumreg1, sumreg2;
    // registers

    always @ (posedge clk)
        begin
            sumreg1 <= sum1;
            sumreg2 <= sum2;
        end

    // 3-bit additions
    assign sum1 = a + b + c;
    assign sum2 = sumreg1 + d + e;
    assign out = sumreg2;
endmodule
```

These examples show pipelined adders, but partitioning your addition operations can help you achieve better results in nonpipelined adders as well. If your design is not pipelined, a ternary tree provides much better performance than a binary tree. For example, depending on your synthesis tool, the HDL code $sum = (A + B + C) + (D + E)$ is more likely to create the optimal implementation of a 3-input adder for $A + B + C$ followed by a 3-input adder for $sum1 + D + E$ than the code without the parentheses. If you do not add the parentheses, the synthesis tool may partition the addition in a way that is not optimal for the architecture.

State Machines

Synthesis tools can recognize and encode Verilog HDL and VHDL state machines during synthesis. This section presents guidelines to ensure the best results when you use state machines. Ensuring that your synthesis tool recognizes a piece of code as a state machine allows the tool to recode the state variables to improve the quality of results, and allows the tool to use the known properties of state machines to optimize other parts of the design. When synthesis recognizes a state machine, it is often able to improve the design area and performance.

To achieve the best results on average, synthesis tools often use one-hot encoding for FPGA devices and minimal-bit encoding for CPLD devices, although the choice of implementation can vary for different state machines and different devices. Refer to your synthesis tool documentation for specific ways to control the manner in which state machines are encoded.



For information about state machine encoding in Quartus II integrated synthesis, refer to the *State Machine Processing* section in the [Quartus II Integrated Synthesis](#) chapter in volume 1 of the *Quartus II Handbook*.


To ensure proper recognition and inference of state machines and to improve the quality of results, Altera recommends that you observe the following guidelines, which apply to both Verilog HDL and VHDL:

- Assign default values to outputs derived from the state machine so that synthesis does not generate unwanted latches.
- Separate the state machine logic from all arithmetic functions and data paths, including assigning output values.
- If your design contains an operation that is used by more than one state, define the operation outside the state machine and cause the output logic of the state machine to use this value.
- Use a simple asynchronous or synchronous reset to ensure a defined power-up state. If your state machine design contains more elaborate reset logic, such as both an asynchronous reset and an asynchronous load, the Quartus II software generates regular logic rather than inferring a state machine.

If a state machine enters an illegal state due to a problem with the device, the design likely ceases to function correctly until the next reset of the state machine. Synthesis tools do not provide for this situation by default. The same issue applies to any other registers if there is some kind of fault in the system. A `default` or `when others` clause does not affect this operation, assuming that your design never deliberately enters this state. Synthesis tools remove any logic generated by a default state if it is not reachable by normal state machine operation.

Many synthesis tools (including Quartus II integrated synthesis) have an option to implement a safe state machine. The software inserts extra logic to detect an illegal state and force the state machine's transition to the reset state. It is commonly used when the state machine can enter an illegal state. The most common cause of this situation is a state machine that has control inputs that come from another clock domain, such as the control logic for a dual-clock FIFO.

This option protects only state machines. All other registers in the design are not protected this way.

 For additional information about tool-specific options for implementing state machines, refer to the tool vendor's documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.


The following two sections, “[Verilog HDL State Machines](#)” and “[VHDL State Machines](#)” on page 6-58, describe additional language-specific guidelines and coding examples.

Verilog HDL State Machines

To ensure proper recognition and inference of Verilog HDL state machines, observe the following additional Verilog HDL guidelines. Some of these guidelines may be specific to Quartus II integrated synthesis. Refer to your synthesis tool documentation for specific coding recommendations.

If the state machine is not recognized and inferred by the synthesis software (such as Quartus II integrated synthesis), the state machine is implemented as regular logic gates and registers and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

- If you are using the SystemVerilog standard, use enumerated types to describe state machines (as shown in the “[SystemVerilog State Machine Coding Example](#)” on page 6-57).
- Represent the states in a state machine with the parameter data types in Verilog-1995 and -2001 and use the parameters to make state assignments (as shown below in the “[Verilog-2001 State Machine Coding Example](#)”). This implementation makes the state machine easier to read and reduces the risk of errors during coding.

 Altera recommends against the direct use of integer values for state variables such as `next_state <= 0`. However, using an integer does not prevent inference in the Quartus II software.

- No state machine is inferred in the Quartus II software if the state transition logic uses arithmetic similar to that shown in the following example:

```
case (state)
  0: begin
    if (ena) next_state <= state + 2;
    else next_state <= state + 1;
  end
  1: begin
    ...
  end
endcase
```

- No state machine is inferred in the Quartus II software if the state variable is an output.
- No state machine is inferred in the Quartus II software for signed variables.

Verilog-2001 State Machine Coding Example

The following module `verilog_fsm` is an example of a typical Verilog HDL state machine implementation (Example 6-44).

This state machine has five states. The asynchronous reset sets the variable `state` to `state_0`. The sum of `in_1` and `in_2` is an output of the state machine in `state_1` and `state_2`. The difference (`in_1 - in_2`) is also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in_1` and `in_2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

Example 6-44. Verilog-2001 State Machine

```

module verilog_fsm (clk, reset, in_1, in_2, out);
    input clk, reset;
    input [3:0] in_1, in_2;
    output [4:0] out;
    parameter state_0 = 3'b000;
    parameter state_1 = 3'b001;
    parameter state_2 = 3'b010;
    parameter state_3 = 3'b011;
    parameter state_4 = 3'b100;

    reg [4:0] tmp_out_0, tmp_out_1, tmp_out_2;
    reg [2:0] state, next_state;

    always @ (posedge clk or posedge reset)
    begin
        if (reset)
            state <= state_0;
        else
            state <= next_state;
    end
    always @ (state or in_1 or in_2)
    begin
        tmp_out_0 = in_1 + in_2;
        tmp_out_1 = in_1 - in_2;
        case (state)
            state_0: begin
                tmp_out_2 <= in_1 + 5'b00001;
                next_state <= state_1;
            end
            state_1: begin
                if (in_1 < in_2) begin
                    next_state <= state_2;
                    tmp_out_2 <= tmp_out_0;
                end
                else begin
                    next_state <= state_3;
                    tmp_out_2 <= tmp_out_1;
                end
            end
            state_2: begin
                tmp_out_2 <= tmp_out_0 - 5'b00001;
                next_state <= state_3;
            end
            state_3: begin
                tmp_out_2 <= tmp_out_1 + 5'b00001;
                next_state <= state_0;
            end
            state_4: begin
                tmp_out_2 <= in_2 + 5'b00001;
                next_state <= state_0;
            end
            default: begin
                tmp_out_2 <= 5'b00000;
                next_state <= state_0;
            end
        endcase
        assign out = tmp_out_2;
    end
endmodule

```

An equivalent implementation of this state machine can be achieved by using ``define` instead of the parameter data type, as follows:

```
`define state_0 3'b000
`define state_1 3'b001
`define state_2 3'b010
`define state_3 3'b011
`define state_4 3'b100
```

In this case, the `state` and `next_state` assignments are assigned a `'state_x` instead of a `state_x`, as shown in the following example:

```
next_state <= `state_3;
```



Although the ``define` construct is supported, Altera strongly recommends the use of the parameter data type because doing so preserves the state names throughout synthesis.

SystemVerilog State Machine Coding Example

The module `enum_fsm` shown in [Example 6-45](#) is an example of a SystemVerilog state machine implementation that uses enumerated types. Altera recommends using this coding style to describe state machines in SystemVerilog.



In Quartus II integrated synthesis, the enumerated type that defines the states for the state machine must be of an unsigned integer type as shown in [Example 6-45](#). If you do not specify the enumerated type as `int unsigned`, a signed `int` type is used by default. In this case, the Quartus II integrated synthesis synthesizes the design, but does not infer or optimize the logic as a state machine.

Example 6-45. SystemVerilog State Machine Using Enumerated Types

```

module enum_fsm (input clk, reset, input int data[3:0], output int o);
enum int unsigned { S0 = 0, S1 = 2, S2 = 4, S3 = 8 } state, next_state;

always_comb begin : next_state_logic
    next_state = S0;
    case(state)
        S0: next_state = S1;
        S1: next_state = S2;
        S2: next_state = S3;
        S3: next_state = S3;
    endcase
end

always_comb begin
    case(state)
        S0: o = data[3];
        S1: o = data[2];
        S2: o = data[1];
        S3: o = data[0];
    endcase
end

always_ff@(posedge clk or negedge reset) begin
    if (~reset)
        state <= S0;
    else
        state <= next_state;
    end
end
endmodule

```

VHDL State Machines

To ensure proper recognition and inference of VHDL state machines, represent the states in a state machine with enumerated types and use the corresponding types to make state assignments. This implementation makes the state machine easier to read and reduces the risk of errors during coding. If the state is not represented by an enumerated type, synthesis software (such as Quartus II integrated synthesis) does not recognize the state machine. Instead, the state machine is implemented as regular logic gates and registers and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

VHDL State Machine Coding Example

The following entity, `vhd1_fsm`, is an example of a typical VHDL state machine implementation ([Example 6-46](#)).

This state machine has five states. The asynchronous reset sets the variable `state` to `state_0`. The sum of `in1` and `in2` is an output of the state machine in `state_1` and `state_2`. The difference (`in1 - in2`) is also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in1` and `in2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

Example 6-46. VHDL State Machine (Part 1 of 2)

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY vhdl_fsm IS
    PORT(
        clk: IN STD_LOGIC;
        reset: IN STD_LOGIC;
        in1: IN UNSIGNED(4 downto 0);
        in2: IN UNSIGNED(4 downto 0);
        out_1: OUT UNSIGNED(4 downto 0)
    );
END vhdl_fsm;

ARCHITECTURE rtl OF vhdl_fsm IS
    TYPE Tstate IS (state_0, state_1, state_2, state_3, state_4);
    SIGNAL state: Tstate;
    SIGNAL next_state: Tstate;

BEGIN
    PROCESS(clk, reset)
    BEGIN
        IF reset = '1' THEN
            state <= state_0;
        ELSIF rising_edge(clk) THEN
            state <= next_state;
        END IF;
    END PROCESS;

    PROCESS (state, in1, in2)
        VARIABLE tmp_out_0: UNSIGNED (4 downto 0);
        VARIABLE tmp_out_1: UNSIGNED (4 downto 0);
```

Example 6-46. VHDL State Machine (Part 2 of 2)

```

BEGIN
  tmp_out_0 := in1 + in2;
  tmp_out_1 := in1 - in2;
  CASE state IS
    WHEN state_0 =>
      out_1 <= in1;
      next_state <= state_1;
    WHEN state_1 =>
      IF (in1 < in2) then
        next_state <= state_2;
        out_1 <= tmp_out_0;
      ELSE
        next_state <= state_3;
        out_1 <= tmp_out_1;
      END IF;
    WHEN state_2 =>
      IF (in1 < "0100") then
        out_1 <= tmp_out_0;
      ELSE
        out_1 <= tmp_out_1;
      END IF;
      next_state <= state_3;

    WHEN state_3 =>
      out_1 <= "11111";
      next_state <= state_4;
    WHEN state_4 =>
      out_1 <= in2;
      next_state <= state_0;
    WHEN OTHERS =>
      out_1 <= "00000";
      next_state <= state_0;
  END CASE;
END PROCESS;
END rtl;

```

Multiplexers


Multiplexers form a large portion of the logic utilization in many FPGA designs. By optimizing your multiplexer logic, you ensure the most efficient implementation in your Altera device. This section addresses common problems and provides design guidelines to achieve optimal resource utilization for multiplexer designs. The section also describes various types of multiplexers, and how they are implemented in the 4-input LUT found in many FPGA architectures, such as Altera's Stratix devices.



Stratix II and other high-performance devices have 6-input ALUTs and are not specifically addressed here. Although many of the principles and techniques for optimization are similar, device utilization differs in the 6-input LUT devices. For example, these devices can implement wider multiplexers in one ALM than can be implemented in the 4-input LUT of an LE.

Quartus II Software Option for Multiplexer Restructuring

Quartus II integrated synthesis provides the **Restructure Multiplexers** logic option that extracts and optimizes buses of multiplexers during synthesis. In certain situations, this option automatically performs some of the recoding functions described in this section without changing the HDL code in your design. This option is on by default, when the **Optimization technique** is set to **Balanced** (the default for most device families) or set to **Area**.

 For details, refer to the *Restructure Multiplexers* subsection in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

Even with this Quartus II-specific option turned on, it is beneficial to understand how your coding style can be interpreted by your synthesis tool, and avoid the situations that can cause problems in your design.

Multiplexer Types

This subsection addresses how multiplexers are created from various types of HDL code. CASE statements, IF statements, and state machines are all common sources of multiplexer logic in designs. These HDL structures create different types of multiplexers including binary multiplexers, selector multiplexers, and priority multiplexers. Understanding how multiplexers are created from HDL code and how they might be implemented during synthesis is the first step towards optimizing multiplexer structures for best results.

Binary Multiplexers

Binary multiplexers select inputs based on binary-encoded selection bits.

[Example 6-47](#) shows Verilog HDL code for two ways to describe a simple 4:1 binary multiplexer.

Example 6-47. Verilog HDL Binary-Encoded Multiplexers

```
case (sel)
  2'b00: z = a;
  2'b01: z = b;
  2'b10: z = c;
  2'b11: z = d;
endcase
```

A 4:1 binary multiplexer is efficiently implemented by using two 4-input LUTs. Larger binary multiplexers can be constructed that use the 4:1 multiplexer; constructing an N-input multiplexer (N:1 multiplexer) from a tree of 4:1 multiplexers can result in a structure using as few as $0.66 \cdot (N - 1)$ LUTs.

Selector Multiplexers

Selector multiplexers have a separate select line for each data input. The select lines for the multiplexer are one-hot encoded. [Example 6-48](#) shows a simple Verilog HDL code example describing a one-hot selector multiplexer.

Example 6-48. Verilog HDL One-Hot-Encoded Case Statement

```

case (sel)
  4'b0001: z = a;
  4'b0010: z = b;
  4'b0100: z = c;
  4'b1000: z = d;
  default: z = 1'bx;
endcase

```

Selector multiplexers are commonly built as a tree of AND and OR gates. Using this scheme, two inputs can be selected using two select lines in a single 4-input LUT that uses two AND gates and an OR gate. The outputs of these LUTs can be combined with a wide OR gate. An N-input selector multiplexer of this structure requires at least $0.66 \cdot (N-0.5)$ LUTs, which is just slightly worse than the best binary multiplexer.

Priority Multiplexers

In priority multiplexers, the select logic implies a priority. The options to select the correct item must be checked in a specific order based on signal priority. These structures commonly are created from IF, ELSE, WHEN, SELECT, and `? :` statements in VHDL or Verilog HDL. The example VHDL code in [Example 6-49](#) probably result in the schematic implementation illustrated in [Figure 6-4](#).


Example 6-49. VHDL IF Statement Implying Priority

```

IF cond1 THEN z <= a;
ELSIF cond2 THEN z <= b;
ELSIF cond3 THEN z <= c;
ELSE z <= d;
END IF;

```

The multiplexers shown in [Figure 6-4](#) form a chain, evaluating each condition or select bit, one at a time.

Figure 6-4. Priority Multiplexer Implementation of an IF Statement

An N-input priority multiplexer uses a LUT for every 2:1 multiplexer in the chain, requiring N-1 LUTs. This chain of multiplexers generally increases delay because the critical path through the logic traverses every multiplexer in the chain.

To improve the timing delay through the multiplexer, avoid priority multiplexers if priority is not required. If the order of the choices is not important to the design, use a CASE statement to implement a binary or selector multiplexer instead of a priority multiplexer. If delay through the structure is important in a multiplexed design requiring priority, consider recoding the design to reduce the number of logic levels to minimize delay, especially along your critical paths.

Default or Others Case Assignment

To fully specify the cases in a CASE statement, include a `default` (Verilog HDL) or `OTHERS` (VHDL) assignment. This assignment is especially important in one-hot encoding schemes where many combinations of the select lines are unused. Specifying a case for the unused select line combinations gives the synthesis tool information about how to synthesize these cases, and is required by the Verilog HDL and VHDL language specifications.

Some designs do not require that the outcome in the unused cases be considered, often because designers assume these cases will not occur. For these types of designs, you can choose any value for the `default` or `OTHERS` assignment. However, be aware that the assignment value you choose can have a large effect on the logic utilization required to implement the design due to the different ways synthesis tools treat different values for the assignment, and how the synthesis tools use different speed and area optimizations.

In general, to obtain best results, explicitly define invalid CASE selections with a separate `default` or `OTHERS` statement instead of combining the invalid cases with one of the defined cases.

If the value in the invalid cases is not important, specify those cases explicitly by assigning the X (don't care) logic value instead of choosing another value. This assignment allows your synthesis tool to perform the best area optimizations.

You can experiment with different `default` or `OTHERS` assignments for your HDL design and your synthesis tool to test the effect they have on logic utilization in your design.

Implicit Defaults

The IF statements in Verilog HDL and VHDL can be a convenient way to specify conditions that do not easily lend themselves to a CASE-type approach. However, using IF statements can result in complicated multiplexer trees that are not easy for synthesis tools to optimize.

In particular, every IF statement has an implicit ELSE condition, even when it is not specified. These implicit defaults can cause additional complexity in a multiplexed design.

The code in [Example 6-50](#) represents a multiplexer with four inputs (a, b, c, d) and one output (z). Altera does not recommend using this coding style.

Example 6-50. VHDL IF Statement with Implicit Defaults

```

IF cond1 THEN
  IF cond2 THEN
    z <= a;
  END IF;
ELSIF cond3 THEN
  IF cond4 THEN
    z <= b;
  ELSIF cond5 THEN
    z <= c;
  END IF;
ELSIF cond6 THEN
  z <= d;
END IF;

```

Although the code appears to implement a 4:1 multiplexer, each of the three separate IF statements in the code has an implicit ELSE condition that is not specified. Because the output values for the ELSE cases are not specified, the synthesis tool assumes the intent is to maintain the same output value for these cases and infers a combinational loop, such as a latch. Latches add to the design's logic utilization and can also make timing analysis difficult and lead to other problems.

The code sample shown in [Example 6-51](#) shows code with the same functionality as the code shown in [Example 6-50](#), but specifies the ELSE cases explicitly. (This is not a recommended coding style improvement, but it explicitly shows the default conditions from the previous example.)

Example 6-51. VHDL IF Statement with Default Conditions Explicitly Specified


```

IF cond1 THEN
  IF cond2 THEN
    z <= a;
  ELSE
    z <= z;
  END IF;
ELSIF cond3 THEN
  IF cond4 THEN
    z <= b;
  ELSIF cond5 THEN
    z <= c;
  ELSE
    z <= z;
  END IF;
ELSIF cond6 THEN
  z <= d;
ELSE
  z <= z;
END IF;

```

[Figure 6-5](#) is a schematic representing the code in [Example 6-51](#), which illustrates that the multiplexer logic is significantly more complicated than a basic 4:1 multiplexer, although there are only four inputs.

Figure 6-5. Multiplexer Implementation of an IF Statement with Implicit Defaults



There are several ways you can simplify the multiplexed logic and remove the unrequired defaults. The optimal method may be to recode the design so the logic takes the structure of a 4:1 CASE statement. Alternatively, if priority is important, you can restructure the code to deduce default cases and flatten the multiplexer. In this example, instead of IF cond1 THEN IF cond2, use IF (cond1 AND cond2), which performs the same function. In addition, examine whether the defaults are don't care cases. In this example, you can promote the last ELSIF cond6 statement to an ELSE statement if no other valid cases can occur.

Avoid unnecessary default conditions in your multiplexer logic to reduce the complexity and logic utilization required to implement your design.

Degenerate Multiplexers


A degenerate multiplexer is a multiplexer in which not all of the possible cases are used for unique data inputs. The cases that are not required tend to contribute to inefficiency in the logic utilization for these multiplexers. You can recode degenerate multiplexers so they take advantage of the efficient logic utilization possible with full binary multiplexers.

Example 6-52 shows a VHDL CASE statement describing a degenerate multiplexer.


Example 6-52. VHDL CASE Statement Describing a Degenerate Multiplexer

```
CASE sel[3:0] IS
  WHEN "0101" => z <= a;
  WHEN "0111" => z <= b;
  WHEN "1010" => z <= c;
  WHEN OTHERS => z <= d;
END CASE;
```

The number of select lines in a binary multiplexer normally dictates the size of a multiplexer required to implement the desired function. For example, the multiplexer structure represented in Figure 6-6 has four select lines capable of implementing a binary multiplexer with 16 inputs. However, the design does not use all 16 inputs, which makes this multiplexer a degenerate 16:1 multiplexer.

Figure 6-6. Binary Degenerate Multiplexer

In [Figure 6-6](#), the first and fourth multiplexers in the top level can easily be eliminated because all four inputs to each multiplexer are the same value, and the number of inputs to the other multiplexers can be reduced, as shown in [Figure 6-7](#).

Figure 6-7. Optimized Version of the Degenerate Binary Multiplexer

Implementing this version of the multiplexer still requires at least five 4-input LUTs, two for each of the remaining 3:1 multiplexers and one for the 2:1 multiplexer. This design selects an output from only four inputs, a 4:1 binary multiplexer can be implemented optimally in two LUTs, so this degenerate multiplexer tree reduces the efficiency of the logic.

You can improve logic utilization of this structure by recoding the select lines to implement a full 4:1 binary multiplexer. The code sample shown in [Example 6-53](#) shows a recoder design that translates the original select lines into the `z_sel` signal with binary encoding.

Example 6-53. VHDL Recoder Design for Degenerate Binary Multiplexer

```
CASE sel[3:0] IS
  WHEN "0101" => z_sel <= "00";
  WHEN "0111" => z_sel <= "01";
  WHEN "1010" => z_sel <= "10";
  WHEN OTHERS => z_sel <= "11";
END CASE;
```

The code sample shown in [Example 6-54](#) shows you how to implement the full binary multiplexer.

Example 6-54. VHDL 4:1 Binary Multiplexer Design


```

CASE z_sel[1:0] IS
  WHEN "00" => z <= a;
  WHEN "01" => z <= b;
  WHEN "10" => z <= c;
  WHEN "11" => z <= d;
END CASE;

```

Use the new `z_sel` control signal from the recoder design to control the 4:1 binary multiplexer that chooses between the four inputs `a`, `b`, `c`, and `d`, as illustrated in [Figure 6-8](#). The complexity of the select lines is handled in the recoder design, and the data multiplexing is performed with simple binary select lines enabling the most efficient implementation.

Figure 6-8. Binary Multiplexer with Recoder



The design for the recoder can be implemented in two LUTs and the efficient 4:1 binary multiplexer uses two LUTs, for a total of four LUTs. The original degenerate multiplexer required five LUTs, so the recoded version uses 20% less logic than the original.

You can often improve the logic utilization of multiplexers by recoding the select lines into full binary cases. Although logic is required to perform the encoding, the overall logic utilization is often improved.

Buses of Multiplexers

The inputs to multiplexers are often data input buses in which the same multiplexer function is performed on a set of data input buses. In these cases, any inefficiency in the multiplexer is multiplied by the number of bits in the bus. The issues described in the previous sections become even more important for wide multiplexer buses.

For example, the recoding of select lines into full binary cases detailed in the previous section can often be used in multiplexed buses. Recoding the select lines may have to be completed only once for all the multiplexers in the bus. By sharing the recoder logic among all the bits in the bus, you can greatly improve the logic efficiency of a bus of multiplexers.

The degenerate multiplexer in the previous section requires five LUTs to implement. If the inputs and output are 32 bits wide, the function could require 32×5 or 160 LUTs for the whole bus. The recoder design uses only two LUTs, and the select lines only have to be recoded once for the entire bus. The binary 4:1 multiplexer requires two LEs per bit of the bus. The total logic utilization for the recoded version could be $2 + (2 \times 32)$ or 66 LUTs for the whole bus, compared to 160 LUTs for the original version. The logic savings become more important with wide multiplexer buses.

Using techniques to optimize degenerate multiplexers, removing implicit defaults are not required, and choosing the optimal `DEFAULT` or `OTHERS` case can play an important role when optimizing buses of multiplexers.

Cyclic Redundancy Check Functions

CRC computations are used heavily by communications protocols and storage devices to detect any corruption of the data. These functions are highly effective; there is a very low probability that corrupted data can pass a 32-bit CRC check.

CRC functions typically use wide XOR gates to compare the data. The way that synthesis tools flatten and factor these XOR gates to implement the logic in FPGA LUTs can greatly impact the area and performance results for the design. XOR gates have a cancellation property which creates an exceptionally large number of reasonable factoring combinations, so synthesis tools cannot always choose the best result by default.

The 6-input ALUT has a significant advantage over 4-input LUTs for these designs. When properly synthesized, CRC processing designs can run at high speeds in devices with 6-input ALUTs.

The following guidelines help you improve the quality of results for CRC designs in Altera devices.

If Performance is Important, Optimize for Speed

Synthesis tools flatten XOR gates to minimize area and depth of levels of logic. Synthesis tools such as Quartus II integrated synthesis target area optimization by default for these logic structures. Therefore, for more focus on depth reduction, set the synthesis optimization technique to speed.

Flattening for depth sometimes causes a significant increase in area.

Use Separate CRC Blocks Instead of Cascaded Stages

Some designers optimize their CRC designs to use cascaded stages, for example, four stages of 8 bits. In such designs, intermediate calculations are used as required (such as the calculations after 8, 24, or 32 bits) depending on the data width. This design is not optimal in FPGA devices. The XOR cancellations that can be performed in CRC designs mean that the function does not require all the intermediate calculations to determine the final result. Therefore, forcing the use of intermediate calculations increases the area required to implement the function, as well as increasing the logic depth because of the cascading. It is typically better to create full separate CRC blocks for each data width that you require in the design, then multiplex them together to choose the appropriate mode at a given time.

Use Separate CRC Blocks Instead of Allowing Blocks to Merge

Synthesis tools often attempt to optimize CRC designs by sharing resources and extracting duplicates in two different CRC blocks because of the factoring options in the XOR logic. As addressed previously, the CRC logic allows significant reductions but this works best when each CRC function is optimized separately. Check for duplicate extraction behavior if you have different CRC functions that are driven by common data signals or that feed the same destination signals.

If you are having problems with the quality of results and you see that two CRC functions are sharing logic, ensure that the blocks are synthesized independently using one of the following methods:

- Define each CRC block as a separate design partition in an incremental compilation design flow
 - For details, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.
- Synthesize each CRC block as a separate project in your third-party synthesis tool and then write a separate Verilog Quartus Mapping (.vqm) or EDIF netlist file for each

Take Advantage of Latency if Available

If your design can use more than one cycle to implement the CRC functionality, adding registers and retiming the design can help reduce area, improve performance, and reduce power utilization. If your synthesis tool offers a retiming feature (such as the Quartus II software **Perform gate-level register retiming** option), you can insert an extra bank of registers at the input and allow the retiming feature to move the registers for better results. You can also build the CRC unit half as wide and alternate between halves of the data in each clock cycle.

Save Power by Disabling CRC Blocks When Not in Use

CRC designs are heavy consumers of dynamic power because the logic toggles whenever there is a change in the design. To save power, use clock enables to disable the CRC function for every clock cycle that the logic is not required. Some designs don't check the CRC results for a few clock cycles while other logic is performed. It is valuable to disable the CRC function even for this short amount of time.

Use the Device Synchronous Load (sload) Signal to Initialize

The data in many CRC designs must be initialized to 1's before operation. If your target device supports the use of the `sload` signal, you should use it to set all the registers in your design to 1's before operation. To enable use of the `sload` signal, follow the coding guidelines presented in *"Secondary Register Control Signals Such as Clear and Clock Enable"* on page 6-38. You can check the register equations in the Timing Closure Floorplan or the Chip Planner to ensure that the signal was used as expected.

- If you must force a register implementation using an `sload` signal, you can use low-level device primitives as described in the *Designing with Low-Level Primitives User Guide*.

Comparators

Synthesis software, including Quartus II integrated synthesis, uses device and context-specific implementation rules for comparators (<, >, or ==) and selects the best one for your design. This section provides some information about the different types of implementations available and provides suggestions on how you can code your design to encourage a specific implementation.

The == comparator is implemented in general logic cells. The < comparison can be implemented using the carry chain or general logic cells. In devices with 6-input ALUTs, the carry chain is capable of comparing up to three bits per cell. In devices with 4-input LUTs, the capacity is one bit of comparison per cell, similar to an add/subtract chain. The carry chain implementation tends to be faster than the general logic on standalone benchmark test cases, but can result in lower performance when it is part of a larger design due to the increased restriction on the Fitter. The area requirement is similar for most input patterns. The synthesis software selects an appropriate implementation based on the input pattern.

If you are using Quartus II integrated synthesis, you can guide the synthesis by using specific coding styles. To select a carry chain implementation explicitly, rephrase your comparison in terms of addition. As a simple example, the following coding style allows the synthesis tool to select the implementation, which is most likely using general logic cells in modern device families:

```
wire [6:0] a,b;
wire alb = a<b;
```

In the following coding style, the synthesis tool uses a carry chain (except for a few cases, such as when the chain is very short or the signals *a* and *b* minimize to the same signal):

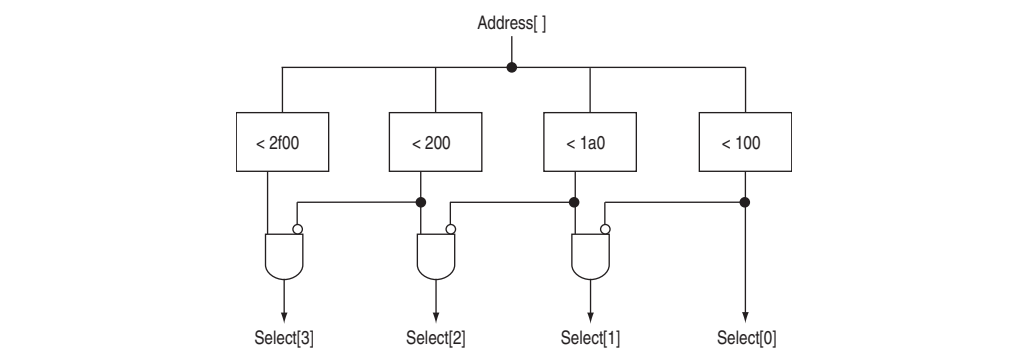
```
wire [6:0] a,b;
wire [7:0] tmp = a - b;
wire alb = tmp[7]
```

This second coding style uses the top bit of the *tmp* signal, which is 1 in twos complement logic if *a* is less than *b*, because the subtraction $a - b$ results in a negative number.

If you have any information about the range of the input, you have “don’t care” values that you can use to optimize the design. Because this information is not available to the synthesis tool, you can often reduce the device area required to implement the comparator with specific hand implementation of the logic.

You can also check whether a bus value is within a constant range with a small amount of logic area by using the logic structure shown in [Figure 6-9](#). This type of logic occurs frequently in address decoders.

Figure 6-9. Example Logic Structure for Using Comparators to Check a Bus Value Range



Counters

Implementing counters in HDL code is easy; they are implemented with an adder followed by registers. Remember that the register control signals, such as enable (*ena*), synchronous clear (*sclr*), and synchronous load (*sload*), are available. For the best area utilization, ensure that the up/down control or controls are expressed in terms of one addition instead of two separate addition operators.

If you use the following coding style, your synthesis tool may implement two separate carry chains for addition (if it doesn't detect the issue and optimize the logic):

```
out <= count_up ? out + 1 : out - 1;
```

The following coding style requires only one adder along with some other logic:

```
out <= out + (count_up ? 1 : -1);
```

In this case, the coding style better matches the device hardware because there is only one carry chain adder, and the -1 constant logic is implemented in the LUT in front of the adder without adding extra area utilization.

Designing with Low-Level Primitives

Low-level HDL design is the practice of using low-level primitives and assignments to dictate a particular hardware implementation for a piece of logic. Low-level primitives are small architectural building blocks that assist you in creating your design. With the Quartus II software, you can use low-level HDL design techniques to force a specific hardware implementation that can help you achieve better resource utilization or faster timing results.



Using low-level primitives is an advanced technique to help with specific design challenges, and is optional in the Altera design flow. For many designs, synthesizing generic HDL source code and Altera megafunctions gives you the best results.

Low-level primitives allow you to use the following types of coding techniques:

- Instantiate the logic cell or `LCELL` primitive to prevent Quartus II integrated synthesis from performing optimizations across a logic cell
- Create carry and cascade chains using `CARRY`, `CARRY_SUM`, and `CASCADE` primitives
- Instantiate registers with specific control signals using `DFF` primitives
- Specify the creation of LUT functions by identifying the LUT boundaries
- Use I/O buffers to specify I/O standards, current strengths, and other I/O assignments
- Use I/O buffers to specify differential pin names in your HDL code, instead of using the automatically-generated negative pin name for each pair



For details about and examples of using these types of assignments, refer to the *Designing with Low-Level Primitives User Guide*.

Conclusion

Because coding style and megafunction implementation can have such a large effect on your design performance, it is important to match the coding style to the device architecture from the very beginning of the design process. To improve design performance and area utilization, take advantage of advanced device features, such as memory and DSP blocks, as well as the logic architecture of the targeted Altera device by following the coding recommendations presented in this chapter.



For additional optimization recommendations, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Referenced Documents

This chapter references the following documents:


- *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*
- *Advanced Synthesis Cookbook: A Design Guide for Stratix II, Stratix III, and Stratix IV Devices*
- *ALTSHIFT_TAPS Megafunction User Guide*
- *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*
- *Designing with Low-Level Primitives User Guide*
- *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*
- *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*
- *Synthesis* section in volume 1 of the *Quartus II Handbook*

Document Revision History

Table 6-2 shows the revision history for this chapter.

Table 6-2. Document Revision History

Date and Document Version	Changes Made	Summary of Changes
March 2009 v9.0.0	<ul style="list-style-type: none"> ■ Corrected and updated several examples ■ Added support for Arria II GX devices ■ Other minor changes to chapter 	Updated for the Quartus II 9.0 software release.
November 2008 v8.1.0	Changed to 8-1/2 x 11 page size. No change to content.	Updated for the Quartus II 8.1 software release.
May 2008 v8.0.0	Updates for the Quartus II software version 8.0 release, including: <ul style="list-style-type: none"> ■ Added information to “RAM Functions—Inferring ALTSYNCRAM and ALTDPRAM Megafunctions from HDL Code” on page 6-13 ■ Added information to “Avoid Unsupported Reset and Control Conditions” on page 6-14 ■ Added information to “Check Read-During-Write Behavior” on page 6-16 ■ Added two new examples to “ROM Functions—Inferring ALTSYNCRAM and LPM_ROM Megafunctions from HDL Code” on page 6-28: Example 6-24 and Example 6-25 ■ Added new section: “Clock Multiplexing” on page 6-46 ■ Added hyperlinks to references within the chapter ■ Minor editorial updates 	Updates and enhancements to subject coverage for the Quartus II software version 8.0 release.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

