# Is Hardware Innovation Over?

Arvind
Computer Science and Artificial Intelligence Laboratory
M.I.T.
*and*
WCU Distinguished Professor, Seoul National University

Seoul National University, Seoul, Korea
October 7, 2009

# The power of numbers

- Last year 950M cell phones were sold as opposed to 100M PC
- India & China are each selling > *7M new cell-phone connections per month*
  - In developing countries cell phone is the only computer most people have
  - In the developed world cell phone is the only computer people carry all the time

A shift in research is underway from PCs to cell phone, not very different from the shift from Mainframes and Minis to PCs in early eighties.

# The future would be dominated by the concerns of

- cheap & powerful handheld devices

*and*

- Powerful infrastructure needed to support services on these devices.

# Current Cellphone Architecture

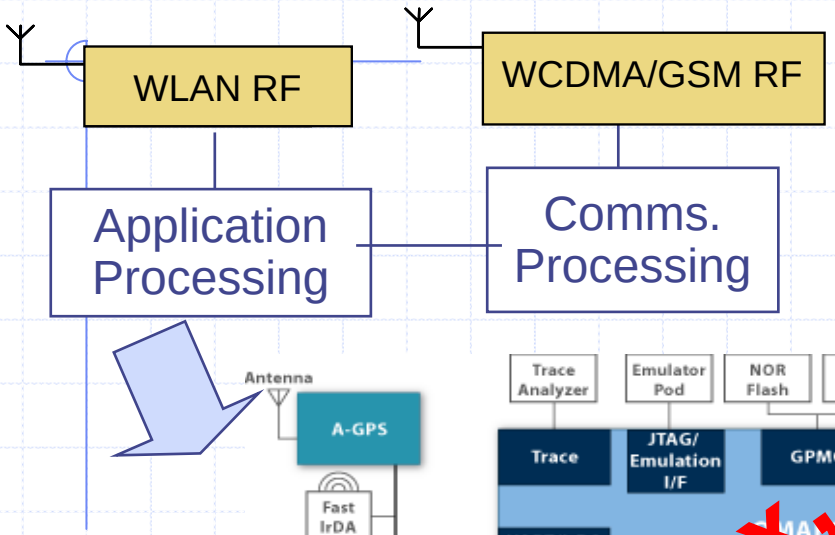


Two chips, each with an ARM general-purpose processor (GPP) and a DSP (TI OMAP 2420)

Many specialized complex blocks


Complex, High Performance but must not dissipate more than 3 watts

# Real power saving implies specialized hardware

- H.264 video decoder implementations in software vs. hardware
  - the power/energy savings could be 100 to 1000 fold

*but our mind set is that hardware design is:*

- Difficult, risky
  - Increases time-to-ma
- Inflexible, brittle, er
  - Difficult to deal with changing standards, …

New design flows and tools can change this mind set

# SoC & Multicore Convergence:

*more application specific blocks*

Application-specific processing units

On-chip memory banks

General-purpose processors

Structured on-chip networks

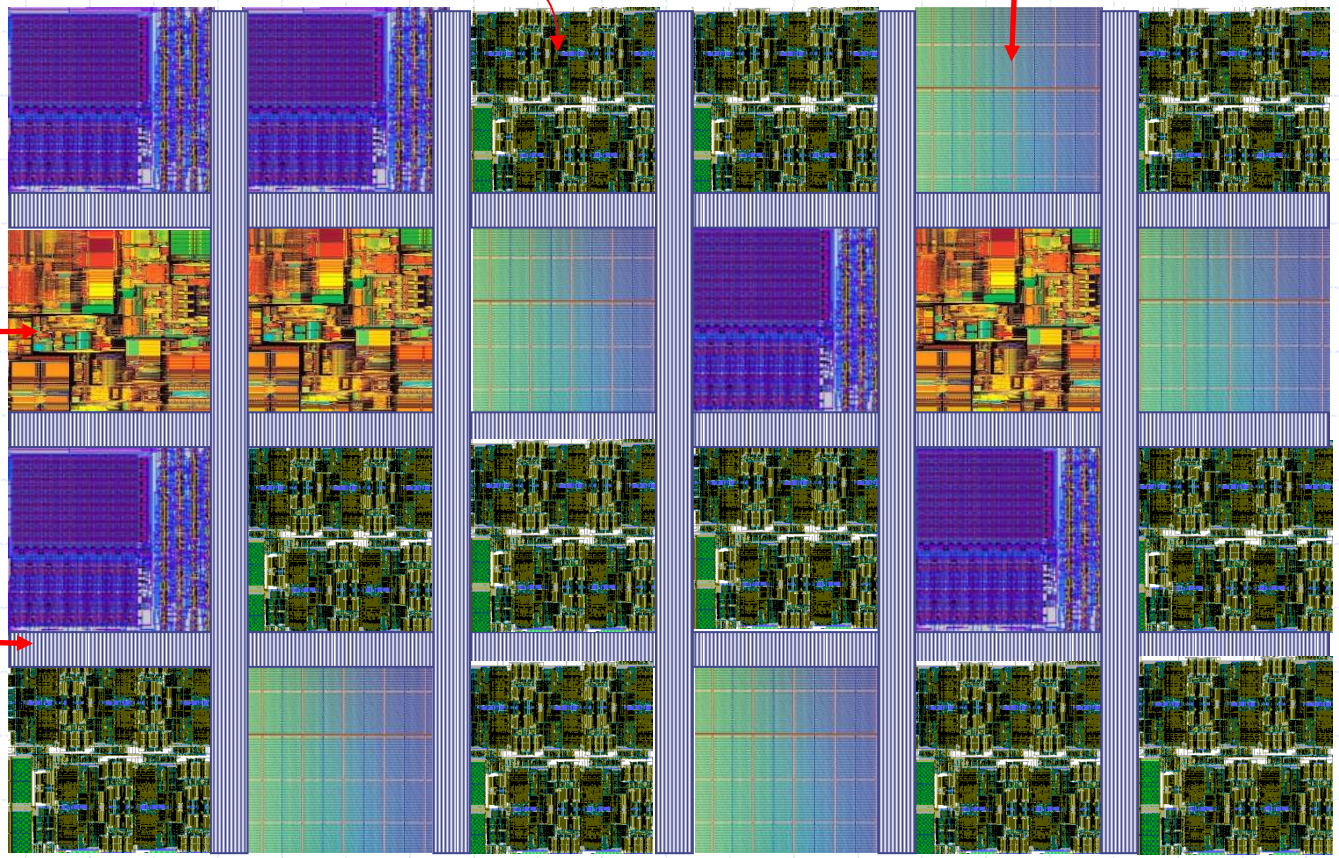

Is consumer space different from enterprise space?

# Server Microprocessors

- Also highly regular multicores with lots of specialized processing capabilities for
  - compression/decompression
  - encryption/decryption
  - intrusion detection and other security related solutions
  - Dealing with spam
  - Self diagnosing errors and masking them
  - ...

*One way to provide these functionalities is via on-chip FPGAs*

# Server Multicore

*more memory, cores, reconfigurable logic…*

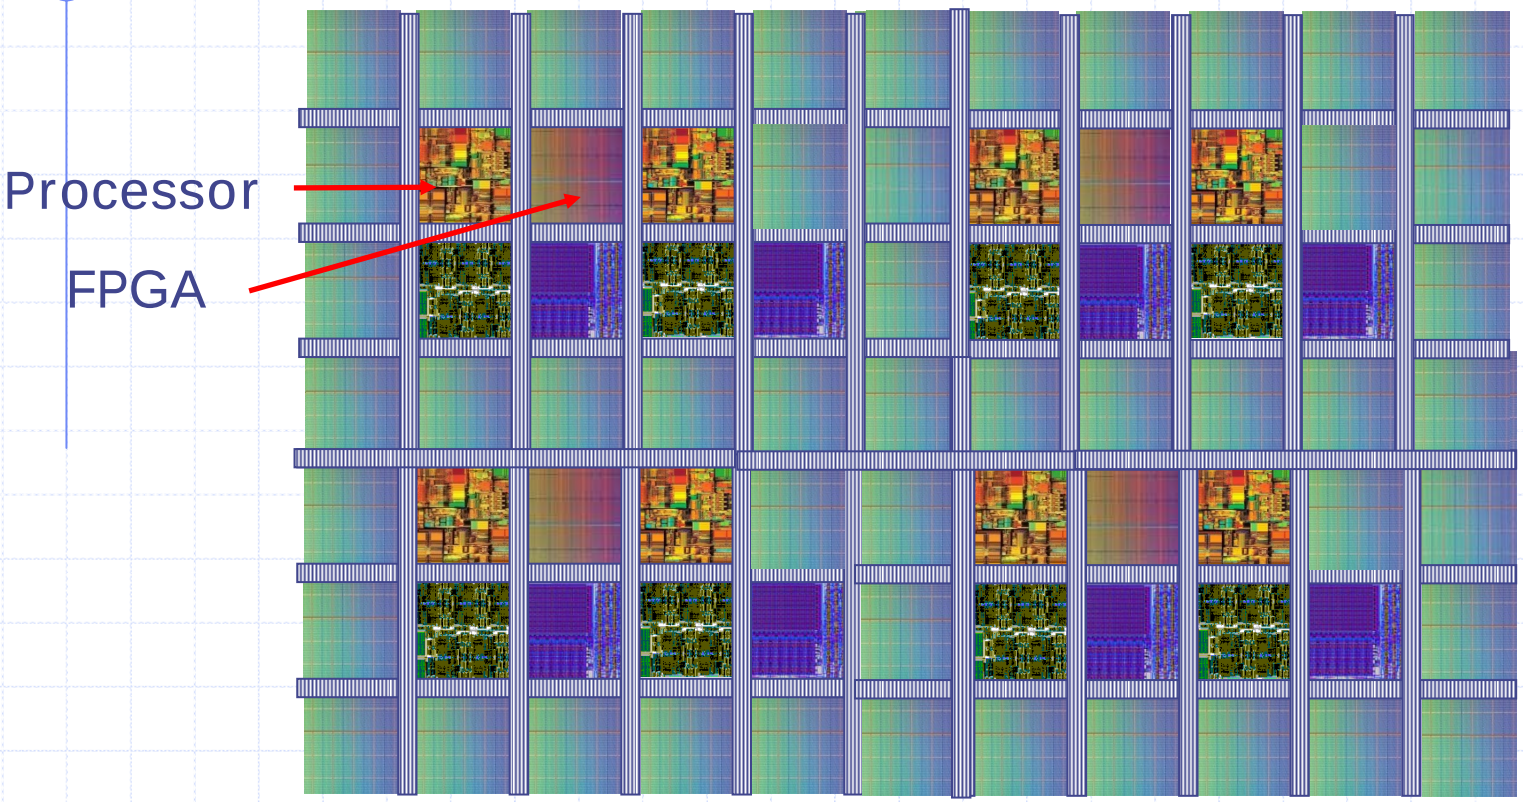


Quality-of-Service (QoS) aware on-chip networks and resource management are essential for guaranteeing performance

# Architectural Renaissance

- Unprecedented opportunity to rethink parallel architectures
- Unprecedented need to design low-power functional blocks
- Unprecedented opportunity to experiment offered by large FPGAs and high-level synthesis tools

# Bluespec A new way of expressing behavior

- A formal method of composing modules with parallel interfaces (ports)
  - Compiler manages muxing of ports and associated control
- Powerful and *zero-cost* parameterization of modules
  - Encapsulation of C and Verilog codes using Bluespec wrappers
  - Helps Transaction Level modeling

➔ *Smaller, simpler, clearer, more correct code*

➔ *not just simulation, synthesis as well*

# High-level Synthesis from Bluespec

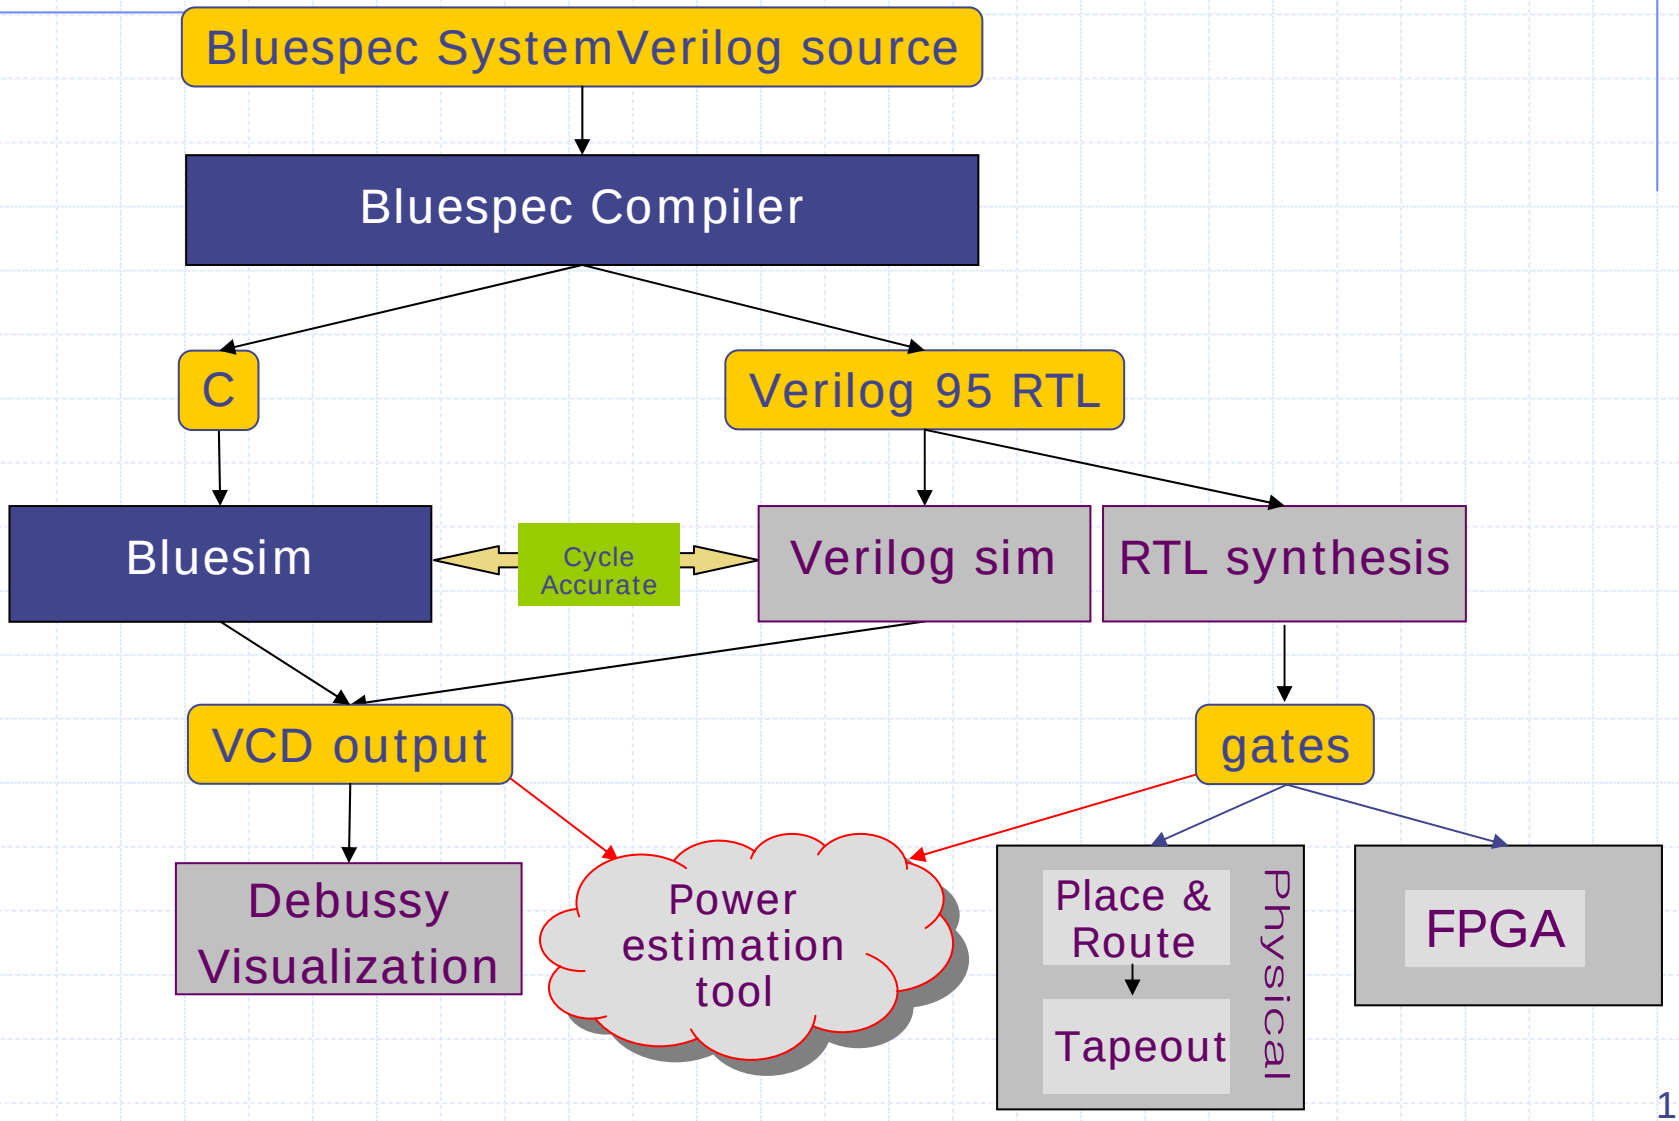

# Bluespec enables

- Extreme IP reuse
  - Multiple instantiations of a block for different performance and application requirements
  - Packaging of IP so that the blocks can be assembled easily to build a large system (black box model)
- Architectural exploration

"Intellectual Property"

*An example*

# Parameterized modules

## *Example OFDM based protocols*

WiFi: 64pt @ 0.25MHz

WiMAX: 256pt @ 0.03MHz

WUSB: 128pt 8MHz

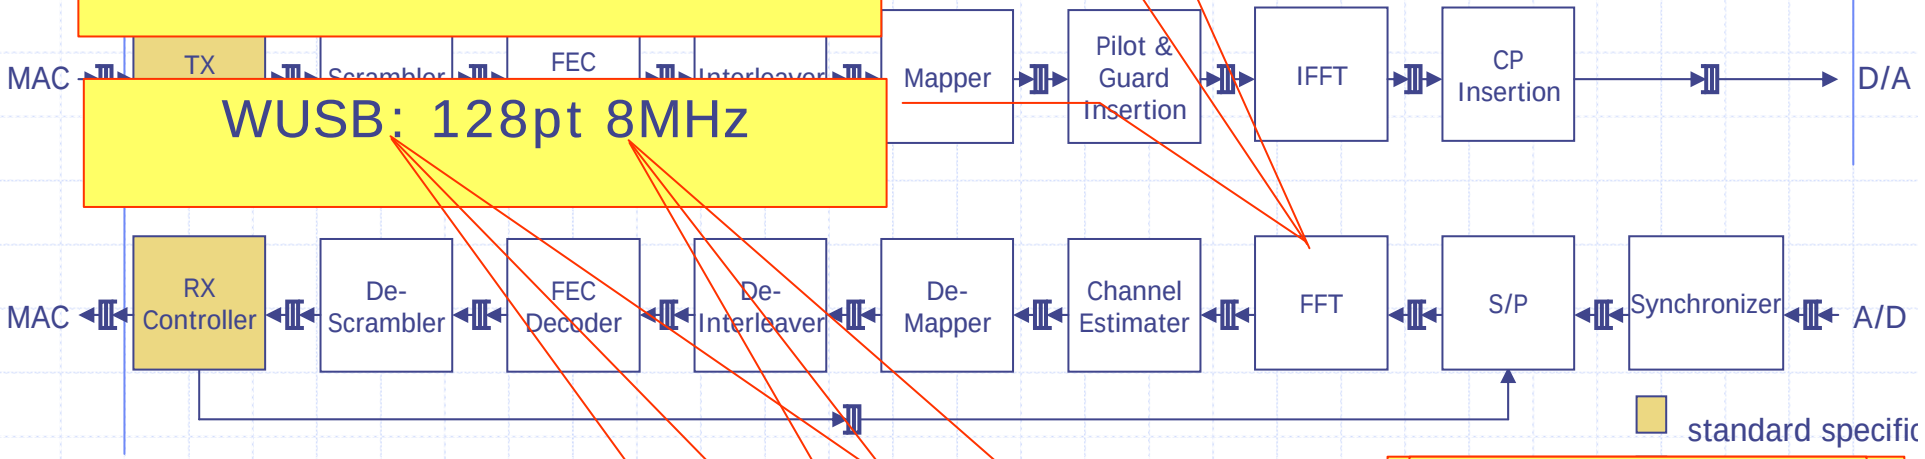


Convolutional

Reed-Solomon

Turbo

- Reusable algorithm with different parameter settings
- [illegible]
- Different algorithms

85% reusable code between WiFi and WiMAX

From WiFi to WiMAX in 4 wee

# 802.11a Transmitter Design:
## *Preliminary results*

| Design Block | Lines of Code (BSV) | Relative Area |
|---|---|---|
| Controller | 49 | 0% |
| Scrambler | 40 | 0% |
| Conv. Encoder | 113 | 0% |
| Interleaver | 76 | 1% |
| Mapper | 112 | 11% |
| IFFT | 95 | 85% |
| Cyc. Extender | 23 | 3% |

Complex arithmetic libraries constitute another 200 lines of code

# FFT – fold to save area

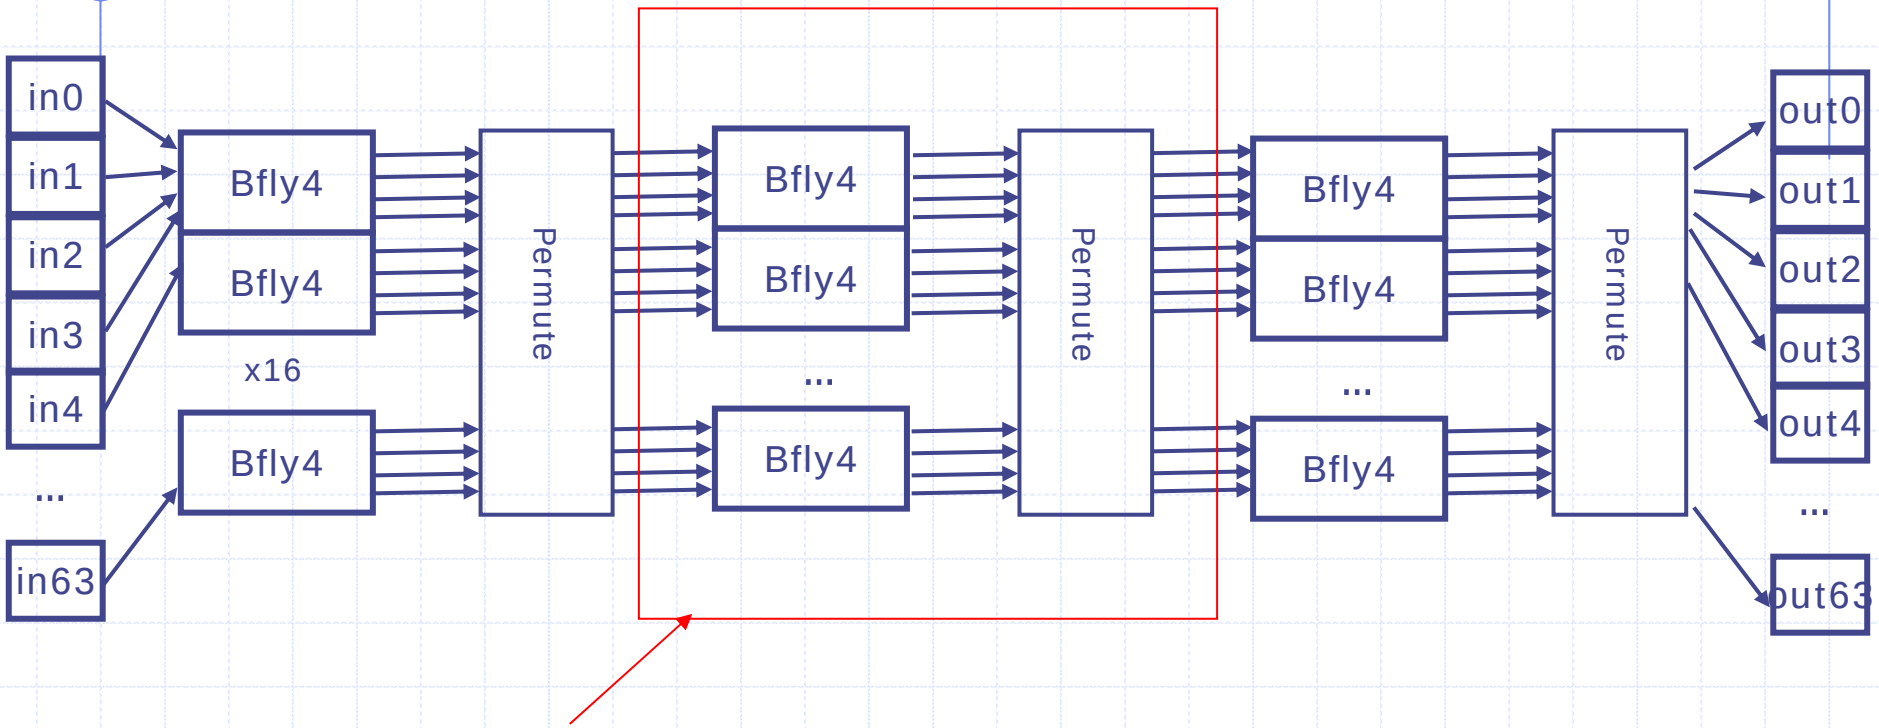


Reuse the same circuit three times to reduce area

# 802.11a Transmitter Synthesis results (Only the IFFT block is changing)

| IFFT Design | Area ($mm^2$) | Throughput Latency (CLKs/sym) | Min. Freq Required |
|---|---|---|---|
| Pipelined | 5.25 | 04 | 1.0 MHz |
| Combinational | **4.91** | 04 | 1.0 MHz |
| Folded (16 Bfly-4s) | **3.97** | 04 | 1.0 MHz |
| Super-Folded (8 Bfly-4s) | 3.69 | 06 | 1.5 MHz |
| SF(4 Bfly-4s) | 2.45 | 12 | 3.0 MHz |
| SF(2 Bfly-4s) | 1.84 | 24 | 6.0 MHz |
| SF (1 Bfly4) | 1.52 | 48 | 12 MHZ |

The same source code

All these designs were done in less than 24 hours!

TSMC .18 micron; numbers reported are before place and route.

# Some cool projects

- Video decoder – H.264
- AirBlue – A new platform to experiment with cross-layer wireless protocols
- IBM PowerPC Prototype and Cycle-accurate performance models
- Hardware software co-generation

# H.264 Video Decoder

Chun-Chieh Lin, K Elliott Fleming [MEMOCODE 2008]

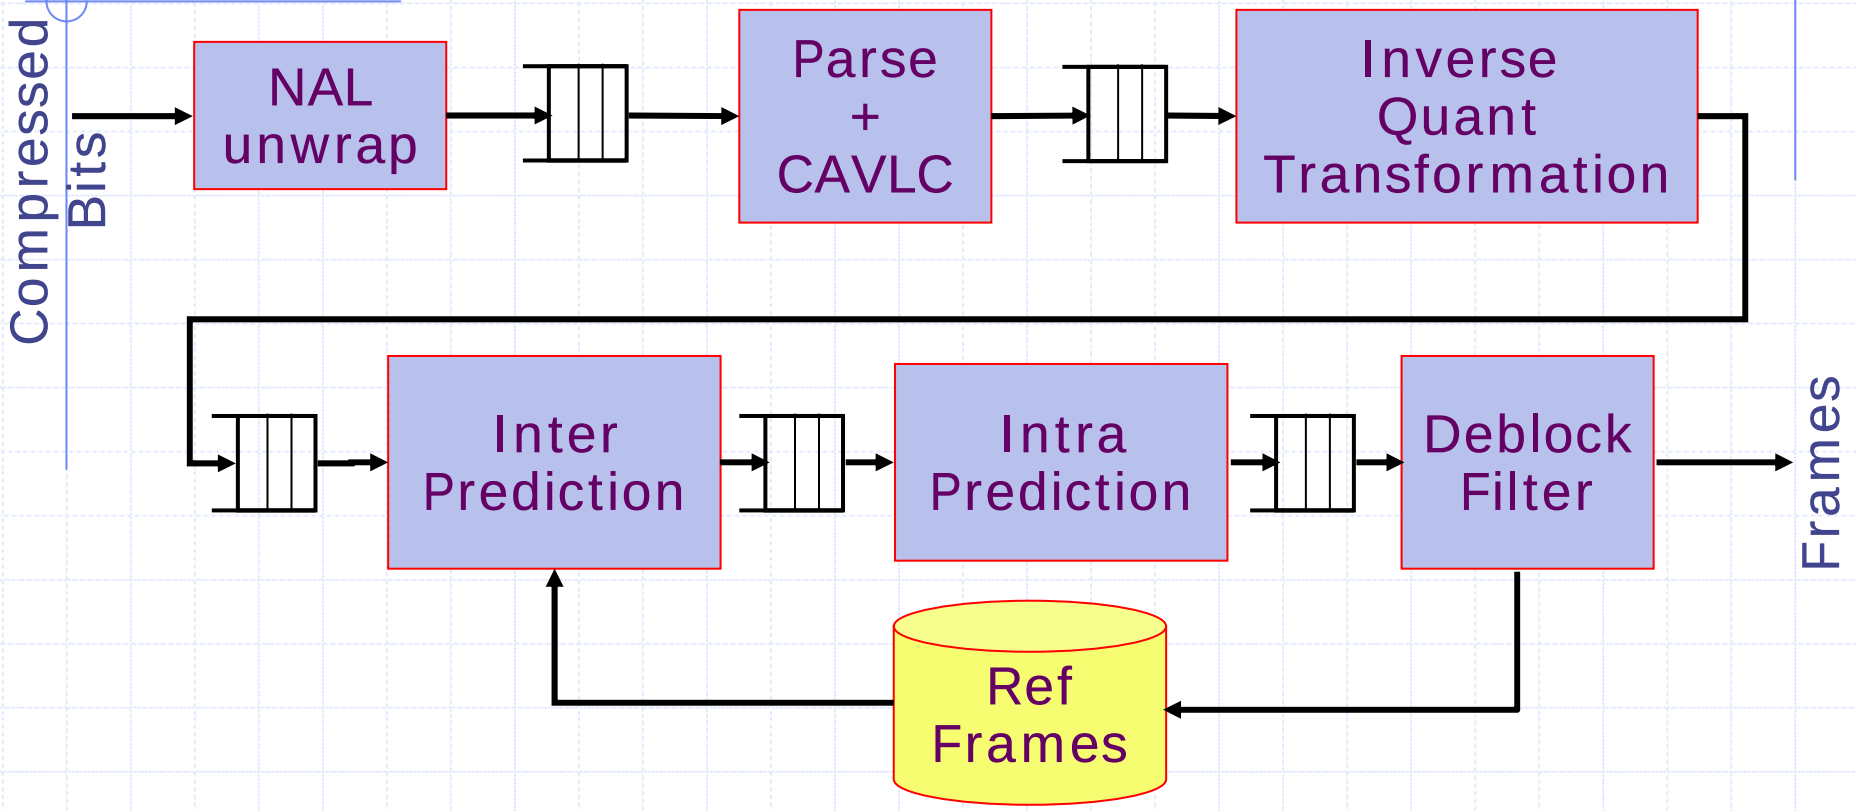


Different requirements for different environments

- QVGA 320x240p (30 fps)
- DVD 720x480p
- HD DVD 1280x720p (60-75 fps)

*May be implemented in hardware or software depending upon ...*

# H.264 in Bluespec

- Initial Design: Base profile
  - Eight man-months
  - 8K lines of Bluespec
    - in contrast to 80K lines of C standard
  - Decoded 720p@32FPS
- Major architectural explorations over 3 months to meet different performance or cost criteria
  - High performance designs (4.2 mm sq in 180nm)
    - 720p@75FPS, 1080p@65FPS,
  - Low cost designs
    - QCIF@15FPS (2.2mm sq), 720p@30FPS (2.4mm sq)

Current focus is on high performance FPGA implementations

# AirBlue: A platform for Cross-Layer Wireless Protocol development

- **Cross-layer protocols** are the hottest area of research in wireless
  - Jointly optimizing PHY, MAC, network layers
- Realistic experimentations are difficult
  - PHY (baseband) layer requires a lot of computation: traditionally in hardware
  - MAC typically done in firmware
  - Higher layers in software

Collaboration with Professor Hari Balakrishnan

# AirBlue Platform: Alfred Ng, Elliott Fleming, Mythli Vutukuku, Ramki Gummadi,

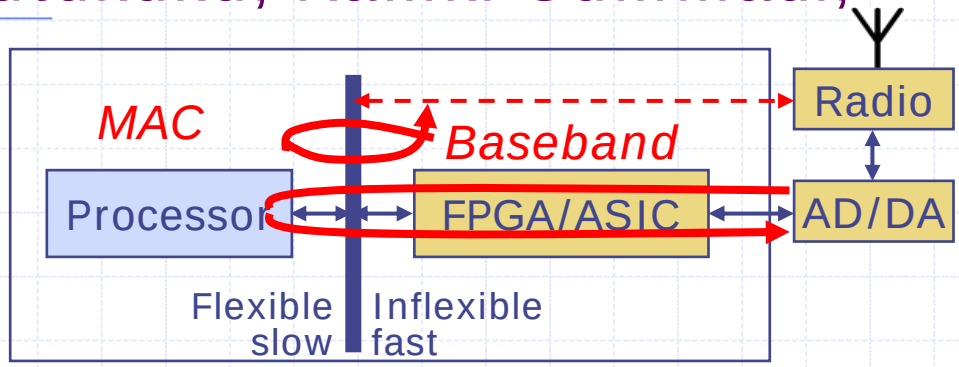


Need short latencies

*Cross-layer wireless protocols* require a platform that offers both flexibility/programmability and performance

Current platforms do not offer both

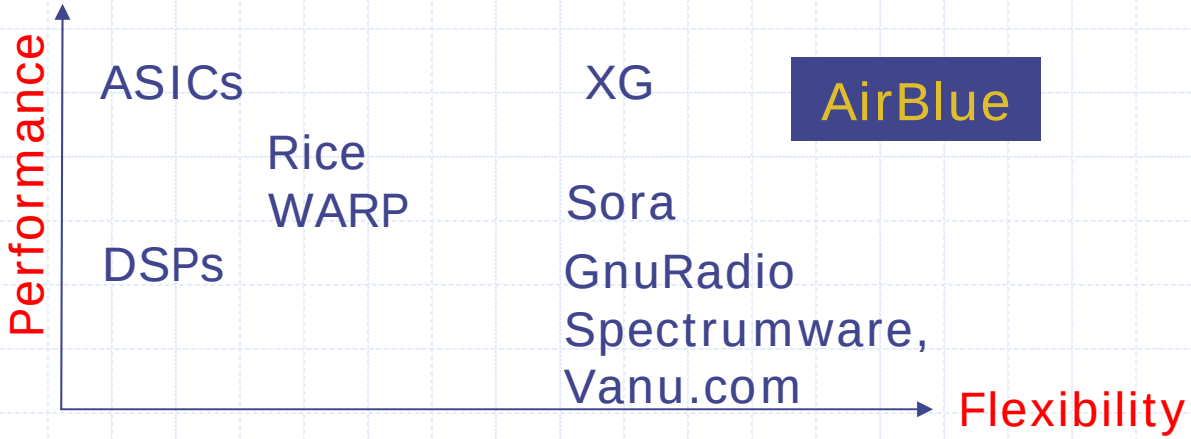

# AirBlue

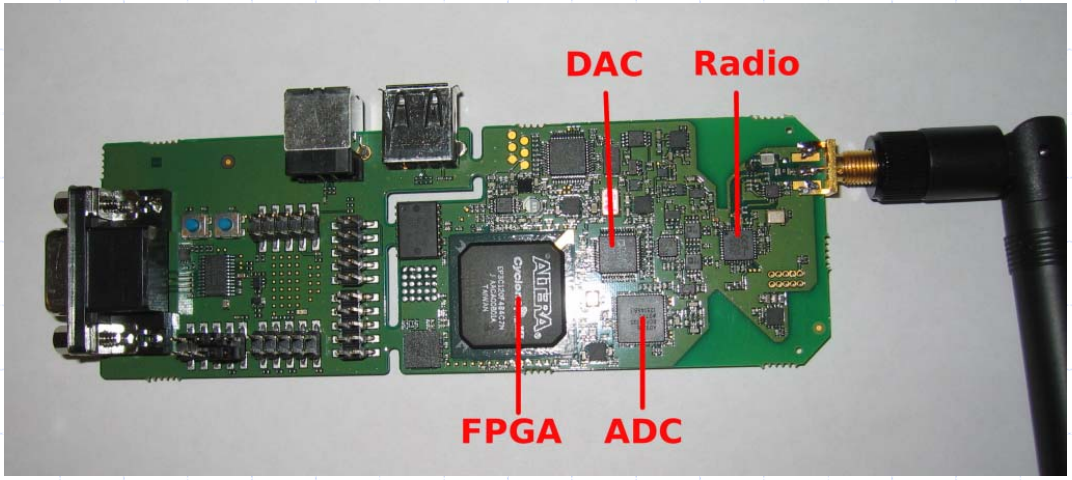


Fits in Nokia N95 phones

- Several cross-layer experiments have already been conducted on full-speed 802.11a/g implementation
  - SoftPHY: Exposes signal quality to higher layers
    - Enables new protocols: MIXIT, PPR, better rate-adaptation
  - Efficient allocation of
    - Variable demands, heterogeneous SNRs

Each new protocol required less than 100 lines of code

# IBM: PowerPC Prototype

K. Ekanadham, Jessica Tseng (IBM)
Asif Khan, M. Vijayaraghavan (MIT)

- Goal: Implement a multithreaded, multicore, in-order PowerPC on an FPGA platform and boot Linux on it in 12 months

- Team:
  - 2(IBM) + 2(MIT) + Linux and FPGA help

The team accomplished the goal (Nov 2008)
- Bluespec PowerPC boots Linux on FPGAs in 10min;
- 100M instructions to reach “Hello World”;
- 15K lines of Bluespec generated 90K lines of Verilog

IBM synthesized the generated Verilog using their tools in 40nm library
– ran at 500MHz on the first try!

# Phase II: IBM/MIT Collaboration

March 2009 –

- Goal: Produce a *cycle-accurate* and parameterized model of multithreaded, multicore PowerPC to run on FPGAs
  - Architecture models in software can be flexible and have high fidelity but tend to be slow
  - Can we gain 1000X speedup by running the models on FPGAs ?
- Use cheaper and widely available FPGA boards
  - Xilix 110 as opposed to 330
- Target open source distribution by summer 2010

Lots of technical challenges

Currently trying to boot linux

# Could we have done these projects in C, C++, SystemC?

# Hardware synthesis from C does not work very well:

## Reed Solomon Results

WiMAX requirement is to support a throughput of 134Mbps

| | Bluespec | C-synthesis | Xilinx IP |
|---|---|---|---|
| Equivalent Gate Count | 267,741 | 596,730 | 297,409 |
| Frequency (MHz) | 108.5 | 91.2 | 145.3 |
| Steady State (Cycles/Block) | 276 | 2073 | 660 |
| Data rate (Mbps) | 701.3 | 89.7 | 392.8 |

Lower is better

Higher is better

For the same area!

# Hardware innovation is *far from over*

- Ubiquitous mobile devices and demand for new services are ushering in a new era of computing
- Large FPGAs are offering an unprecedented opportunity to experiment
- High-level synthesis tools like Bluespec are making architecture exploration and SoC development much easier
  - High quality synthesis
  - Modules with formal interfaces (not just wires)
  - Parameterized modules (higher-order functions)
  - Strong type system
  - Ability to interact with modules written in C, Verilog, …

*Thanks!*

# Exploiting Multiple Clock Domains in Bluespec for Hw/Sw cogeneration

- MCD allows us to run parts of the design at different speeds
- Each GAA/Method is associated with a clock
- Special Module to Cross Clocks
- The idea works even if some of the domains are implemented in software