

# FPGA Code Portability for Signal Processing in Software Defined Radios

Jim Hwang  
Xilinx, Inc.  
2100 Logic Drive  
San Jose, CA  
(408) 879-4883

jim.hwang@xilinx.com

## ABSTRACT

*There has been considerable recent interest in defining a hardware abstraction layer to facilitate code reuse in the signal processing subsystems of software-defined radios. Code portability for FPGA-based signal processing is a significant aspect of such HAL efforts. In this paper, we show how a platform-based approach to FPGA design that provides a high level of design abstraction can also provide the ability to target multiple FPGA families from a single source model. The approach combines direct mapping of a Simulink model with code generation of register-transfer level HDL. By exploiting retiming and other optimizations available through logic synthesis, it is possible to obtain very efficient realizations of signal processing functions. This work complements HAL recommendations by focusing on mechanisms, guidelines, and methodologies for constructing signal processing functions in FPGAs. It helps to address requirements for executable specifications, as well as providing source that can be compiled through automatic code generation.*

## 1. INTRODUCTION

Field-programmable gate arrays (FPGAs) are widely used to implement physical layer signal processing functions for software-defined radios (SDRs [1] [2]). FPGAs provide very high performance custom hardware solutions, and can be reconfigured in system, and when bringing up a new waveform in the modem. Despite their reprogrammability, they have historically been considered part of the “hardware” within a modem, rather than part of the “software”. Consequently, the SDR software control layer, or Software Communications Architecture (SCA [12]), has largely ignored issues related to the specification, configuration, signal transport, or inter-component interfaces that are important to the platform provider of an SDR that exploits FPGAs..

The U.S. government has been the primary end customer for SDRs, with significant investment in the technologies and products, e.g., as part of the Joint Tactical Radio System (JTRS) program run by the Joint Program Office (JPO) of the U.S. Department of Defense. Whereas such government programs have lifetimes on the order of ten or more years, FPGA vendors continue to provide new devices roughly every 12-18 months. The increased signal processing capability of new families of FPGAs has remained sufficiently compelling that most platforms require retooling to incorporate new devices as they become available. Recognizing that the current SCA standard does not sufficiently address the design and deployment of the FPGA portion of the modem, the JPO has recently embarked upon a concentrated effort in the SDR community to extend the SCA to

provide guidance and ideally, standardization, for the use of FPGA technologies within SDRs [13][12].

More or less concurrently with the JTRS JPO efforts to recommend SCA extensions, the Software Defined Radio Forum formed a working group devoted to providing recommendations for a hardware abstraction layer to assist in the development, maintenance, and cost management of SDRs [9]. Although there are many viewpoints being brought forth by design tool, component, and platform vendors, as well as system integrators and subcontractors for the JTRS program, there is general agreement that FPGA code portability is an important, but to date, largely neglected aspect of design methodologies.

In this paper, we describe a platform-based approach for obtaining portable FPGA source code, whilst simultaneously providing executable specifications, test harnesses, and “golden” test vectors (i.e., providing accurate input/output relations for establishing conformance to specification through simulation). Our approach treats a high-level system model specified in Simulink [10] as the source code for an FPGA implementation. A block in the model may map onto a set of intellectual property blocks provided by the vendor that exploit vendor-specific device resources to implement the block’s function efficiently in a number of FPGA families. Alternatively, a block may map onto a behavioral description in a hardware description language that is inherently portable. It is on the latter case that we focus in this paper. The approach extends widely used FPGA design techniques, using industry standard design tools. Although described in terms of proprietary (though commercially available) tools for Xilinx FPGAs, our approach is equally applicable to other devices.

In Section 2, we present several definitions of code portability, and comment on their feasibility with current device technologies and design tools. Section 3 provides a brief introduction to a platform-based design methodology for implementing DSP systems in FPGAs that underlies our approach to code portability. Section 4 describes a case study of the approach, building a fractionally spaced equalizer (FSE) for a QAM system that is relatively simple, but a representative example of a modem function well-suited to an FPGA.

## 2. CODE PORTABILITY

Major FPGA vendors have multiple device product lines, each of which may be further divided into different families, each of which is further divided into different part types that differ in available resources, speed grade, and package. For example, Xilinx has two primary FPGA product lines: Virtex, which targets

highest performance and gate density, and whose most recent families include Virtex-4 and Virtex-II Pro [16], and the Spartan family, which targets high volume and lower cost applications. Spartan-3 is the most recently introduced family member [14].

Because a new FPGA family is introduced roughly every 12-18 months, and the design cycle for a major SDR design can be a significant fraction of this period, the implications of code portability (or more accurately, non-portability) are clear. Often a system must be built to target a family in advance of broadly available silicon.

*Bitstream portability* means that a bitstream for FPGA family  $v(i)$  will run directly, possibly via an intermediate run-time software layer, on a  $v(i+1)$  part. In terms of cost reduction, it is also desirable that a bitstream for FPGA family  $v(i)$  run directly on a different family device  $s(j)$ . However desirable it may be, at the current time no FPGA vendor supports bitstream portability.

*Source-level portability* implies that source code written for device  $v(i)$  will run after recompilation (but otherwise without change) in device  $v(i+1)$ . In addition, it is desirable to have source level portability between families  $v(i)$  and  $s(i)$ . Many FPGA users adopt a set of internal guidelines to facilitate full or near source-level portability. In this paper, we describe one way in which source level portability can be achieved using existing devices and design tools.

## 2.1 Register Transfer Level HDL

The prevailing abstraction in hardware description languages for FPGA design is register transfer level (RTL), which can be synthesized into device-specific logic resources [4]. At this level of abstraction, a design is a network of combinational circuits separated by registers. Registers and other circuit elements are represented behaviorally through idioms inferable by commercial synthesis tools. This style of coding allows the user to specify for example an addition operation with the operator '+', with the synthesis tool mapping this appropriately to device specific architecture primitives.

Considerable progress has been made over recent years in commercial synthesis tools to efficiently target FPGAs. In addition to technology mapping, synthesis tools also apply optimization algorithms to a circuit that preserve behavior, while improving the circuit quality under well-defined criteria (typically logic area or performance). Of particular interest is retiming, which is the reallocation of unit delays (e.g. registers) throughout a circuit, in order to reduce the number of combinational logic levels [4]. There is a close correlation between the largest number of logic levels and the frequency with which bounding registers can be clocked without setup or hold time violations, so retiming is a particularly effective synthesis optimization.

## 3. A PLATFORM-BASED APPROACH TO FPGA DESIGN

Design methodologies for FPGAs traditionally lag those for application specific integrated circuits (ASICs) by roughly a decade, in large part because until recently, the design complexity lagged by roughly the same factor. Of late however, the gap in complexity has narrowed as device geometries have continued to shrink, and the relative complexity of FPGA designs has increased more rapidly than that of ASICs. Ideas relating to

platform-based design, originally motivated by systems-on-chip ASICs [6] have been increasingly adopted for FPGA design [5].

We interpret a platform to be an articulation point between abstract behavior and realizable function. Viewed from above, the platform is a restriction on the space of all realizable systems, but one that can be usefully employed to capture the behavior of end applications. Viewed from below, the platform is a restriction on the space of all possible applications, but one that can be readily realized.

More specifically, we consider the platform to be a set of arithmetic, logic, memory, and other functional abstractions that allow a user to specify an FPGA-based SPS in a natural way. At the same time the functions in the platform are defined so that they can each be implemented, possibly in a number of distinct ways according to additional constraints. As "platform provider", we implement a library of operators, functions, and objects that can be composed within a high-level framework to implement DSP systems. To the application programmer, the library can be used (indeed extended) to specify a rich set of DSP systems.

In this paper, we address one aspect of platform-based design, namely, how this approach can be used in a commercially available framework to obtain portable, yet highly efficient FPGA code.

### 3.1 System Generator for DSP

System Generator for DSP is a software framework for modeling and implementing systems in FPGAs using Simulink [5]. Simulink provides a powerful component-based computing model that is well suited for specifying the concurrency in a custom signal processing architecture. System Generator provides libraries of functions and hardware related abstractions that can be used to model a signal processing system suitable for FPGAs. Such models are bit and cycle accurate to FPGA hardware; System Generator ensures this by providing automatic code generation from Simulink to a combination of synthesizable HDL and intellectual property (IP) cores. In addition, System Generator extends Simulink to include event-driven HDL semantics, hardware co-simulation, and rich customization interfaces traditionally associated with modern programming languages [2] [15].

In this paper, we focus on an aspect of System Generator that is not widely appreciated: it has the ability to create generic RTL that is extremely efficient, and is portable. There are three ways to obtain RTL out of System Generator:

- Importing HDL modules as Black Boxes – although a trivial "mapping", this capability is powerful and should not be ignored;
- Targeting blocks that have RTL implementations, e.g., Expression block, Register, Delay, up/down samplers;
- Using the MCode block, which maps MATLAB .m code to synthesizable VHDL.

Because of its importance and utility, in this paper we concentrate on the MCode block and its application.

#### 3.1.1 M-Code Block

The System Generator MCode block provides a simple interface for interpreting a MATLAB .m function in the context of a Simulink simulation. The block is a convenient and flexible way



using Synplify Pro 7.6, with retiming and pipelining options enabled. The designs were run through mapping, placement, and routing using the Xilinx ISE 6.2.02i software, with highest placer and router effort levels. This process was run repeatedly with different clock frequency constraints in order to determine the highest frequency obtainable. The results, summarized in the following Table are somewhat surprising.

Part Type	MHz Cores	MHz Synth	LUTs Cores	LUTs Synth	DFFs Cores	DFF Synth
xc2vp20	88.9	104.1	1636	1595	1868	2195
xc3s1500	79.3	85.5	1636	1531	1868	2193

Although one might expect the generic RTL implementation to run at a lower clock frequency than the original that employed IP cores, in fact the reverse was true. For both Virtex-II Pro and Spartan-3 devices, the RTL version ran at an appreciably higher clock frequency. The RTL implementations used more DFF registers than the core-based implementations (recall the MCode did not specify SRL16 resources, although the synthesis tool was free to map onto them when it could infer them correctly).

Current synthesis tools treat IP cores as black boxes, and no optimizations are available that cross module boundaries. The RTL design in contrast allowed the synthesis tool to freely move registers and optimize logic across module boundaries. It should be noted however, that in all prior versions of the Synplify Pro (as well as all versions of the Xilinx XST tool), retiming did not provide significant speed-up for this design. One concludes that logic synthesis optimizations available (and necessary) for high-performance design are still in an early stage of development.

## 5. CONCLUSIONS

With the increased adoption of FPGAs as signal processors comes an increased expectation for design flows and methodologies that support similar programming models to general purpose and DSP processors. Code portability, at least at the source level (i.e., admitting recompilation) is of fundamental importance. Although FPGA source code is not as widely portable as code for general-purpose microprocessors, we have demonstrated how System Generator and similar design tools provide considerable progress towards this end. Using an adaptive FSE as an example, we have shown how a single System Generator model can be used to specify both behavior and implementation, producing a generic RTL implementation suitable for an FPGA. The design exploits retiming and logic synthesis optimizations in order to achieve high performance.

## 6. ACKNOWLEDGEMENTS

The author gratefully acknowledges the contributions of, and numerous fruitful discussions with colleagues Brent Milne, Haibing Ma, and Brad Taylor.

## 7. REFERENCES

- [1] Berkeley Design Technology, Inc. "FPGAs for DSP," Focus Report, July 2002, [http://www.bdti.com/products/reports\\_focus.html](http://www.bdti.com/products/reports_focus.html)
- [2] C. Dick and J. Hwang, "FPGAs: A Platform-Based Approach to Software Radios," in *Software Defined Radio: Baseband Technologies for 3G Handsets and Basestations*, (W.H.W. Tuttlebee, Editor), Wiley, 2004.
- [3] C. Dick, "Design and implementation of high-performance FPGA signal processing datapaths for software-defined radios", *VMEbus Systems*, August 2001.
- [4] G. DiMicheli, *Synthesis and Optimization of Digital Circuits*, McGraw-Hill, 1994.
- [5] J. Hwang, B. Milne, N. Shirazi, J. Stroemer, "System Level Tools for FPGAs," *Proceedings FPL 2001*. Springer-Verlag 2001.
- [6] K. Keutzer, A.R. Newton, J.M Rabaey, and A. Sangiovanni-Vincentelli, "System-Level Design: orthogonalization of concerns and platform-based design", *IEEE Trans. On Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, pp.1523-43, 2000.
- [7] I. Page and W. Luk, "Compiling occam into FPGAs", in *FPGAs*, W. Moore and W. Luk (editors), Abingdon EE&CS Books, 1991, pp. 271-283.
- [8] V. Singh, A. Root, E. Hemphill, N. Shirazi, J. Hwang, "Accelerating a Bit Error Rate Tester with a System Level Tool," *Field-Programmable Custom Computing Machines, FCCM 2003*, Proceedings, IEEE 2003.
- [9] Software Defined Radio Forum, Hardware Abstraction Working Group, [http://www.sdrforum.org/tech\\_comm/halwg.html](http://www.sdrforum.org/tech_comm/halwg.html).
- [10] The Mathworks, Inc., *Using Simulink*, 2002.
- [11] J.R. Treichler, I. Fijalkow, and C.R. Johnson, Jr., "Fractionally Spaced Equalizers", *IEEE Signal Processing Magazine*, May 1996, pp. 65-81
- [12] U.S. Department of Defense, Joint Tactical Radio System, Software Communication Architecture, Technical Overview. [http://jtrs.army.mil/sections/technicalinformation/fset\\_technical\\_sca.html](http://jtrs.army.mil/sections/technicalinformation/fset_technical_sca.html)
- [13] U.S. Department of Defense JTRS Joint Program Office, SCA Extensions Workshop, Arlington, Va., April 2004, [http://jtrs.army.mil/sections/programinfo/fset\\_programinfo.html?programinfo\\_industry](http://jtrs.army.mil/sections/programinfo/fset_programinfo.html?programinfo_industry).
- [14] Xilinx, Inc., *Spartan-3 Data Sheet*, <http://direct.xilinx.com/bvdocs/publications/ds099.pdf>.
- [15] Xilinx, Inc., *System Generator for DSP User Guide*, [http://www.support.xilinx.com/products/software/sysgen/app\\_docs/user\\_guide.htm](http://www.support.xilinx.com/products/software/sysgen/app_docs/user_guide.htm).
- [16] Xilinx, Inc., *Virtex-II Pro Platform FPGA User Guide* <http://direct.xilinx.com/bvdocs/userguides/ug012.pdf>.