Some Thoughts on Process Retargettable and Reusable IC Intellectual Property

Neil Weste

Electronics Dept.

Macquarie University

Sydney 2109 NSW Australia

1. Introduction

Recently there has been a dramatic rise in the interest in the ability to provide process portable IC designs (intellectual property - IP). This has occurred for two main reasons. Firstly, as systems get more complicated, they start to involve considerable engineering effort, effort that managers do not want to squander when the next project starts. Secondly, integrated "systems on a chip" combine the disparate technology normally found on a single printed wiring board. So rather than the integration happening on the board, it happens on the chip. As chips are usually designed by a single group or company, all of the functions formally available in chip form have to be available as "soft" IP.

At a first glance the above problem doesn't seem to be too hard in today's CAD tool environments. The process to incorporate a new module might be as follows:

a. Have breakfast

b. Check out and purchase behavioral designs from Internet using NetCredit

c. Combine and Synthesize

d. Place and Route design

e. Ship tp Manufacturing

f. Have lunch

Of course, this is currently a fantasy (But who knows for how long?). This talk will examine some of the issues in making the above scenario work as close as possible to ideal. This will be based on the experience of the author in designing and supplying process independent IP over the last 17 years.

The ability to mix and match IC IP assumes some standard interchange formats and standards. In addition, the market is already clearly dividing into IP sources and IP sinks. The former tend to be small companies while the latter tend to be large companies. The legal and business issues surrounding this emerging industry are therefore quite interesting.

2. Where to start?

An obvious place to start looking for solutions is as high up the design process as possible. That leads us to Behavioral HDLS. And, as VHDL and Verilog are the virtual standards they are the obvious targets. While ultimately, these specifications might well be the solution, at present they do not seem appropriate where high performance is required. Emerging tools such as combined datapath and control synthesis from algorithmic descriptions might well be a viable solution to a large class of IP. Current HDLs are poor at specifying and generating parameterizable designs. Designers currently use meta-languages to generate HDLs that are then used further in the design cycle. So these (C, C++, Lisp) specifications might well become the specification level for IP providers.

Specifying a design at the RTL level exerts testing of final chip more control over the speed of the logic but is still subject to the vagaries of place and route. Controlling the layout can put tight bounds on timing and area and is the lowest level where process

ASP-DAC '97
0-89791-851-7/$5.00

independence can be maintained. The idea of "virtual" floorplans seems to have some merit.

In summary, the design needs to be captured at the highest level that is consistent achieving the required performance (where performance might involve cycle time, power dissipation, die size, design time or a combination of all of these metrics). In reality, IP has and will continue to be captured at various levels and engineers will find ways of efficiently combining this IP.

3. What deliverables do you need?

What comprises a piece of "soft" IP? Industry in the US has already made a start on this by working on a "virtual socket" interface. This work is ongoing but we can make a guess at what might be required by considering how current packaged ICs are documented. The list of items would include:

a. documentation on theory of operation

b. an I0 pin specification

- names
- port descriptions

c. a behavioral model

- for higher level simulation including this module
- possibly for synthesis
- for comparison with generated version

d. an RTL description

- to synthesize, place and route

e. a functional verification suite

- to verify design and be used in testing of final chip

f. a timing shell (for a P&R'd design)

- input loadings
- delays
- setup and hold times
- output drive

g. manufacturing test vectors

h. power dissipation, DC ratings, etc.

And perhaps:

i. a virtual floorplan

- to provide a targeted placement

j. size, aspect ratio, port information (size,location,layer) for a physical instantiation

4. What tools do you need?

While we can all benefit from better and higher level synthesizers, current tools provide the basis for providing process portable IP. What is required is more automation of the design flows that are used with those tools. The following gives a flavor of what we have implemented in the Electronics Department at Macquarie University.

4.1 A High Performance HDL Design Flow

Macquarie has been working on a design flow that allows the fast design of high performance DSP chips for wireless LANs and other DSP applications. This work follows on tools developed by CADENCE Design Systems and circuit and design technology developed by the author at TLW, Symbolics and AT&T Bell Labs. The flow may be summarized as follows:

a. capture behavioral design in Verilog

- EMACs
- CADENCE Verilog-XL
- Veriwell

(PC/Mac Verilog simulation)

- UNIX make, etc.

b. verify functionality - provide test suite

- EMACs
- CADENCE Verilog-XL
- Veriwell

(PC/Mac Verilog simulation)

- UNIX make, etc.

c. convert behavioral Verilog to RTL Verilog

- EMACs
- UNIX make, etc.

d. reverify functionality

- EMACs

- CADENCE Verilog-XL
- Veriwell
(PC/Mac Verilog simulation)
- UNIX make, etc.

e. place and route design
- CADENCE CeII3 Place and Route
- CADENCE SmartPath datapath generator
- Macquarie Process Independent Standard Cell Engineering System (MacPISCES) - (built on top of CADENCE DFII)
- CADENCE ACC - Automatic Cell Characterization
- CADENCE SPECTRE circuit simulator
- CADENCE compactor
- CADENCE Skill Environment
- UNIX make, etc.

f. timing analyze design - CADENCE Pearl Timing Analyzer
- UNIX make, etc.

g. iterate until timing specification achieved
- CADENCE delay and capacitance extractors

h. generate timing shell
- CADENCE Pearl Timing Analyzer
- UNIX make, etc.

i. document design
- EMACs

j. ship design

The design flow can support the "synthesis" (really procedural generation) of high performance datapaths (as well as control) in a process independent manner. The two key pieces of technology used here are a datapath generation tool available from CADENCE called SmartPath (SmartGen) and a CMOS process independent library and characterization system developed at Macquarie (MacPISCES). With the ability to have datapaths synthesized directly from Verilog with the same placement as would have been achieved by hand, a large degree of control over timing is available. The process retargettable library allows the generation and characterization of a new library within a few hours in a completely "hands off" manner. MacPISCES completes the following steps:

a. generate virtual grid symbolic layout for cells (Macquarie)

b. compact cells, create geometry, create abstracts (CADENCE)

c. generate Verilog (Macquarie)

d. generate simulations scripts (Macquarie)

f. characterize cells with ACC (Automatic Cell Characterization -CADENCE)

g. annotate Verilog with timing (Macquarie)

h. generate timing analyzer (Pearl) views (Macquarie)

i. generate synthesis library (Macquarie scripts with CADENCE's Synergy)

The system is currently being used in the design of a high speed (25Mb/s) WLAN for the 5GHz NII/SuperNet band. The following modules are under development:

- 5OMHz 16 point complex FFT/IFFT
- 50 Mhz Viterbi coder/decoder
- SPARC V8 processor core for the MAC and other system functions
- 50 MHz 16 bit Cordic processor for rectangular->polar conversion and other functions
- Interleaver/Deinterleaver
- Differential Phase modulator/demodulator

Most of these modules have been implemented previously as FPGAs, ASICs or full custom chips. This gives us the opportunity to compare the new methodology with the older methodologies. We have already found that the new methodology provides rapid capture of functionality combined with sizes that are close to previously done custom designs.

In particular, the SPARC core, which we intend to use for a wide range of embedded applications, provides a

good testbed for datapath, control and memory structures.

The problem of process independent analog design is also an interesting problem which requires a great deal of research.

5. Legal and Business Issues

As with any emerging technology new legal and business issues arise. Issues include compensation, copy protection, nonperformance recourse and patent infringement. The main difference with current design outsourcing is that potentially the IP provided in a technology independent manner may be used for many different designs and design variations for a long time into the future. It is hard at the beginning of a development effort to judge how successful a particular piece of IP might be, and the designer certainly should share in the success (or failure) of the design. The observation was made that IP providers tend to be small while IP users tend to be large. This asymmetry in size is paralleled by the size of legal staffs. Thus a degree of trust has to be built up between source and sink to enable smooth running of the business. Pricing models also need to be worked out. Various schemes exist today - upfront payment, royalties or a mixture of both. Royalties are not favored by many manufacturers and so there is an opportunity for imaginative new compensation schemes.

Of particular concern is the scenario that might occur if a piece of fabricated IP fails to work or perform. This clearly requires a "service after sales" model, unless the IP user has enough engineering resources to deal with the problem (a bit like redesigning a chip you used on a PWB).

6. Summary

The need to combine disparate IP onto large systems chips is providing an opportunity for both design automation providers and users to hone their skills. Interesting things are going to happen in this area.

An undercurrent in this talk maybe needs to be stated explicitly. While traditional DA and CAD algorithmic problems are largely solved, a myriad of research problems in the DA area remain in the *use* of tools. To contribute to this dynamic area, DA researchers need to couple with active design projects and solve the ensuing design problems.

7. Acknowledgements

The author would like to acknowledge the support of the CSIRO Division of Telecommunications and Industrial Physics, CADENCE Design Systems and Macquarie University.