VLSI for Multimedia U-NII WLANs

N. Weste, D. Skellern and T. Percival*

Electronics Department, Macquarie University, Sydney, Australia
Tel. + 61 2 850 9148 Fax. + 61 2 850 9128
email: new@mpce.mq.edu.au
*Division of Telecommunications and Industrial Physics, CSIRO, Sydney, Australia

Abstract - This paper summarizes aspects of the VLSI development of a high-speed wireless local area network (WLAN). The implications for system-on-a-chip designs are summarized.

I. INTRODUCTION

This paper describes an ongoing research program that has prototyped a high speed (18Mbps) 5 GHz WLAN and is due to produce a highly integrated 30Mbps system in early 1998. The aim of the paper is to summarize the research and development of this system from initial concept to the current level of integration and relate this to system on a chip design issues.

The CSIRO and Macquarie University have been working on high-speed wireless communication for about seven years[1]. Initially, the work comprised of simulations aimed at selecting a modulation technique that was multipath resistant. This lead to a communications test-bench which was used to conduct on-air tests at 2 GHz and 40 GHz with a software based modulator and demodulator. After a few years of in-building measurements, it was decided that an orthogonal frequency division multiplex (OFDM) or discrete multi-tone (DMT) modulation technique was desirable. The mm-wave band was selected as it was thought at the time that no spectrum would be allocated in the microwave area. The recent allocation in the USA by the FCC of the 5 GHz U-NII[2] (Unlicensed National Information Infrastructure) bands has revised that opinion and we are now working exclusively on the problems at 5 GHz.

II. SYSTEM DESIGN

Figure 1 shows the basic block diagram of a broadband data radio. It comprises an RF section, and IF section, a baseband modem (DMT-50), a MAC interface and a bus interface (to a serial Utopia bus). As previously mentioned the RF section operates at 5 GHz. The IF operates at 600 MHz. The baseband modem is clocked at 50 MHz and achieves a throughput of around 30Mbps. An option is available to operate the modem at a speed of 100 MHz to achieve a throughput of 60Mbps. A conventional double conversion superheterodyne receiver has been used.

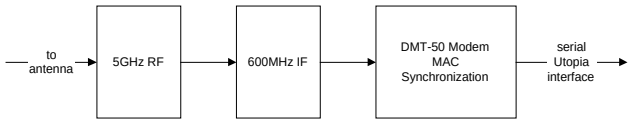

Figure 1. 5 GHz Data Radio

The architecture of the RF front-end is shown in Figure 2. The first local oscillator is at 2.4 GHz and an active doubler provides a 4.8 GHz LO to the bi-directional mixer. This radio is being manufactured by M/A-COM with a chip on board technology in a module that measures 45mm on a side. The cumulative noise figure is around 8dB with a four-way antenna diversity switch. The power output ranges from +10dBm to +25dBm.

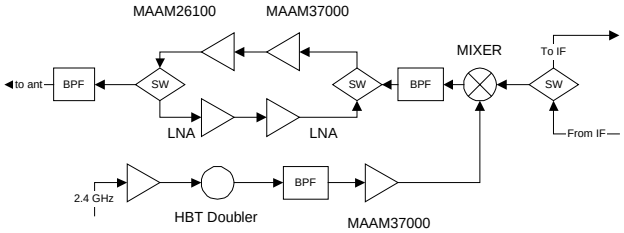

Figure 2. 5GHz front-end

The IF nominally operates at 600 MHz (see Figure 3). The receive section consists of a variable gain amplifier, followed by an IQ demodulator. This is passed through a 7th order antialiasing filter and then to dual 8-bit ADCs operating at 50 MHz. The transmit section comprises two 8-bit current mode DACs driving a reconstruction filter (the antialiasing filter reused) with the filtered output feeding an IQ quadrature modulator. The RX and TX paths are bandpass filtered at 600 MHz. At present we plan to use a SAW filter.

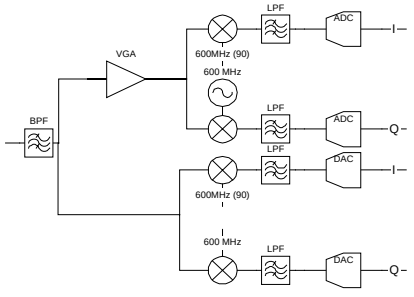

Figure 3. IF Architecture

The first implementation of the IF uses commercial chips for the amplifiers, ADCs, DACs and IQ up-down converters. We have designed a 5th order continuous time tunable filter in CMOS technology that will form the basis for the filter in the fully integrated version[3]. The IF is completely implementable in CMOS and it is our plan to complete this integration in 1998.

The baseband modem architecture is shown in Figure 4. This is the outline of a DMT modem that we have implemented as a single chip (the DMT-50). The main signal-processing block responsible for frequency to time conversion (and vice-versa) is a 16 point FFT operating at 50 MHz.

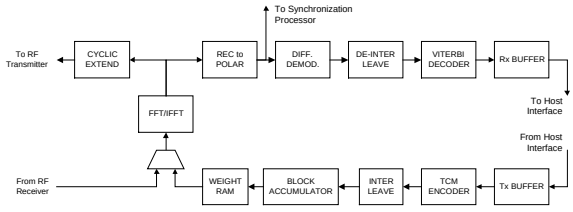

Figure 4. DMT-50 Modem

Initially, the FFT was implemented as a single chip. It works up to 126 MHz in a 0.6µm CMOS process[4]. At 50 MHz and 2.5 volts it dissipates 50 mW. The complete modem (237K transistors) is currently in fabrication and occupies 5.6mm*5.1mm in a 0.6µm TLM CMOS process and is packaged in a 132-pin PGA package. We estimate the power dissipation for the full modem to be around 200 mW at 2.5 volts.

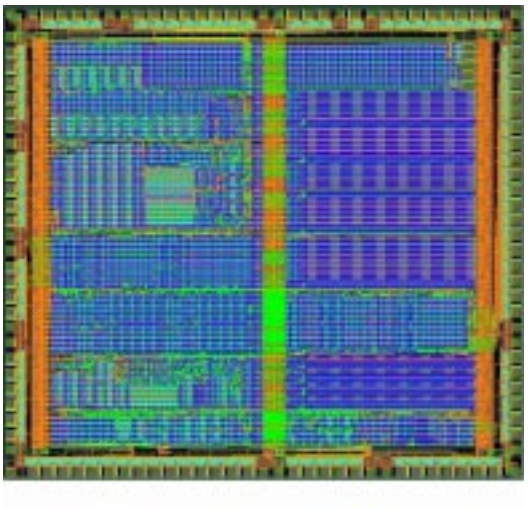

Figure 5. DMT-50 Modem Chip

The modem has to communicate at high speed with a bus or network. Currently, we are planning on using an FPGA to implement this interface. In addition, we will initially use an adjunct processor or host processor for the MAC and network interface. Eventually, we plan to co-locate a SPARC V8 processor with the modem on a single chip. We have a process-independent version of the SPARC V8 that works at 70 MHz in a 0.5µm process and occupies around 5mm^2.

III. DESIGN

We have implemented a range of hardware related to this project . This has ranged from software based modeling to FPGA prototypes to VLSI modules. In this section we summarize our approach to the VLSI design.

There are two key components of our VLSI design approach. Both are based on a mixture of CADENCE tools with custom written software[5].

MacFLOW is a Verilog driven design flow that produces high performance custom layouts with the minimum of effort. A standard automatic P&R/synthesis flow is augmented by a datapath system called SmartPath[6] (a CADENCE product). This is accompanied by MacPISCES, a fully automatic standard cell generation and characterization system. This system provides both regular logic, latches, registers and adders in addition to special cells used in the construction of RAM and ROM. The combination of all components allows a single scriptable design flow to be used for the entire chip.

Individual Verilog test benches are written for major modules. This might be for a memory, a multiplier, a Viterbi decoder, FFT or the complete modem. This task is usually performed by the module designer.

Also, we generate an almost exhaustive set of functional verification code that uses a C language golden model as the reference. This allows the setting of various control registers and application of large amounts of data to probe the boundary points in the design. In the design of the modem, the transmitter and receiver can be connected together to perform a complete loopback test. Internally in the chip we have a loop-back path prior to the FFT (it is shared so can not be used both for RX and TX at the same time).

IV. IMPLICATIONS FOR SYSTEM-ON-A-CHIP

This project has implemented a new piece of functionality (i.e. a high speed WLAN DMT modem) from essentially pure research. It is interesting to note what seems of relevance, in the project to the design of future systems on a chip.

The following points seem to be important:

- Use an HDL approach to design
- Automate the design flow as much as possible
- Reduce the "edit-compile-debug" loop
- Prototype the design in FPGAs
- Provide copious testing points for embedded modules
- Don't design the VLSI too soon
- Don't design the VLSI too late
- When you do design the VLSI, design defensively

The first point is obvious. Using HDLs allows a wide range of portability - from FPGAs to ASICs. We have found that even if the lower level details change, the HDL structure allows quick and easy porting from one implementation to another.

No matter what design flow is used, it has to be automated. That is, the flow should be capable of running

unattended without user input. This is necessary due to the long run times of many of the design steps such as synthesis, place and route and full chip extraction. One needs full 24-hour design operations and thus the design flow has to be capable of running overnight unattended. In addition to this automation, the complete chip assembly/verification procedure should be as fast as possible. This requirement arises when the VLSI design has to track a dynamically changing specification and frequent changes to the specification cause changes to the HDL description of the VLSI component (see comments later).

With communication based designs, one method of testing them is to perform BER tests on the complete system. The only way to do this in order to measure small BERs is to have a prototype running in nearly real time. Luckily FPGAs are at a density where this is often possible.

We have found that especially with error correcting modules (such as Viterbi decoders) special test points need to be inserted in order to adequately test the module and not have the module under test mask internally generated errors.

Choosing the right time to commit to silicon is a fine art. In essence, we followed an approach in which the VLSI design followed the prototype FPGA design as closely as possible (see above). In addition, the chip design should be done with a view to predicting hardware limitations and designing around them. For instance, control functions should be made programmable rather than being hardwired.

We have approached this project with design reuse in mind. The use of Verilog to specify the design should yield a degree of process independence. In addition, by using SmartPath, we have attempted to define "firm" modules in the VSI Alliance[7] nomenclature - that is, modules with a predefined placement. At this point we can see some points that need to be addressed. In particular, clocking is an issue. We decided on a standard cell library that included an internal clock line. Moreover, at the chip level, we were sure to route a mesh connected clock to ensure low clock skew to all devices. This departs from the "buffer insertion" strategy favored by the CADENCE tools. So by using a "non-standard" clocking methodology we have made the design less portable (but probably capable of higher performance). We are currently re-evaluating our P&R approach to determine what benefits our current methodology provides over existing methods.

As the desire to improve the level of integration increases, it is natural to determine what sections of the complete modem should co-locate on the same die. The RF, IF and base-band analog sections contain circuitry that strains CMOS digital processes (especially the RF at 5 GHz). Further integration really depends on the type of products that will be built with this technology. It is unlikely in the near future that this type of circuitry can be adapted to the VSIA model.

Careful partitioning of designs will continue to keep overall system performance up and cost down.

V. SUMMARY

This paper has summarized the some of the design aspects of the hardware for a high-speed wireless LAN. The implications of this for systems-on-a-chip applications have also been addressed. While virtual socket components for digital sections are a relatively straight forward problem, the analog and RF sections will prove more taxing to the system designer.

REFERENCES

[1] Skellern, D. J., et. al., "A mm-Wave High Speed Wireless LAN for Mobile Computing - Architecture and Prototype Modem/Codec Implementation," Hot Interconnects IV Stanford University, Stanford, CA, August 1996.

[2]FCC, Notice of Proposed Rule Making, "Amendment of the Commission's Rules to Provide for Unlicensed NII/SUPERNet Operations in the 5 GHz Frequency Range," ET Docket No. 94-102, RM-8648 and RM-8653, Document FCC96-193, 25 April 1996.

3. Weste, N. and Graham, J., "A Tunable CMOS Filter and Tuning Method for a WLAN Modem" - Proceedings 14th Australian Microelectronics Conference, Sept.Oct, Melbourne, Australia.

4. N. Weste, M. Bickerstaff, T. Arivoli, P.J. Ryan, D.J. Skellern, and T.M. Percival, "A 50 MHz 16-Point FFT Processor for WLAN Applications", Custom Integrated Circuits Conference, Santa Clara, CA, May, 1997.

5. Weste, N., Bickerstaff, M., Arivoli, T., and Foyster, G. "MacFLOW: A CMOS Technology Independent HDL Design Flow" - Proceedings 14th Australian Microelectronics Conference, Sept.Oct, Melbourne, Australia.

[6] CADENCE Design Systems, Inc., Product Literature.

[7] Virtual Socket Initiative Alliance, http://www.vsi.org