
ECE 747

Digital Signal Processing Architecture

DSP Implementation Architectures

Spring 2006
W. Rhett Davis
NC State University



My Goal

- Challenge you to use the techniques you have learned in this class to design the next generation of DSP hardware
- When you undertake a new design, the most important question for you to answer is whether or not it will work better than an existing design.
 - » Faster
 - » Longer Battery Life
 - » Cheaper

Today's Lecture

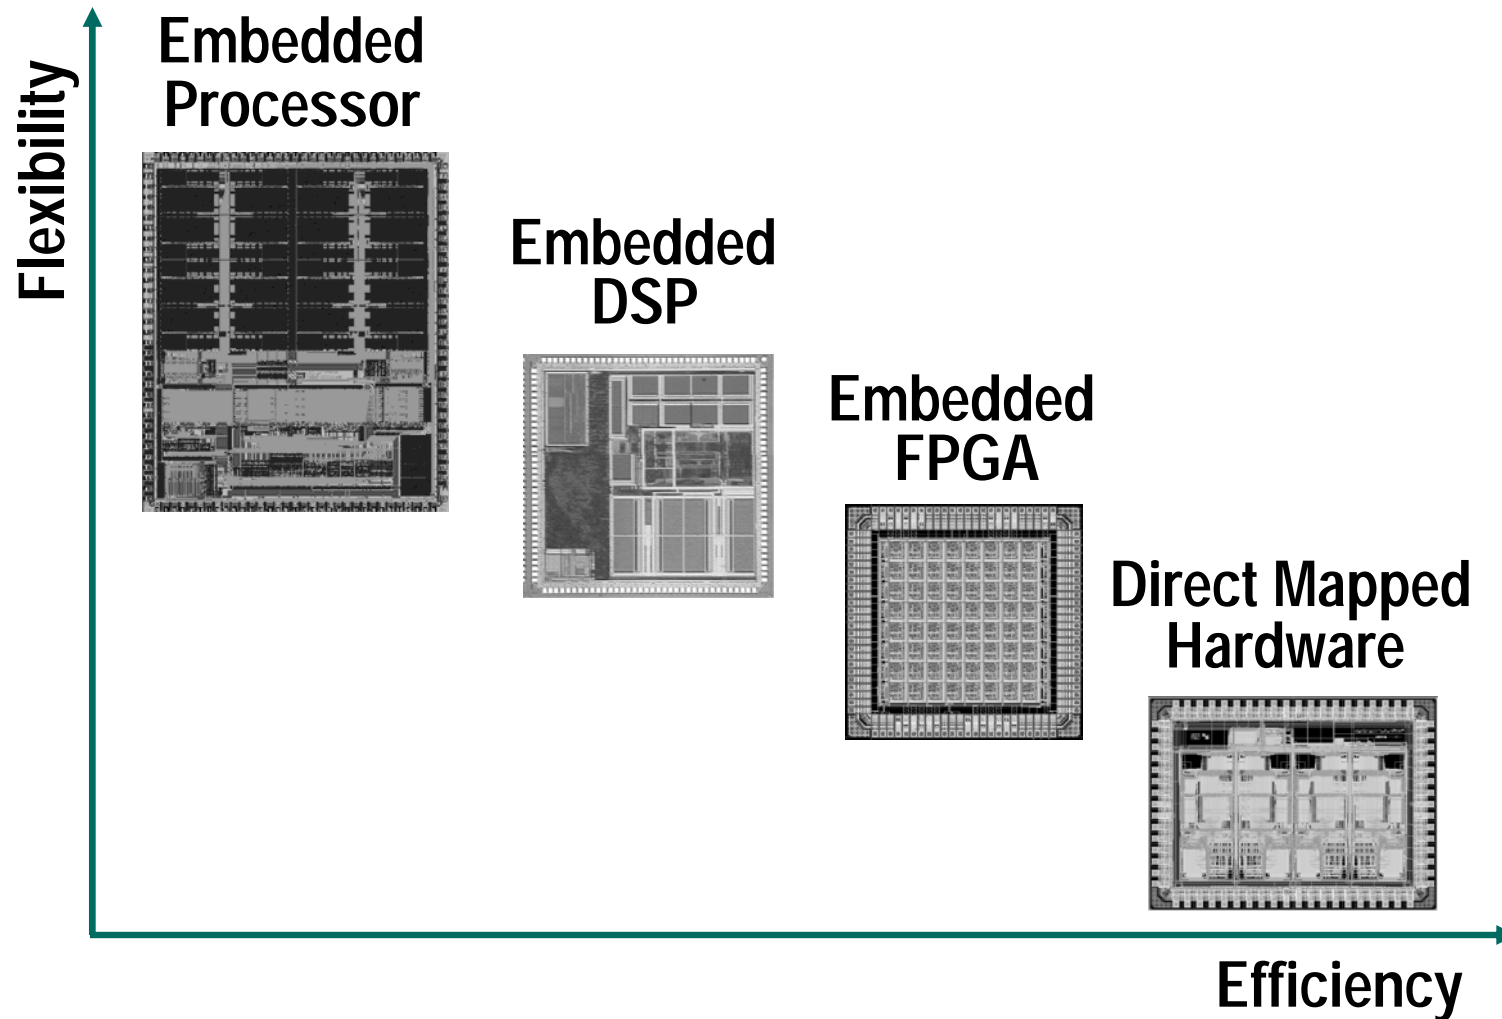
- Types of DSP Implementations
- Comparison of Hardware Efficiency
- The Promise of Systems-on-Chip
- What's keeping us from getting there?

ECE 747-Style Design



- Up to now, you have been designing signal-flow graphs and converting them into hardware, through a design process some call direct-mapping of algorithms
- But what are the other choices?

Efficiency-Flexibility Trade-Off




Computational Efficiency Metrics

- **Definition: MOPS**
 - » Millions of algorithmically defined arithmetic operations (e.g. multiply, add, shift) – in a GP processor several instructions per “useful” operation
- **Figures of merit**
 - » MOPS/mW - Energy efficiency (battery life)
 - » MOPS/mm² - Area efficiency (cost)

Optimization of these “efficiencies” is the basic goal assuming functionality is met

Dedicated Designs 10X-100X More Efficient

- Low Energy Efficiency / Battery Life (MOPS/mW)
due to overhead, lack of parallelism, high supply voltages



(Brodersen, ISSCC 2002)

Potential of Direct Mapping


In .25 micron a multiplier requires .05 mm² and 7pJ per operation at 1 V. Adders and registers are about 10 times smaller and 10 times lower energy

Lets implement a 50mm² , .25 micron chip using adders, registers and multipliers


- We can have 2000 adders/registers and 200 multipliers in less than 1/2 of the chip, also assume 1/3 of power goes into clocks
- 25 MHz clock (1 volt) gives ~50 Gops at 100mW
- **500 MOPS/mW** and **1000 MOPS/mm²**

Why is Direct Mapping Better?

- Low Area Efficiency / High Cost (MOPS/mm²) due to large on-chip memories
- Low Energy Efficiency due to long wires and overhead of multiplexing the datapath



**16x16 multiplier
(.05 mm²)**



**DSP processor with 1 multiplier
(25 mm²)**

Results in Fully Parallel Solutions


Reducing supply voltage saves energy: $E = CV^2$

	Energy		Area	
	64-point FFT Energy per Transform (nJ)	16-State Viterbi Decoder Energy per Decoded bit (nJ)	64-point FFT Transforms per second per unit area (Trans/ms/mm ²)	16-State Viterbi Decoder Decode rate per unit area (kb/s/mm ²)
Direct-Mapped Hardware	1.78	0.022	2,200	200,000
FPGA	683	5.5	1.8	100
Low-Power DSP	436	19.6	4.3	50
High-Performance DSP	1700	108	10	150


(numbers taken from vendor-published benchmarks) (N. Zhang)

Orders of magnitude lower efficiency
even for an optimized processor architecture

Reconfigurability can get the best of both worlds




Energy per Transform
vs. FFT size



Transforms per Second per mm²
vs. FFT size

* All results are scaled to 0.18 μ m

Enter the Era of Systems-on-Chip




Cellular Phone Baseband SOC

(Courtesy Mike McMahon, Texas Instruments)

What is a System-on-Chip?

- “a complex IC that integrates the major functional elements of a complete end-product into a single chip... incorporates at least one programmable processor, on-chip memory, and accelerating function-units....”
 - » *Winning the SoC Revolution*,
Martin & Chang 2003
(paraphrasing from Dataquest)

Problem: Design Productivity Gap



“The main message in 2001 is this: *Cost of design is the greatest threat to continuation of the semiconductor roadmap*” – ITRS

\$20M Average cost for an SOC (includes only software licenses & salaries) – ITRS

Next Lectures

- Survey of System-Level Design Techniques
 - » What tools can I use to get performance estimates faster, with less work?
- Methods of Scaling
 - » How do I convert my 180 nm performance estimates into the latest technology?