

Propagating Constants Past Software to Hardware Peripherals in Fixed-Application Embedded Systems

Frank Vahid

Dept. of Computer Science and Engineering
University of California, Riverside
Also with the Center for Embedded Computer
Systems, UC Irvine

Rilesh Patel

Aristo Technology
San Jose, California
rilesh@aristotech.com

Greg Stitt

Dept. of Computer Science and
Engineering
University of California, Riverside
gstitt@cs.ucr.edu

vahid@cs.ucr.edu, <http://www.cs.ucr.edu/~vahid>

ABSTRACT

Many embedded systems include a microprocessor that executes a single program for the lifetime of the system. These programs often contain constants used to initialize control registers in peripheral hardware components. Now that peripherals are often purchased in intellectual property (core) form and synthesized along with the microprocessor onto a single chip, new optimization opportunities exist. We introduce one such optimization, which involves propagating the initialization constants past the microprocessor to the peripheral, such that synthesis can further propagate the constants inside the peripheral core. While constant propagation in synthesis tools is commonly done, this work illustrates the benefits of recognizing initialization constants from the software as really being constants for hardware. We describe results that demonstrate 23 times reductions in peripheral size, and 10-30% savings in power, on several common peripheral examples.

Keywords

Cores, system-on-a-chip, embedded systems, synthesis, low power, constant propagation, platforms, tuning, intellectual property.

1. INTRODUCTION

Embedded system designers are increasingly composing their designs from pre-designed intellectual-property cores, integrating those cores into a single chip model as shown in Figure 1, and then fabricating a chip [3]. A core is a description of a system-level component, like a microprocessor, memory, or peripheral component like a direct-memory access (DMA) controller or universal asynchronous receiver/transmitters (UART). Cores come in three forms. A soft core is a synthesizable hardware description language (HDL) model. A firm core is a structural HDL model. A hard core is a technology-specific layout. Many commercial core libraries now exist, e.g., [4], and core standards are evolving rapidly [9].

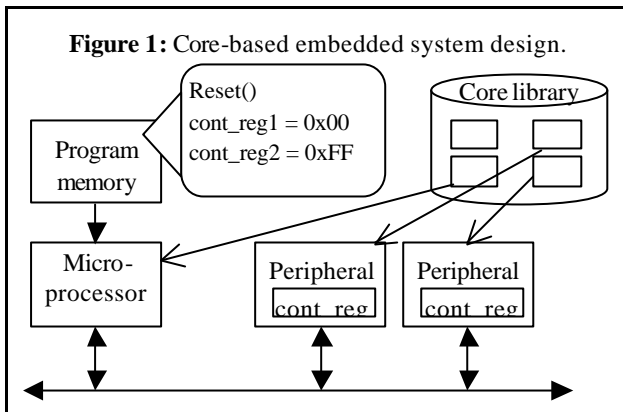
A designer gains many advantages from building a system from standard cores, such as a standard DMA controller or UART. Most importantly, the designer gains improved time-to-market due to familiarity with the standard core and

compatibility with development tools. Such standard cores typically come with parameters [8]. Some are pre-fabrication parameters, which are set by a designer before synthesis, thus influencing the synthesis results. Such parameters are typically achieved using generics or constants in a hardware description language (HDL), but can also be achieved using module generators, which generate unique HDL models depending on the parameter selection. For example, a JPEG decompression core might be synthesizable to have either 12 or 16-bit resolution. Synthesizing for 12-bit resolution would yield a smaller core.

Other parameters, in contrast, are post-fabrication parameters, set only after the core has been synthesized. Such parameters' settings are typically stored in registers or non-volatile memory inside the core. They are more commonly referred to as *software configurable* parameters. For example, a DMA controller will have a base register to indicate the starting address in memory from which the controller should move data, and a block size register to indicate the number of words that should be moved. An arbiter core might have a register whose setting determines whether arbitration uses a fixed or rotating priority scheme.

We make the observation that an embedded system typically runs a single program that never changes – the application is *fixed*. In fact, in many cases that program *cannot* be changed, because it may be burned into ROM (using mask-programmed ROM) that appears with the microprocessor and peripherals on a single chip to reduce chip cost, size and power (at the expense of less flexibility).

A typical embedded system will execute a boot program upon system reset, and this program will, among other things, set these software configurable parameters in the system's peripherals, as shown in Figure 1. However, if the embedded system's program never changes, then those register values never change during the execution of the embedded system. For example, a particular embedded system may use a DMA controller to repeatedly send data directly from an array, of size 48 and starting from memory location 100, to a display device. The system's boot program may set the DMA controller base register to 100,



and the block size register to 48. These values will never change for the life of the embedded system.

Previously, when systems were built using discrete off-the-shelf integrated circuits, such software configuration was necessary. However, since today's systems are being built with cores, we now have an optimization opportunity that did not previously exist. Specifically, for an embedded system whose program does not change, the values to which the software configurable peripheral parameters are being set *are really constants*. As compiler writers are well aware, constants provide excellent optimization capability, through the well-known compiler optimization known as constant propagation [1][10]. Such propagation consists of replacing a variable holding a constant by the constant itself. This replacement can result, for example, in branch conditions that always evaluate to false, resulting in turn in dead code that can then be eliminated. It can also enable compile-time evaluation of expressions.

Such dead code resulting from constant propagation is especially common when propagating constants into subroutines through the subroutine's parameters. While the subroutine may have been designed to handle a variety of sets of parameters, a particular program may only call the subroutine with certain constant values for those parameters, resulting in much dead code in the subroutine.

We can think of a peripheral core as similar to a subroutine, in fact, as a subroutine that has been implemented using additional hardware. The core may have been designed to handle a variety of sets of software configurable parameters. However, a particular program may only use the core with certain constant values for those parameters, resulting in much "dead code" in the core. We therefore propose a deeper propagation of constants than performed by compilers. In particular, we propose to propagate those constants beyond the microprocessor's program, to the microprocessor's peripheral cores – essentially propagating those constants all the way to peripheral hardware. Those constants would then be fed into the synthesis tool being

used to synthesize the cores. The synthesis tool could then perform constant propagations and dead code elimination during synthesis, where the code here refers to the core's HDL description. Most commercial synthesis tools already include such compiler optimizations, but those optimizations are only applied to the pre-fabrication parameter constants. We will show that much benefit would come from enabling the synthesis tool to recognize the post-fabrication parameter values as constants also.

The end result of such propagation is that the synthesized core will be optimized for the particular program that is using the core, something we refer to as *architecture tuning* [8]. By optimized, we mean that the core will have fewer gates, and consume less power, than a standard version of the same core. Reducing size is important since such reduction can increase chip yield and reduce chip cost, and many embedded systems are extremely cost sensitive, especially those being manufactured in high volumes. Reducing power is important since many embedded systems operate on batteries or draw power from very limited sources, and so power reduction is an important design criterion.

In this paper, we introduce the concept of propagating constants past software to hardware peripheral cores. After an introductory example, we'll describe common core parameters that are candidates for constant propagation, discuss methods for achieving such propagation, and highlight experiments showing the size and power reductions possible. The results motivate future work on developing tools that introduce some cooperation between the compilers and the synthesis tools being used in developing a system-on-a-chip from cores.

2. EXAMPLE

As a simple illustration of propagating constants to hardware, let us consider a trivially simple peripheral core that has two parallel ports. Each port can be configured to be an input port or an output port. A VHDL description of part of the core is shown in Figure 2(a). The core description declares a control register *cont_reg* with two bits. The first bit makes port A an output port when set to 0, and an input port when set to 1. Likewise, the second bit makes port B an input or output port. The VHDL description begins with initialization of the control register during a reset. Next, it would describe the synchronous monitoring of the bus for an address corresponding to the control register, and the writing of the control register in this case – this code is omitted from the figure. Next, the VHDL description describes the control logic for the tri-state buffers that implement the port direction functionality. Finally, other behavior of the core would be described.

Synthesis converts this soft core to hardware structure, shown in Figure 2(b). Note that logic is generated to handle

Figure 2: A simple example of propagating constants to hardware: (a) soft core, (b) synthesized core structure, (c) synthesized core structure after propagating constants $cont_reg(0)=0$ and $cont_reg(1)=1$.

```

signal cont_reg: UNSIGNED(1 downto 0);
-- declarations for A, B and buffers omitted.
process(clk, reset)
begin
if (reset) then
  cont_reg = "00";
  A_out <= '1'; A_in <= '0';
  B_out <= '1'; B_in <= '0';
end if;
if rising_edge(clk) then
  -- Code to detect write request
  -- from bus to cont_reg, and

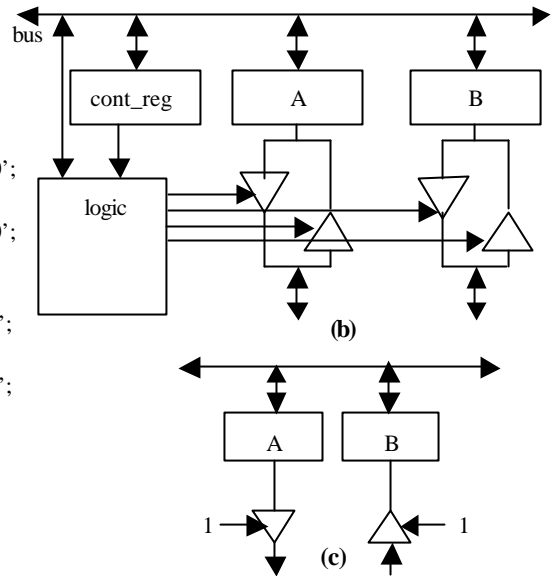
```

(a)

```

if (cont_reg(0) = '0') then
  A_out <= '1'; A_in <= '0';
else
  A_in <= '1'; A_out <= '0';
end if;
if (cont_reg(1) = '0') then
  B_out <= '1'; B_in <= '0';
else
  B_in <= '1'; B_out <= '0';
end if;
end if;
-- Other behavior omitted

```



the bus monitoring and the control of the four required buffers.

Now, consider the situation where this core is used in an embedded system and controlled by a microprocessor executing a fixed C program. We might see the following assembly code embedded in the reset routine of the C program:

```
OUT cont_reg, #"00000010"
```

Assuming $cont_reg$ is the address of the control register in the microprocessor's I/O address space, then this code would write the constant "00000010" onto the peripheral bus, resulting in a 0 being written into $cont_reg(0)$ and a 1 into $cont_reg(1)$. The peripheral core would thus be configured with port A as an output port, and port B as an input port. The rest of the C program would then access these ports appropriately.

Now, suppose we could somehow propagate the constant "10" into the VHDL description of the core, before the core were synthesized, letting the synthesis tool know that $cont_reg$ would be written by that constant and only that constant. If we did this in a way that our synthesis tool could make use of that information, then the synthesis tool would find much "dead code" in the VHDL description. First, the control register would not be needed, since a constant can be derived directly from power and ground in hardware. Second, the logic to monitor the bus for the control register address and then write the register would not be needed. Third, each buffer control signal if statement would have one branch that was always true and the other always false. Finally, the reset code of the core would not

be needed. After all of this dead code is eliminated, the synthesis tool would output the structure shown in Figure 2(c). The resulting structure in this case requires less hardware, and would also consume less power due in part to elimination of the bus monitoring.

3. PARAMETERS IN CORES

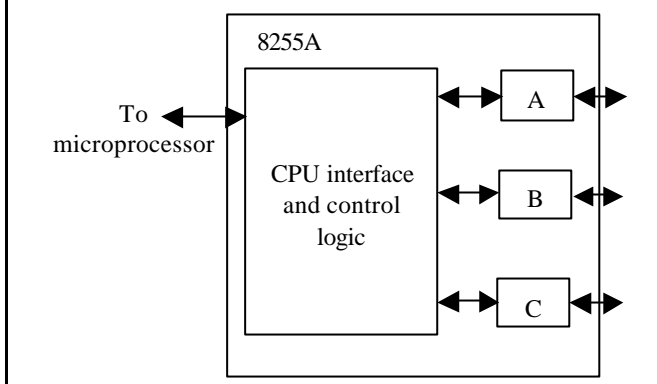
We examined a number of common peripheral cores, and found many software configurable parameters that could be candidates for constant propagation. Some common peripheral cores include the Intel 8255A (programmable peripheral interface), the 8237A (DMA controller), and the M16550A (UART – Universal Asynchronous Receiver-Transmitter).

Figure 3 is the block diagram of the 8255A. The 8255A interfaces with a microprocessor on one side, and provides three configurable ports on the other side. Its software configurable parameters include mode of operation, number of ports in use, and direction of each port (input or output). These parameters are set by a microprocessor by writing an 8-bit control word into a control register in the 8255A.

The 8237A includes even more software configurable parameters, including the number of channels, the type of priority scheme (fixed or rotating) being used to arbitrate between channels, whether each channel operates in single transfer mode or block transfer mode, the starting address and block size for each channel, etc. There are thus several control registers in the 8237A.

Likewise, the UART's parameters include the baud rate, parity type, mode of communication, etc.

Figure 3: The Intel 8255A parallel peripheral interface.



In general, peripheral cores tend to have several types of configurable parameters, related to features such as:

- size of internal data or external data bus
- number of channels or ports
- modes of operation
- I/O direction
- rate of data transfer
- resolution

Supporting numerous parameters is necessary in order for a peripheral to be applicable in a variety of systems and thus to sell in large quantities. While some parameters appear in a core as user-settable constants or generics, others appear as software configurable control registers. Such software configurability is used in peripherals for several reasons. One reason is that, before the advent of cores, software configuration was the only way to configure a peripheral integrated circuit (IC). A core may thus be modeling a widely-used standard peripheral that was defined in the time of ICs, such as UART and DMA controller cores. A second reason is that, even for cores representing new peripherals, the core designer does not know if the peripheral will be controlled by a microprocessor whose application will not change. If the core were used in a system whose application did change, then constant or generic-based parameters would not be appropriate. Thus, support of software configurable parameters is very common, but results in extra hardware size as well as power consumption.

4. PROPAGATING CONSTANTS FROM SOFTWARE TO HARDWARE

We now describe a method for manually propagating constants across the software/hardware boundary in a core-based synthesis methodology, and discuss potential approaches for automating this method. The method is summarized in Figure 4. For each core, the first step is to determine all of the registers in the core that serve as control registers for the various parameters listed in the

Figure 4: Method for propagating constants to peripheral cores.

```
for each peripheral core P
  for each control register C in P
    find all writes, W, to C in processor's program
    if W consists of a single write, of a constant X,
      in a reset or boot routine, then
        replace C in P by a constant declaration set to X
        delete behavior related to writing C in P
      end if
    end for
  end for
end for
run synthesis as usual
```

previous section. Next, for each such control register, we must look for all references to that register in the driving microprocessor program. If the only access to that register is a write with a constant, and this write occurs during the reset or boot routines, as is often the case in embedded systems, then we have a candidate for constant propagation to peripheral cores. We replace the register's declaration in the core by a constant declaration. We delete any behavior that involves detecting and carrying out a write to that register from the peripheral bus. We can then run the synthesis tool on this modified core. A synthesis tool will then detect and eliminate the dead code created by the constants we introduced in the model, and thus result in a simpler synthesized structure. Most modern synthesis tools already carry out standard compiler optimizations like constant propagation, constant folding, and dead code elimination.

We can also eliminate the behavior in the microprocessor's program relating to writing the control register, but this is not always necessary. If we do choose to leave it, then we must ensure that the lack of a response from the core is acceptable. If a response is needed, like an acknowledgement, then we leave such behavior in the core.

The above method has the advantage of being immediately applicable in any existing core-based design process, without any modification to existing tools. Of course, the constant propagation across the software/hardware boundary must be performed manually in the above case. Thus, we describe a potential approach to automating the method. A big help to such automation is if a core design framework is being used. Such frameworks, many of which are commercially available, manage the retrieval and instantiation of cores (e.g., [2][5]). They typically already have support for instantiating cores with specific values for

constants or generics (a generic is essentially a parameter whose value must be chosen before instantiation) and for keeping track of all register address assignments in a system of cores. Thus, modifying such frameworks to handle software configurable parameters can be seen as an extension of an existing method.

One approach to automation would be to extend the software compiler to output a list of external I/O addresses that are assigned a single constant by the program in a reset or boot routine, along with each address' associated constant. This requires that the compiler be aware of the location of those reset or boot routines. Next, each core must have its control registers known to the core framework – this can be done by the framework developer, or the framework user, without too much effort. Furthermore, the framework must know where in the core to find the code that writes the register. Given this setup, the framework can read the contents of the file output by the software compiler, and for each address the framework can then replace the corresponding register declaration by a constant declaration, and delete write behavior from the core, before instantiating the core into the design. Then, synthesis can be run on the instantiated core, and the constants will result in dead code that can be eliminated.

A second approach is possible, and in fact even simpler than the above. In particular, we observe that modern core-based frameworks actually generate the reset or boot code themselves, including the code for initializing peripherals [5]. In other words, suppose a user wishes to instantiate a DMA controller into a system already having a microprocessor and memory. The framework will query the user to ask for the values of software configurable parameters, like transfer mode, base address and block size. The framework then generates the necessary driver software on the microprocessor. The second approach extends the above by having the framework also ask if the software configurable parameter values will ever change, or if instead they are in fact constants. If constants, then the framework can withhold generation of the related driver software, and instead directly proceed to instantiate the core with the corresponding register declaration replaced by a constant, and with the register-write behavior deleted.

5. EXPERIMENTS

We performed several experiments to evaluate the size and power savings possible by using our method of propagating constants to peripheral cores. We modeled three popular peripherals as register-transfer level VHDL soft cores: the 8255A programmable peripheral interface, the 8237A DMA controller, and the 16550A UART. Each core model is nearly a fully-functional model. The three soft-core models required 1045, 920 and 1063 lines of VHDL code, respectively. We also obtained a discrete cosine transform

(DCT) core (Free-DCT-L) from <http://www.opencores.org>, which consisted of 910 lines of code. We manually modified these models to eliminate dead code that would have resulted from constant propagation of the software-configurable parameters described below. We synthesized the cores twice, once before and once after dead code elimination, using the Synopsys Design Compiler. Area and power were measured using Synopsys analysis tools, with power measured while running a suite of test vectors for each core. Because we wanted to see first-hand the impact of the constant propagation on the size of the VHDL code, we performed the propagation of the constants and the dead code elimination manually, so we could measure the resulting lines of code.

The 8255A had only one configuration register used for selecting the modes of various ports. We examined the impact of propagating constants for three different configurations of this register. Mode0 corresponded to a configuration where port A of the device was used as an output port. Mode1 corresponded to port A being used as an output port with handshaking I/O. Mode2 corresponded to port A being used as a bi-directional port with interrupt I/O. Each situation resulted in a reduction of the number of lines in the model from 1045 to an average of only 415 lines.

The 8237A had several configuration registers, including those that select the arbitration mode, the number of active channels, and the transfer mode, base address, and block size of each channel. We examined the situation of using only a single channel, in single transfer mode. This reduced the model from 920 to 435 lines.

The PC16550 also had several configuration registers, including those that enable transmit and receive, select the interrupt mode, and select the baud rate. We examined two situations, one where the device was configured for transmit only at a specific baud rate, and the other where it was configured for receive only at a specific baud rate. Each reduced the model's lines of code from 1063 to roughly 625.

The DCT core had configuration registers for selecting between forward and inverse DCT, and for selecting among 8, 9, 10, or 12-bit resolution. We tested the configuration of forward DCT with 8-bit resolution. This configuration reduced the size of the code from 910 lines to 867 lines.

Note that these parameters were not represented by constants or generics in the VHDL source. Rather, the cores were designed to be synthesized to support software configuration of these parameters, as is common.

The size and power data is summarized in Table 1. We see that size after synthesis was reduced by an average of 58%, and power by an average of 22%. The reason that power is not reduced as much as size is because many of the gates eliminated through constant propagation were not used during a core's execution even when present, so didn't

Table 1: Comparison of cores before and after constant propagation.

Cores	Gates, original	Gates, with constant propagation	% size savings	Power, original (micro-watts)	Power, with constant propagation (micro-watts)	% power savings
8255A mode-0	3069	834	73%	2772	1902	31%
8255A mode-1	3069	918	70%	2915	2098	28%
8255A mode-2	3069	953	69%	2952	2124	28%
8237A single transfer	7276	2344	68%	2453	2097	15%
PC16550 Tx	2503	1169	53%	1461	1249	15%
PC16550 Rx	2503	1188	53%	1449	1307	10%
DCT Forward 8-bit	2295	1861	19%	1391	962	31%
Average savings			58%			22%

consume much power. The power reductions that do occur result we believe from less switching activity occurring due to simpler control and datapath switching logic.

These reductions come of course at the cost of not being able to reprogram the configurable parameters of the core once the system has been implemented. Thus, if modifying the microprocessor's program is a possibility, then propagating constants across the software/hardware boundary should either not be done, or should be done only to the extent that the designer is certain that particular constants won't change. However, as mentioned earlier, many embedded systems have their programs fixed in mask-programmed ROM, and thus the configurable parameters could never have been modified anyways, meaning our approach would have no impact on flexibility in those cases.

6. CONCLUSIONS

As core-based design methodologies grow in popularity, cores will be heavily parameterized to increase applicability and hence sales. Pre-fabrication parameters, specified using HDL generics or constants, can result in optimized hardware. However, post-fabrication parameters, known as software configurable parameters, until now had not been exploited similarly. We introduced the idea of propagating constants beyond the microprocessor software, to the peripheral hardware. We showed that such propagation yielded reductions in size by 23 times, and good power reductions of between 10-30%, using several standard peripheral examples. This work is part of the UCR Dalton project, which seeks to develop techniques for parameterized core-based system-on-a-chip design [7]. This work motivates the need for future work on system-on-a-chip frameworks whose compilers are able to detect "constants" in the sense of software configurable register values, and are able to coordinate between compilers and synthesis tools to propagate those constants to hardware.

7. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grant number CCR-9876006.

8. REFERENCES

- [1] Aho, A.V., R. Sethi, J.D. Ullman. "Compilers: Principles Techniques, and Tools," Reading, Addison-Wesley Publishing Company, March 1998.
- [2] Escalade Corporation, <http://www.escalade.com/>.
- [3] Gupta, R., and Y. Zorian. Introducing Core-Based System Design. IEEE Design & Test, Vol. 14, No. 4, Oct-Dec 1997, pp. 15-25.
- [4] Inventra core library, Mentor Graphics, <http://www.mentor.com/inventra/>.
- [5] Platform Express. Mentor Graphics, http://www.mentor.com/soc/platform_ex/.
- [6] Stitt, G., F. Vahid, T. Givargis, and R. Lysecky. A First-step Towards an Architecture Tuning Methodology for Low Power. Compilers, Architectures, and Synthesis for Embedded Systems (CASES'00), November 2000, pp. 187-192.
- [7] The UCR Dalton project: <http://www.cs.ucr.edu/~dalton>.
- [8] Vahid, F., and T. Givargis. Platform Tuning for Embedded Systems Design. IEEE Computer, Vol. 34, No. 3, March 2001, pp. 112-114.
- [9] Virtual Socket Interface Association, Architecture Document, <http://www.vsi.org>, 1997.
- [10] Wegman, M., and F.K. Zadeck. Constant Propagation with Conditional Branches. ACM Transactions on Programming Languages and Systems, Vol 18, No 2, April 1991, pp. 181-210.