A 167-processor 65 nm Computational Platform with Per-Processor Dynamic Supply Voltage and Dynamic Clock Frequency Scaling

Dean Truong, Wayne Cheng, Tinoosh Mohsenin, Zhiyi Yu, Toney Jacobson, Gouri Landge, Michael Meeuwsen, Christine Watnik, Paul Mejia, Anh Tran, Jeremy Webb, Eric Work, Zhibin Xiao and Bevan Baas

VLSI Computation Lab
University of California, Davis

Outline

- **Background and the First Generation AsAP**
- The Second Generation AsAP
  - Processors and Shared Memories
  - On-chip Communication
  - DVFS
- Analysis and Summary

Project Motivation

- Fully programmable and reconfig. architecture
- High energy efficiency and performance
- Exploit task-level parallelism in:
  - Digital Signal Processing
  - Multimedia

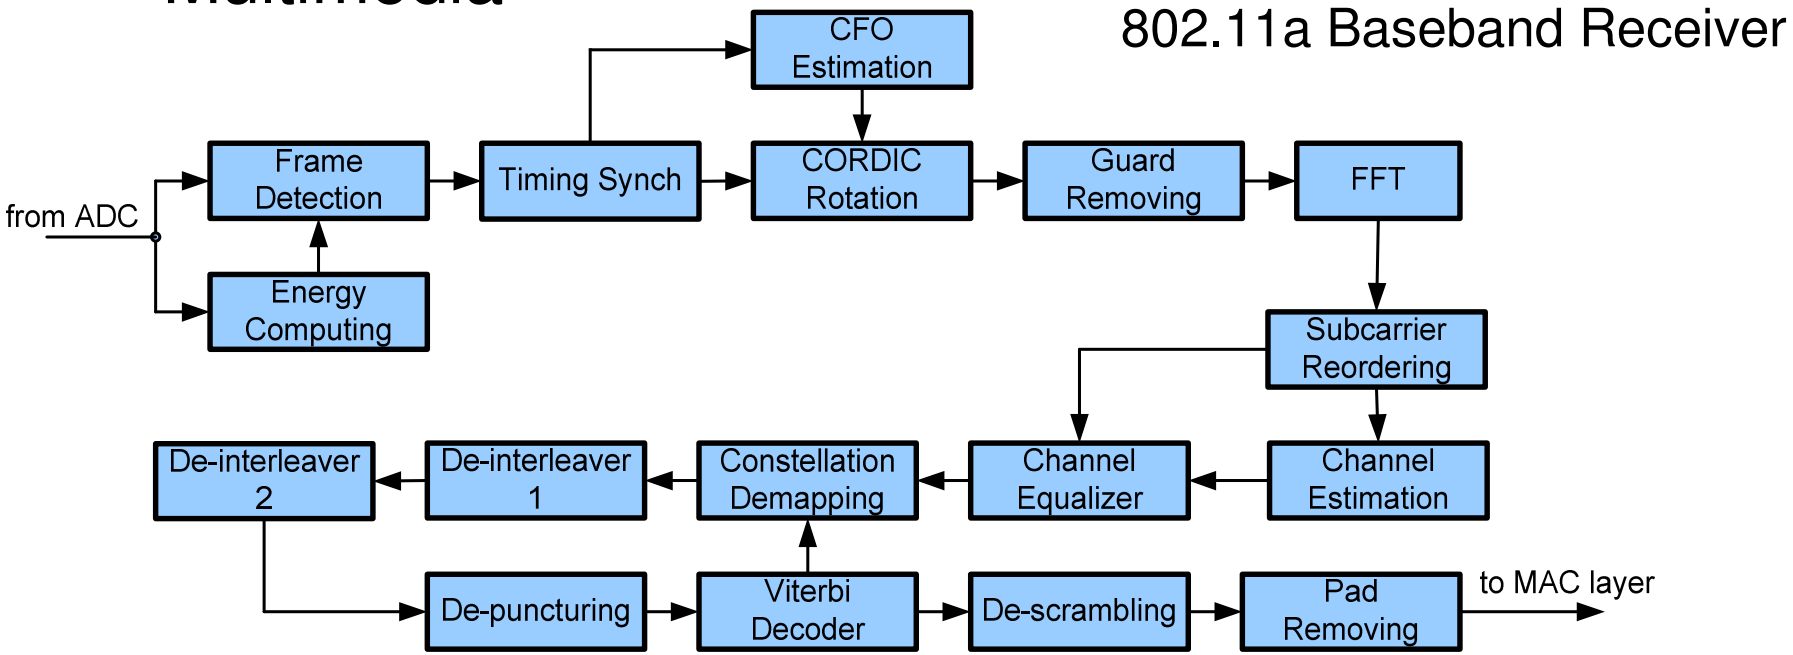

Asynchronous **A**rray of Simple **P**rocessors (AsAP)

- Key Ideas:
  - Programmable, small, and simple fine-grained cores
  - Small local memories sufficient for DSP kernels
  - Globally Asynchronous and Locally Synchronous (GALS) clocking
    - Independent clock frequencies on every core
    - Local oscillator halts when processor is idle

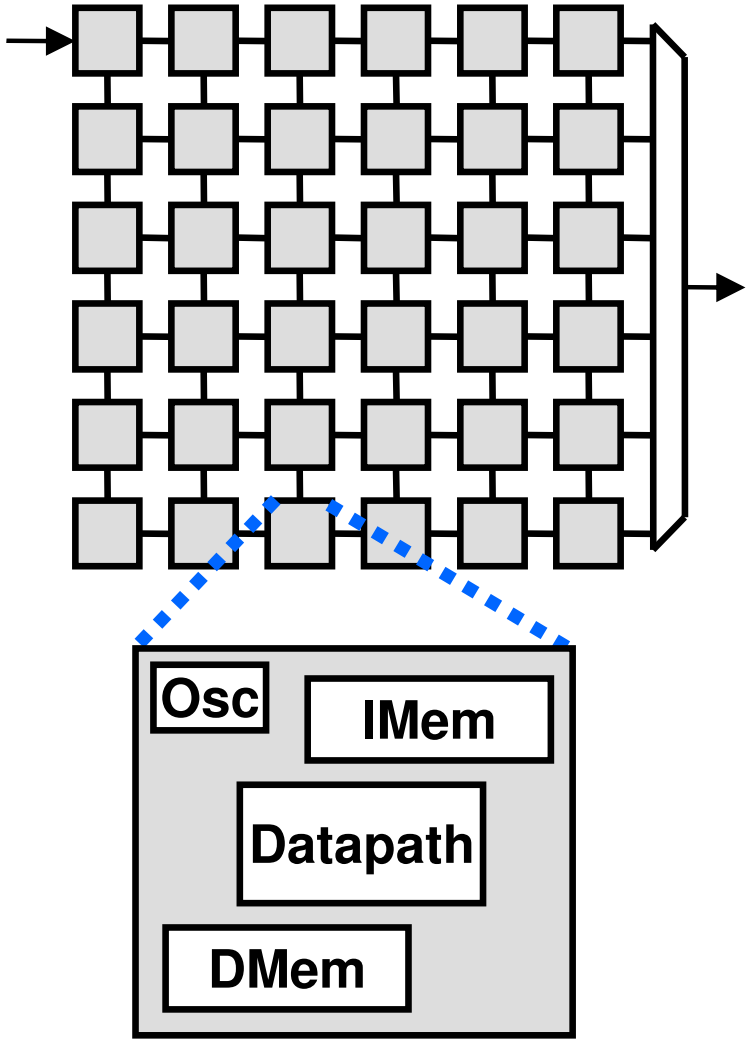

Asynchronous **A**rray of Simple **P**rocessors (AsAP)

- Key Ideas, con't:
  - 2D mesh, circuit-switched network architecture
    - Nearest-neighbor communication only
    - Low area overhead
    - Easily scalable array
  - Increased tolerance to process variations
- 36-processor fully-functional chip, 0.18 µm, 610 MHz @ 2.0 V, 0.66 mm^2 per processor, 802.11a/g tx consumes 407 mW @ 300 MHz
  [ISSCC 06, HotChips 06, IEEE Micro 07, TVLSI 07, JSSC 08,…]

Outline

- Background and the First Generation AsAP
- **The Second Generation AsAP**
  - **Processors and Shared Memories**
  - **On-chip Communication**
  - **DVFS**
- Analysis and Summary

New Challenges Addressed

- Reduction in the power dissipation of
  - Lightly-loaded processors (lowering *Vdd*)
  - Unused processors (leakage)
- Achieving very high efficiencies and speed on common demanding tasks such as FFTs, video motion estimation, and Viterbi decoding
- Larger, area efficient on-chip memories
- Efficient, low overhead communication between distant processors

167-processor Computational Platform

- Key features
  - 164 Enhanced programmable processors
  - 3 Dedicated-purpose processors
  - 3 Shared memories
  - Long-distance circuit-switched communication network
  - Dynamic Voltage and Frequency Scaling (DVFS)

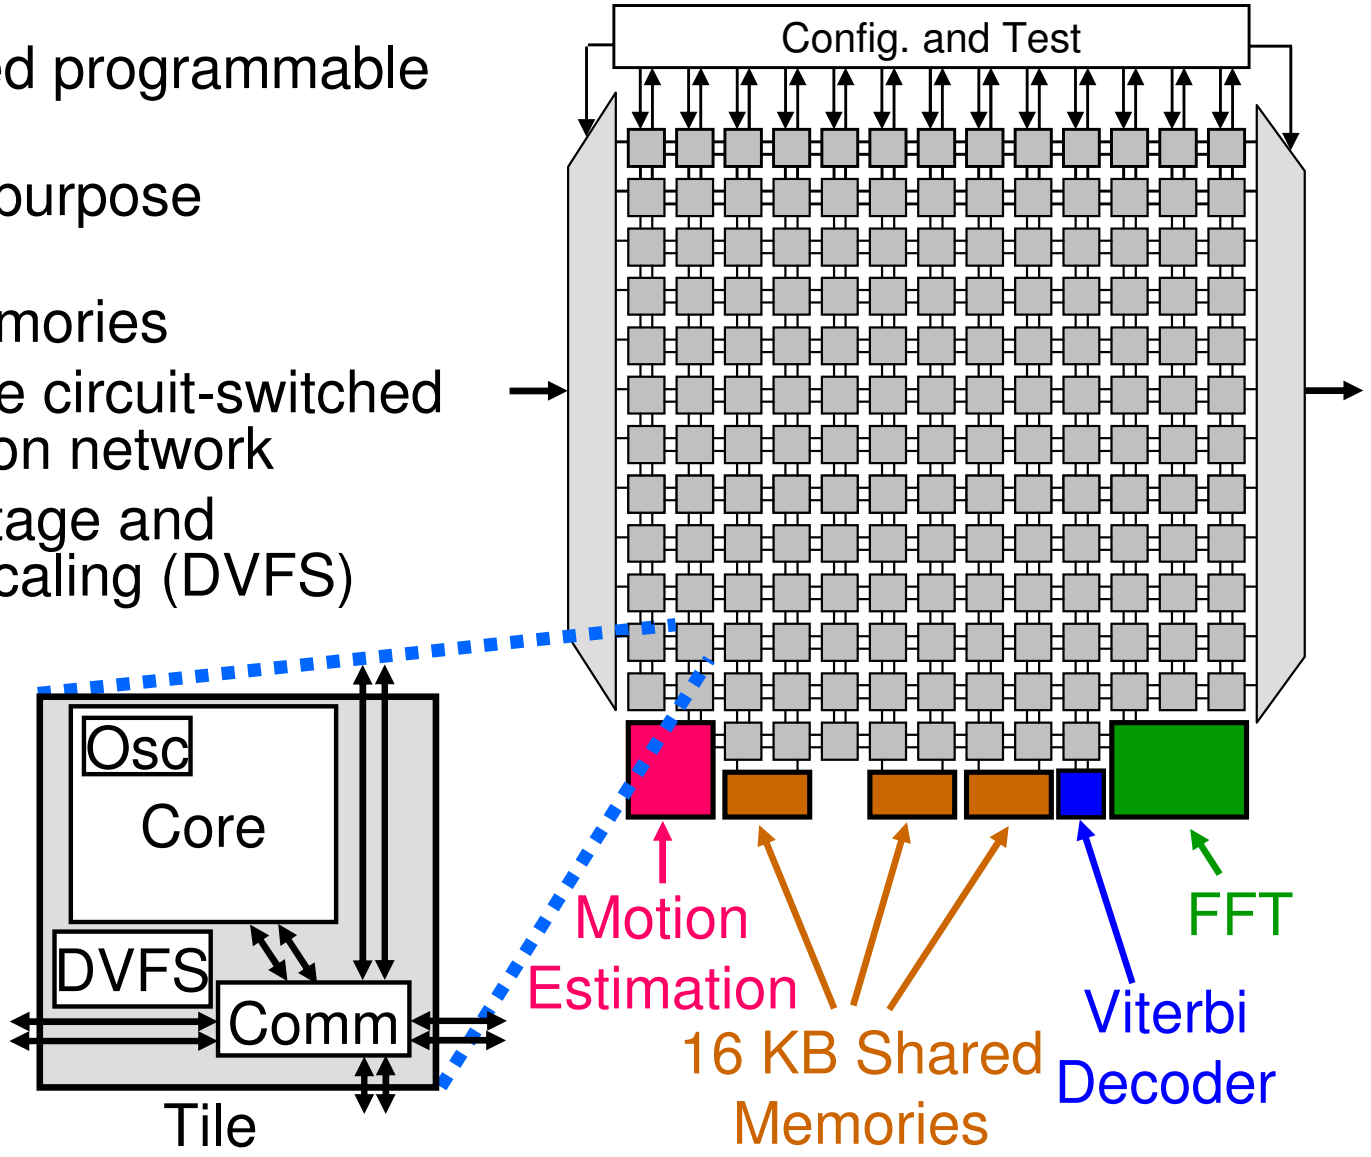

Homogenous Processors

- 164 in-order, single-issue, 6-stage processors
  - 16-bit datapath with RISC-like instructions
  - 128x16-bit data memory (DMEM)
  - 128x35-bit instruction memory (IMEM)
  - Two 64x16-bit FIFOs for inter-processor communication

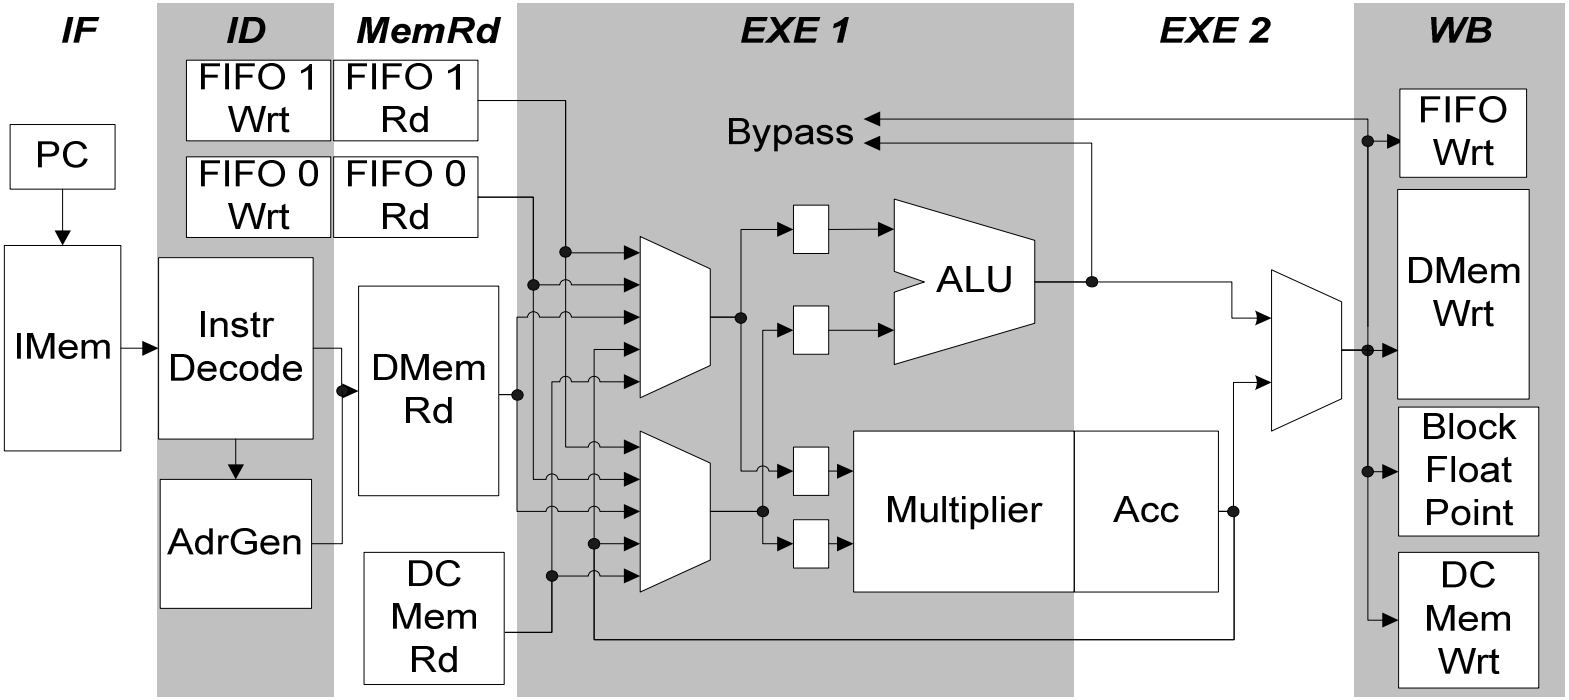

Homogenous Processors

- Over 60 basic instructions
  - Add, Sub, Logic, Multiply, MAC, Branch, …
- New instructions and features
  - Min/Max, Byte-Sub/Add, Absolute value, Fixed-to-Float conversion assist
  - Jump/Return (function support)
  - Zero Overhead Looping (block repeat)
  - Conditional Execution (predicating)
  - Block Floating Point
- Floating point CORDIC square root requires 2.9x fewer cycles compared to first generation AsAP
- Preliminary results from one chip: 1.2 GHz, 59 mW, 1.3 V, 100% active MAC/ALU operations

Fast Fourier Transform (FFT)

- Continuous flow architecture with a single radix-4,2 butterfly
- Runtime configurable from 16-pt to 4096-pt transforms, FFT and IFFT
- 760,000 1024-pt complex FFTs/sec @ 989 MHz, 1.3 V
- 1.01 mm^2
- Preliminary measurements functional at 866 MHz, 34.97 mW @ 1.3 V

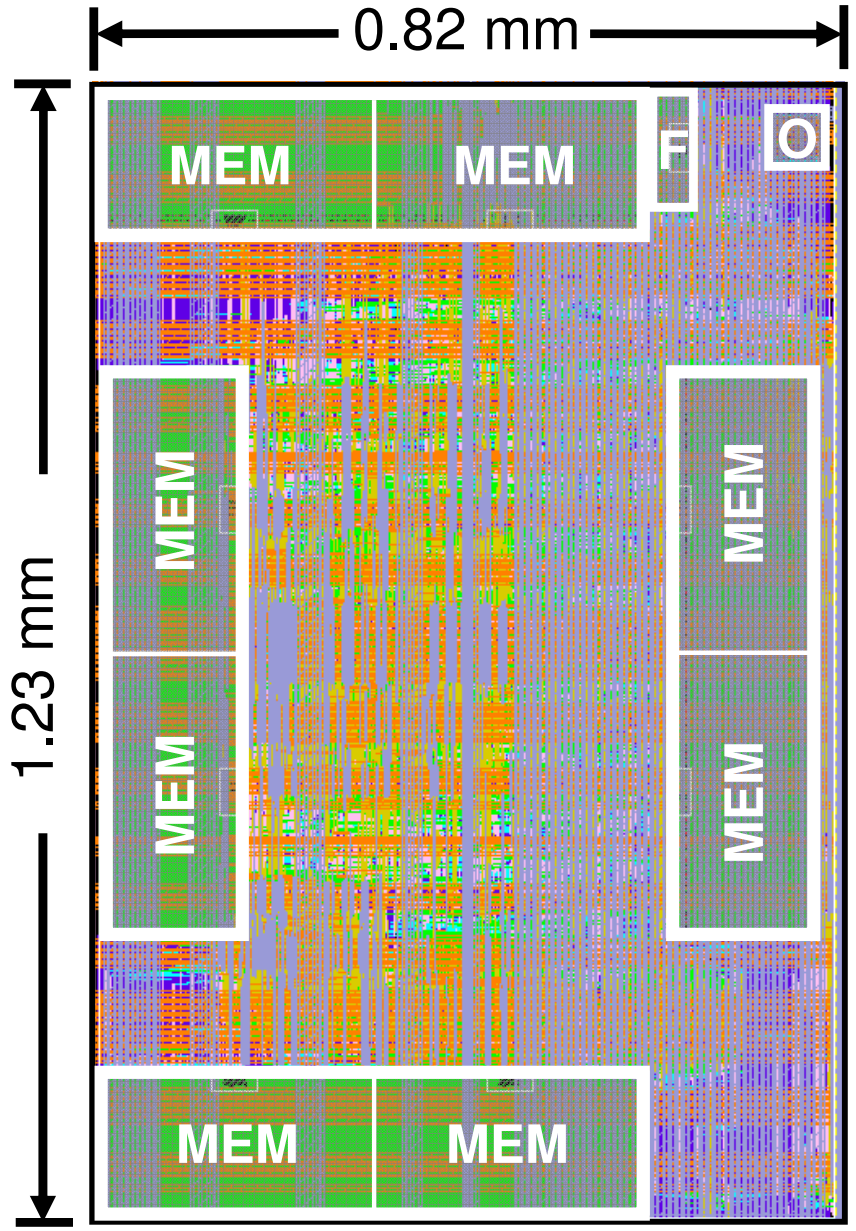

Viterbi Decoder

- 8 Add-Compare-Select (ACS) units
- Highly configurable
  - Up to 32 different rates, including 1/2 and 3/4
  - Decode codes up to constraint length 10
- 72 Mbps @ 789 MHz, 1.3 V for rate = 1/2
- 0.17 mm^2
- Preliminary measurements functional at 894 MHz, 17.55 mW @ 1.3 V

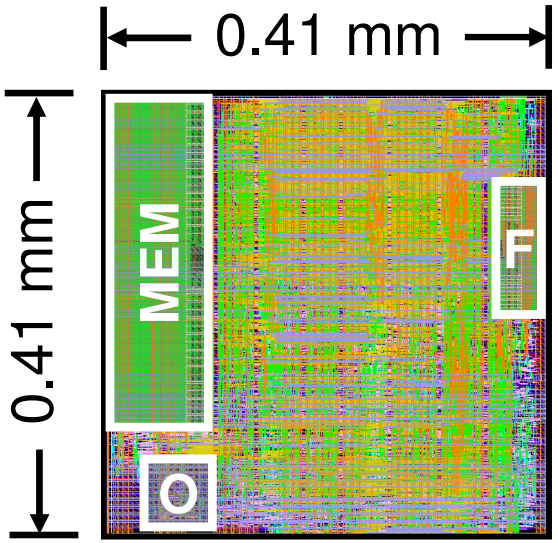

Motion Estimation for Video Encoding

- Supports a number of fixed and programmable search patterns
- Supports all H.264 specified block sizes within a 48x48 search range
- 14 billion SADs/sec @ 880 MHz, 1.3 V; supports 1080p HDTV @ 30fps
- 0.67 mm^2
- Preliminary measurements functional at 938 MHz, 196.17 mW @ 1.3 V

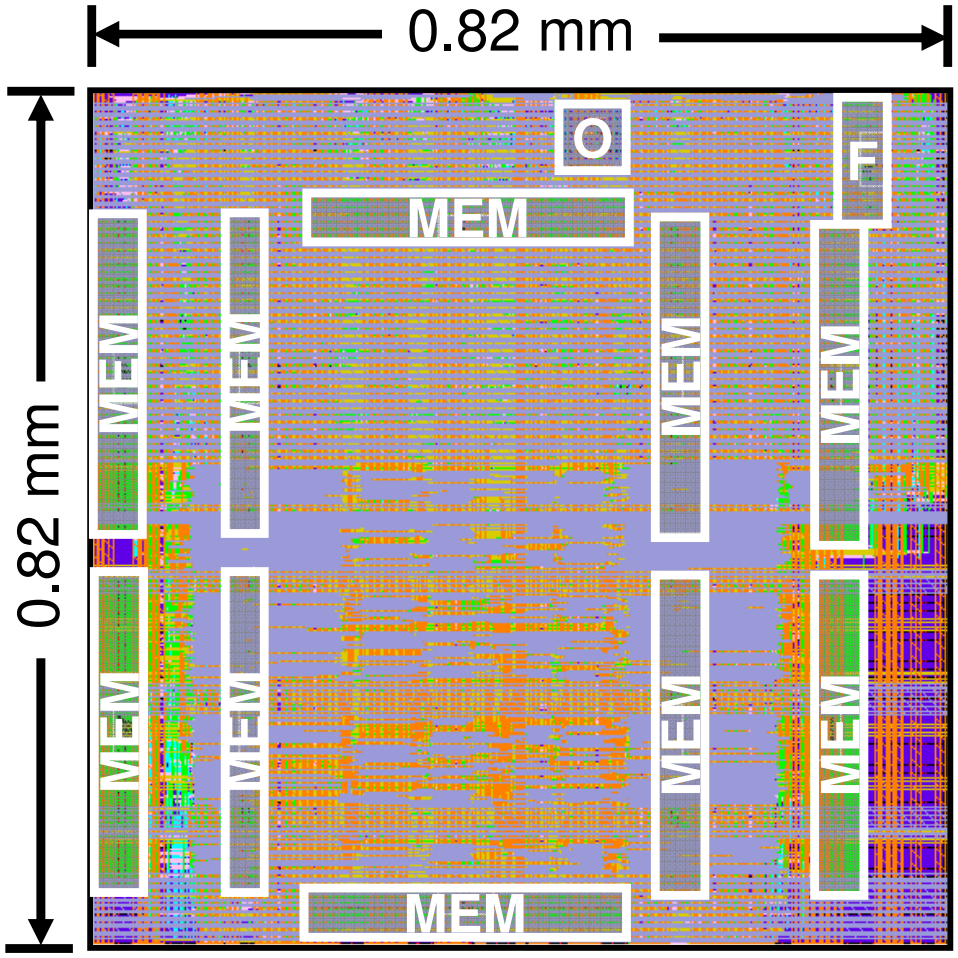

Shared Memories

- Ports for up to four processors (two connected in this chip) to directly connect to the block, which provides
  - Port priority
  - Port request arbitration
  - Programmable address generation supporting multiple addressing modes
- Uses a 16 KByte single-ported SRAM
- 1.28 GHz operation, 1.3 V
  - One read or write per cycle
  - 20.5 Gbps peak throughput
- 0.34 mm^2
- Preliminary measurements functional at 1.3 GHz, 4.55 mW @ 1.3 V

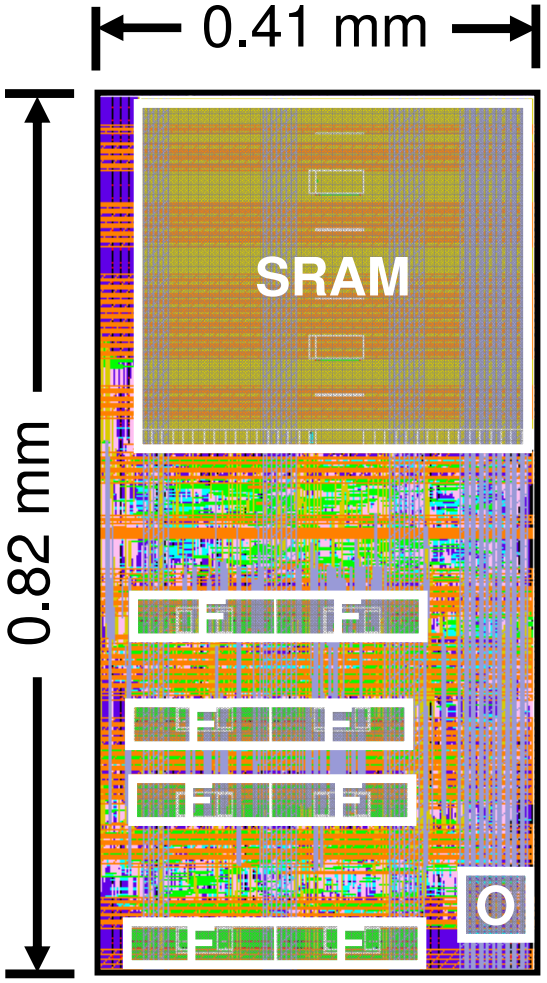

Inter-Processor Communication

- Circuit-switched source-synchronous communication
  - Each link has a *clk*, 16-bit *data* bus, *valid*, and *request*
  - Core can
    - Write to any combination of the 8 outputs under software control
    - Read from any 2 of the 8 inputs using statically configured FIFOs

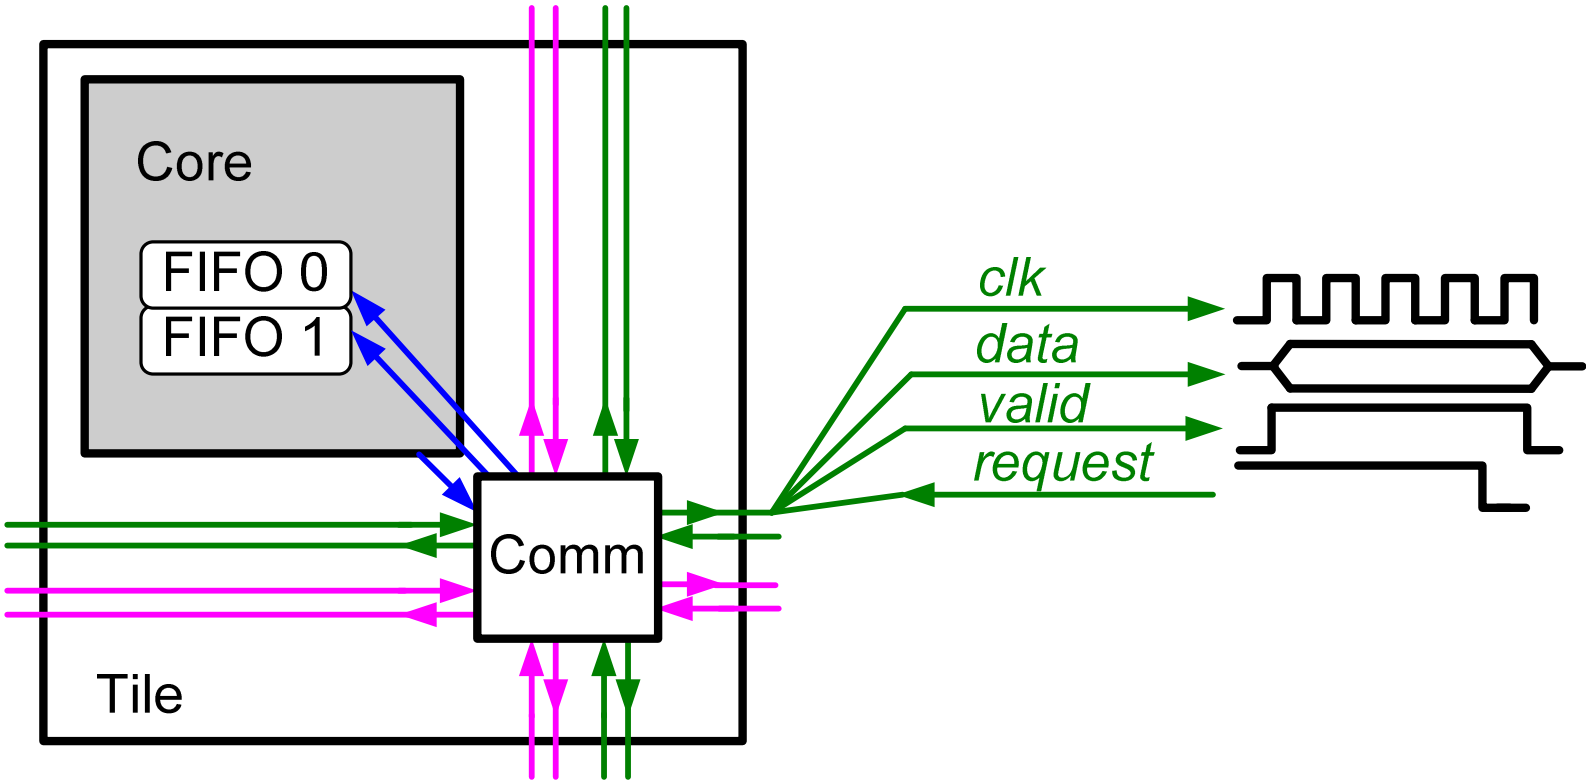

Long-Distance Communication

- Allows communication across tiles without disturbing cores
  - Long-distance links may be pipelined or not
    - Depending on: source clock frequency, distance, and latency

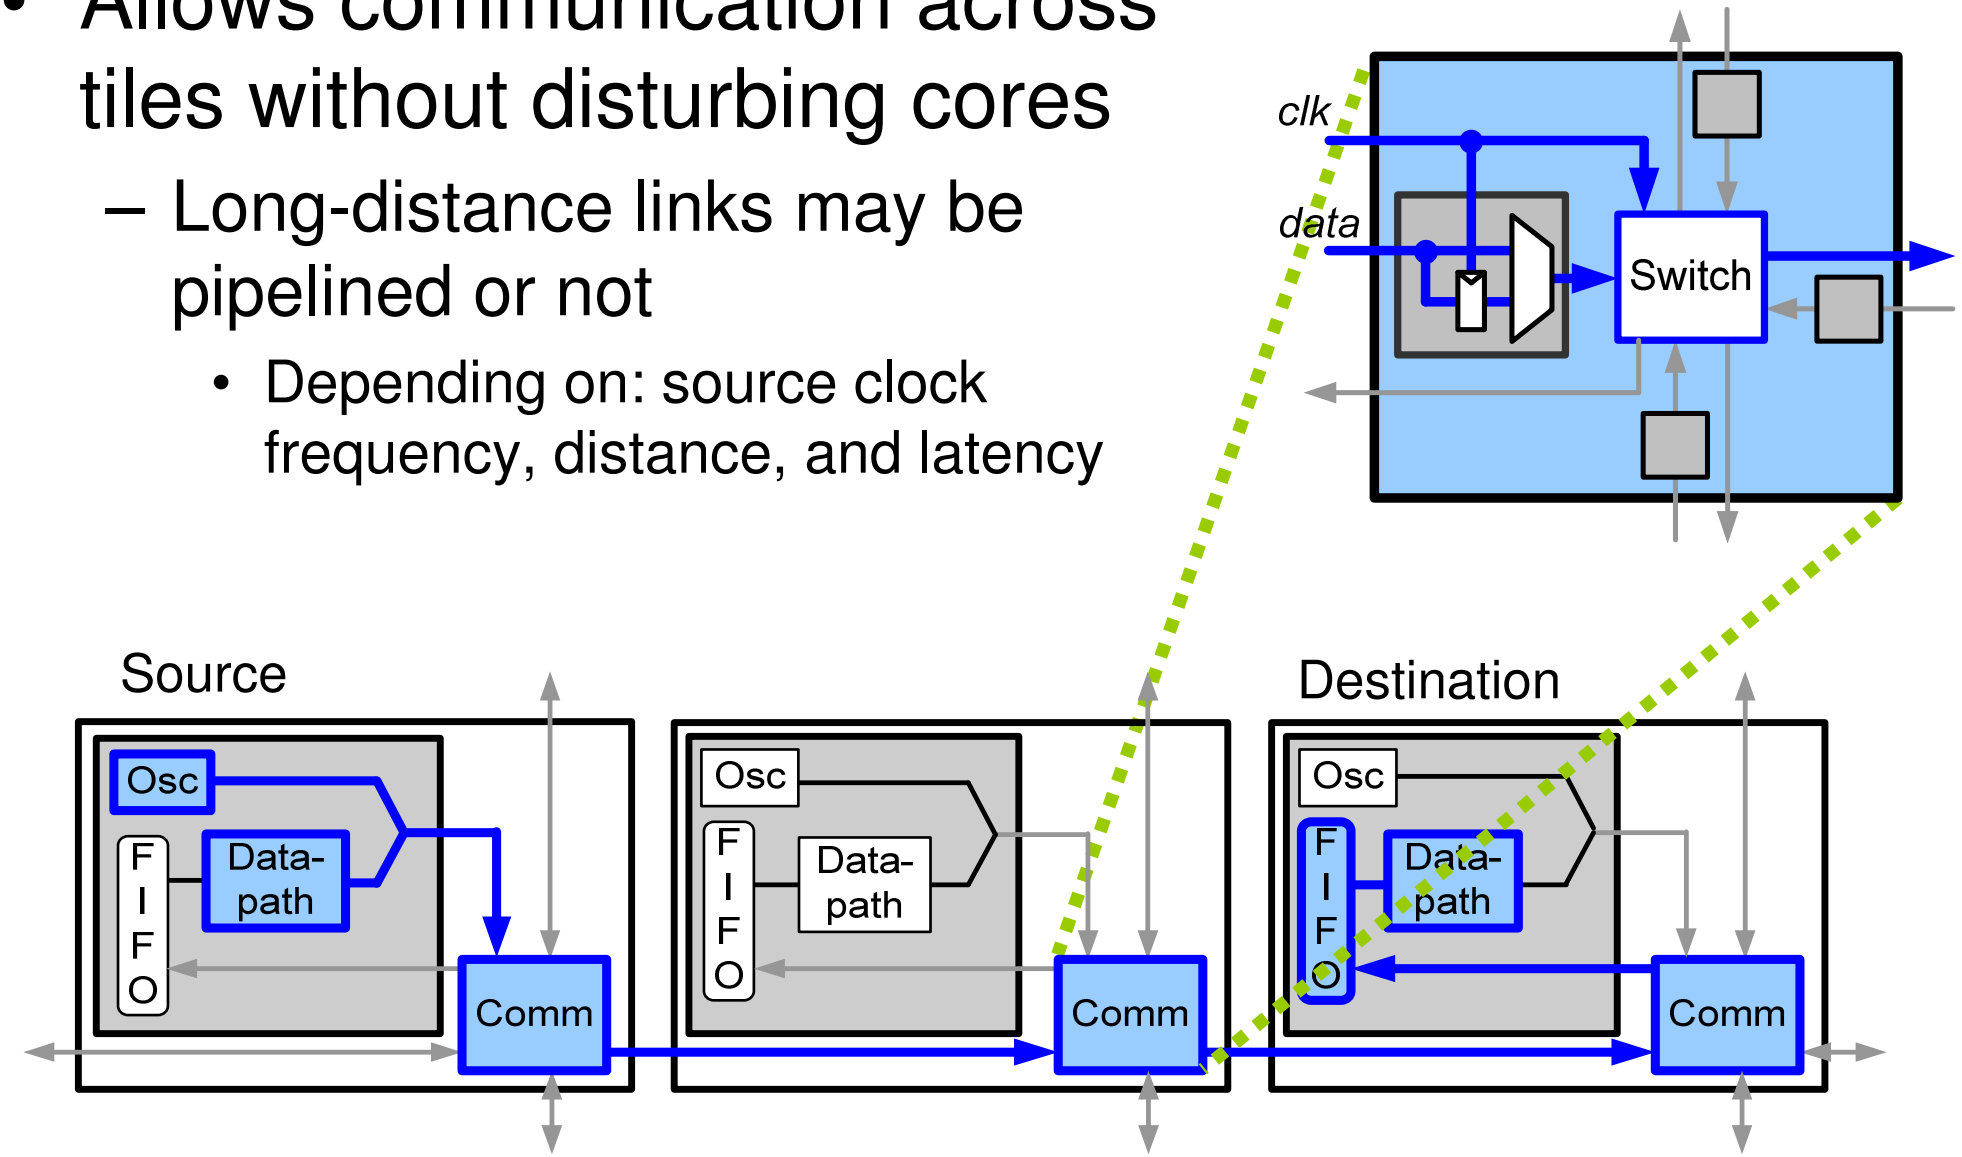

Per-Processor DVFS

- Each processor tile contains a core that operates at:
  - A fully-independent clock frequency
    - Any frequency below maximum
    - Halts, restarts, and changes arbitrarily
  - Dynamically-changeable supply voltage
    - *VddHigh* or *VddLow*
    - Disconnected for leakage reduction
- *VddAlwaysOn* powers DVFS and inter-processor communication

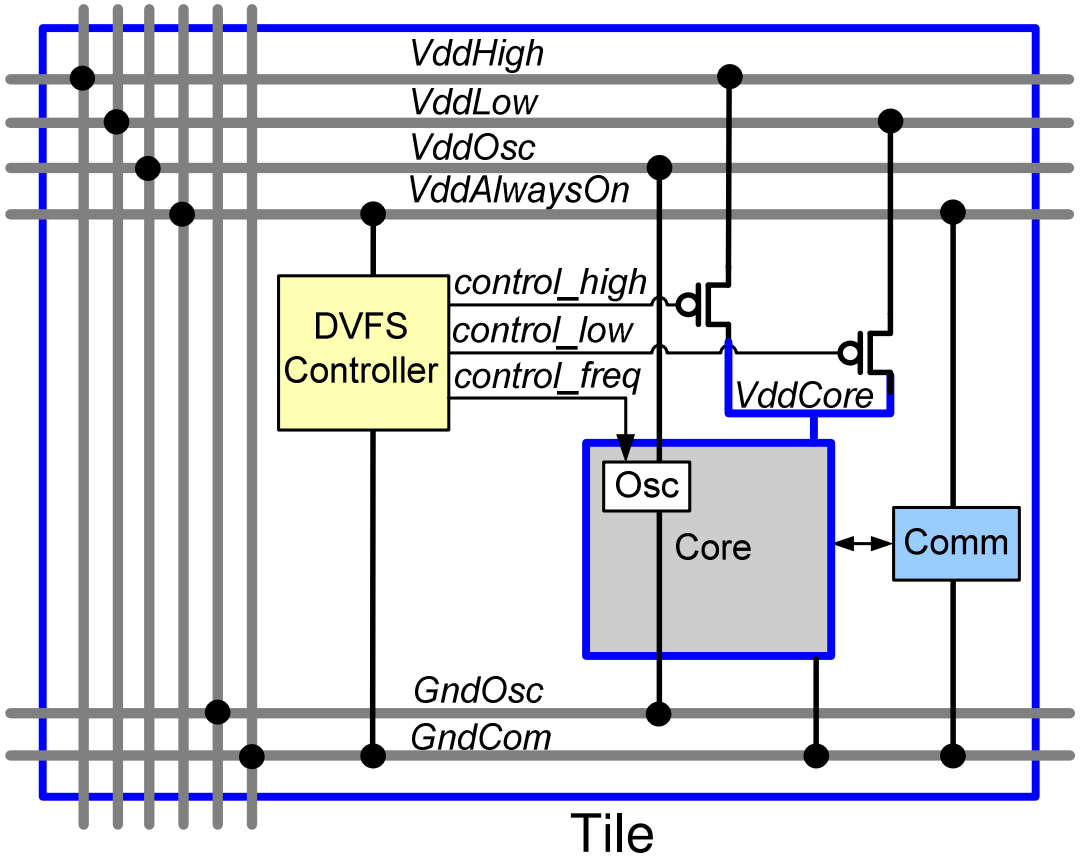

DVFS Controller

- Voltage and frequency is set by:
  - Static configuration
  - Software
  - Hardware (controller)
    - FIFO "fullness"
    - Processor "stalling frequency"

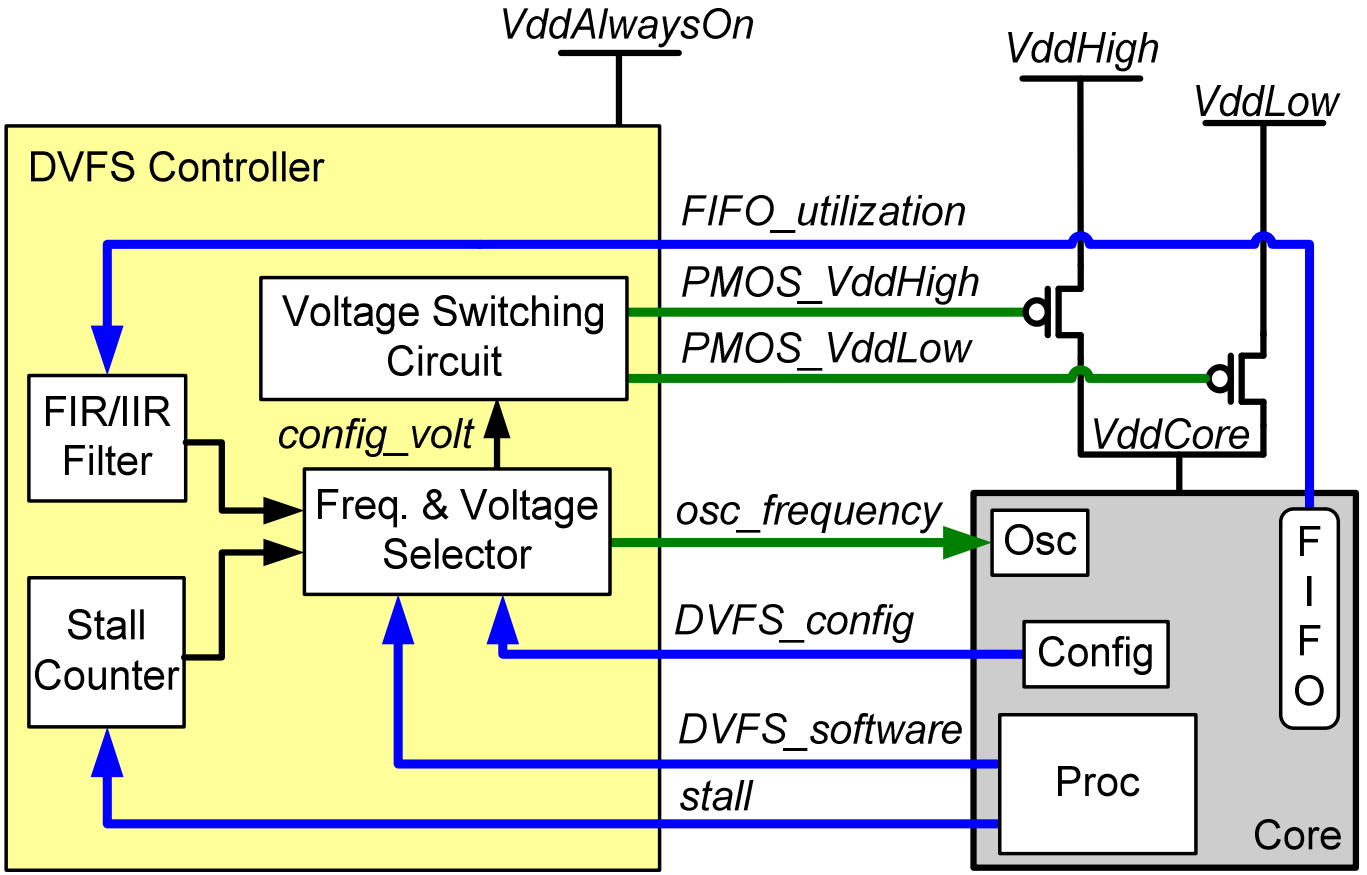

Tile Layout and Power Grids

- 48 power gates surround core
- Metal 6 and 7 are devoted to power distribution—global and local
  - 5 *Vdds*: 79% utilization
  - 2 *Gnds*: 21% utilization
- Vdd/Gnd metal 6 and 7 usage:

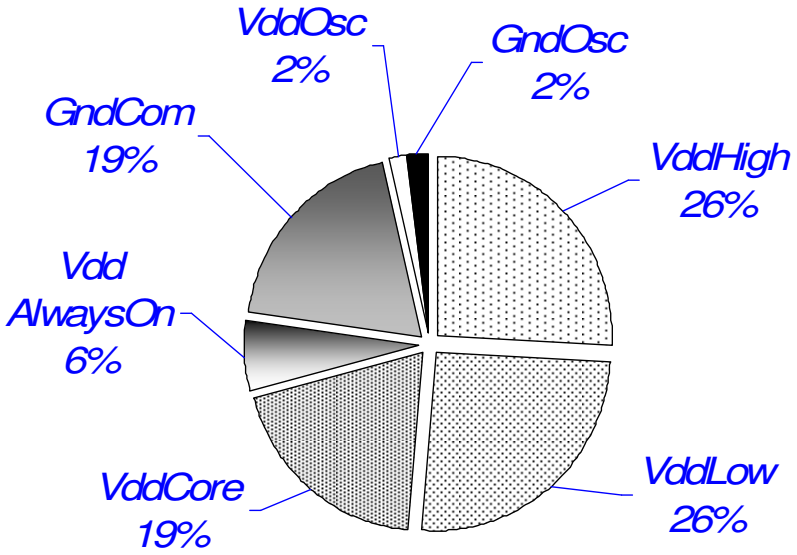

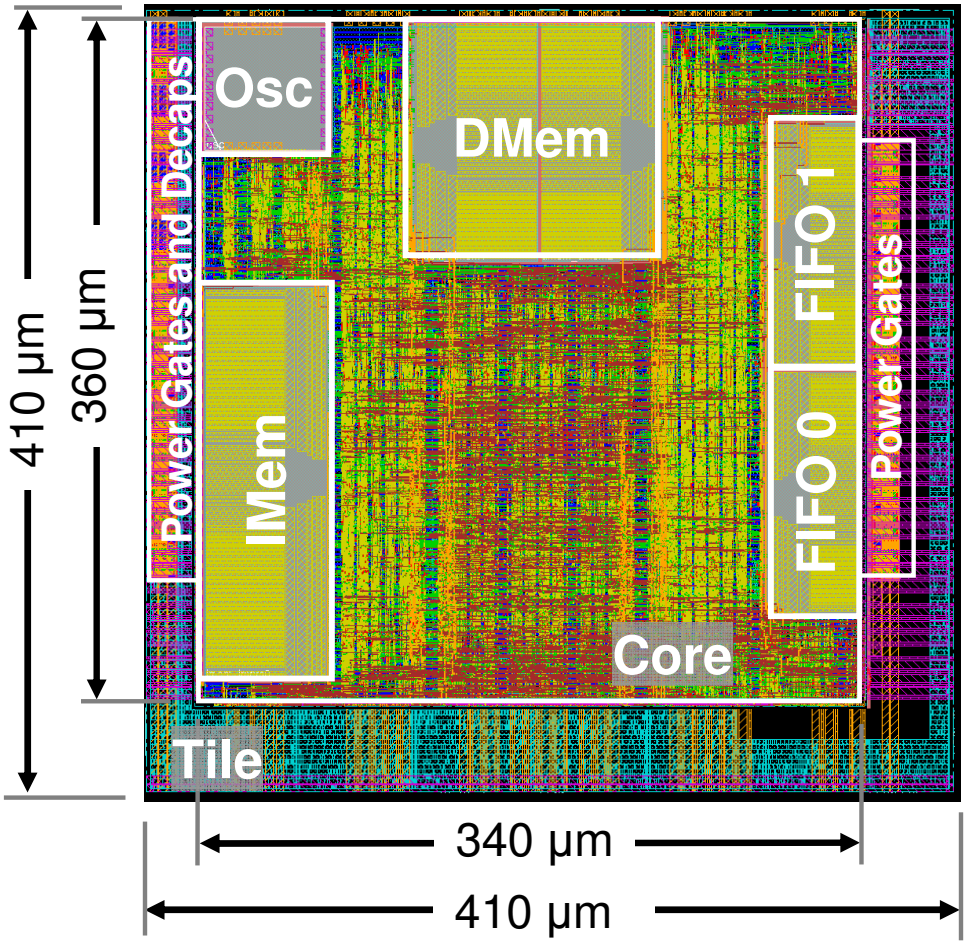

Supply Voltage Switching

- The switching speed and profile “shapes” supply currents while switching to tradeoff switching time versus power grid noise
- Processor cores normally halt during a switch

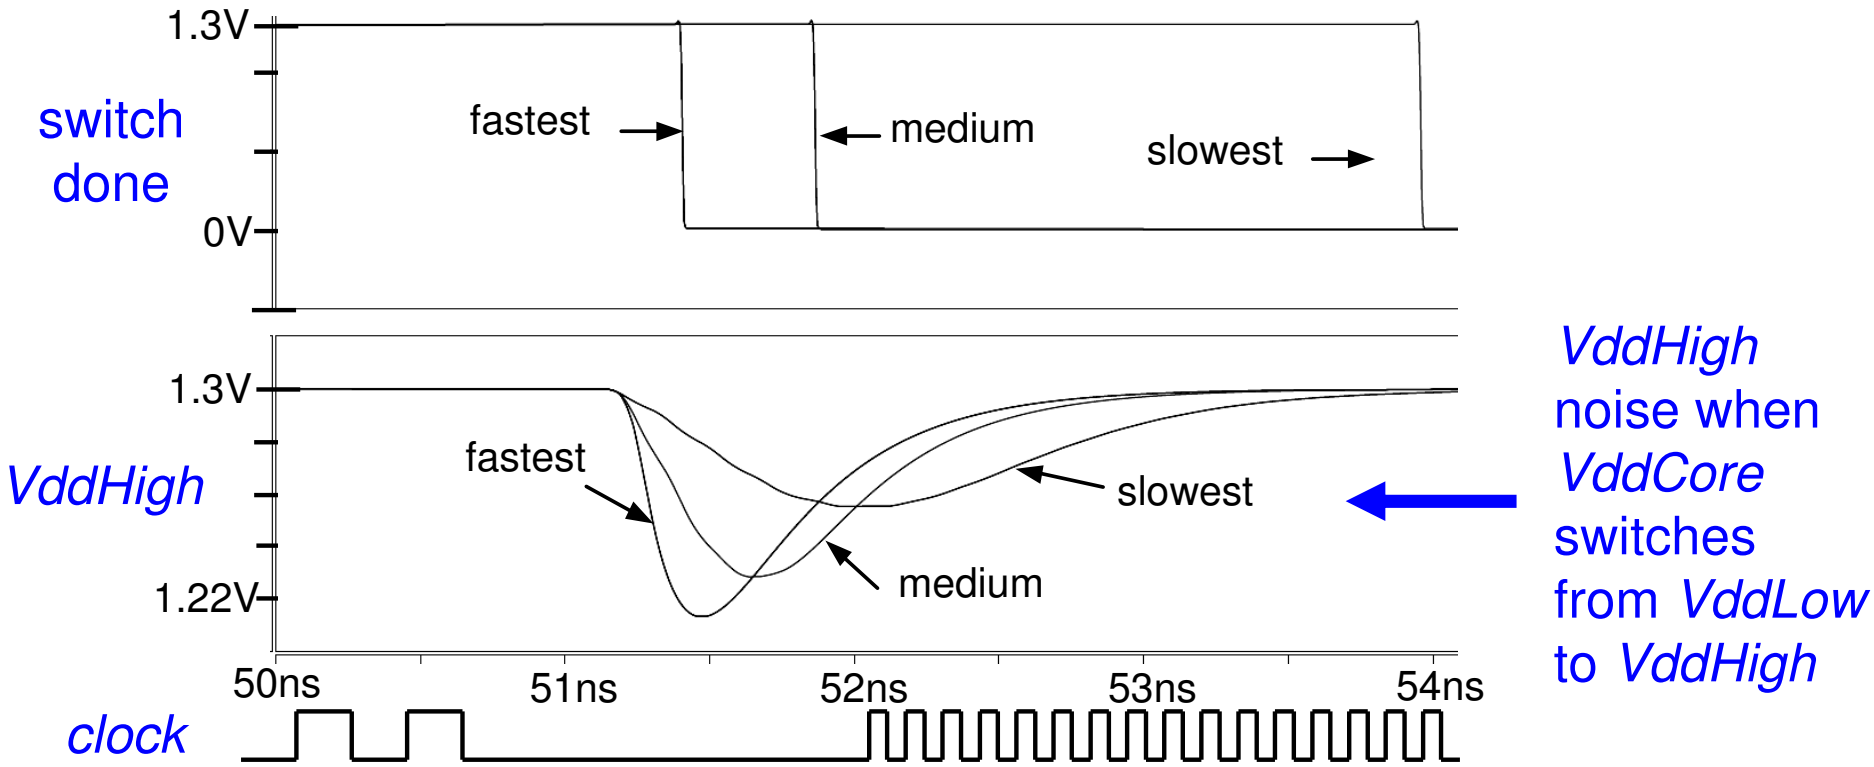

Supply Voltage Switching

- Slow switching results in negligible power grid noise
- Early ***VddCore*** disconnect from ***VddLow*** results in momentary core voltage drop (circled below)

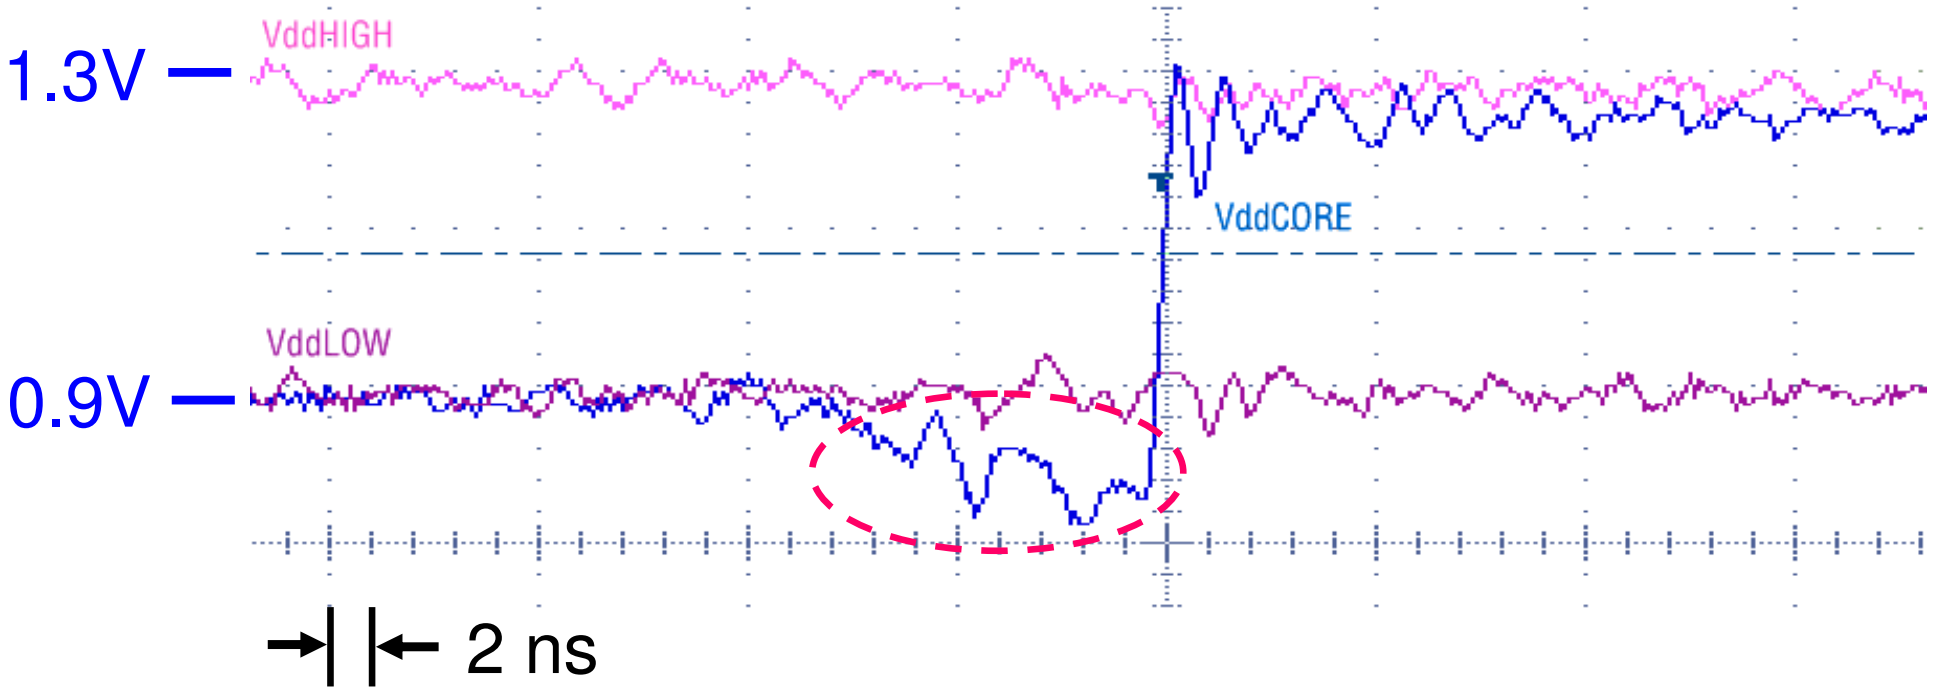

Outline

- Background and the First Generation AsAP
- The Second Generation AsAP
  - Processors and Shared Memories
  - On-chip Communication
  - DVFS
- **Analysis and Summary**

Tile and Core Area Breakdowns

- Communication area approximately 7%
- DVFS area approximately 8%
- Routing complexity results in 27% for gaps and fillers

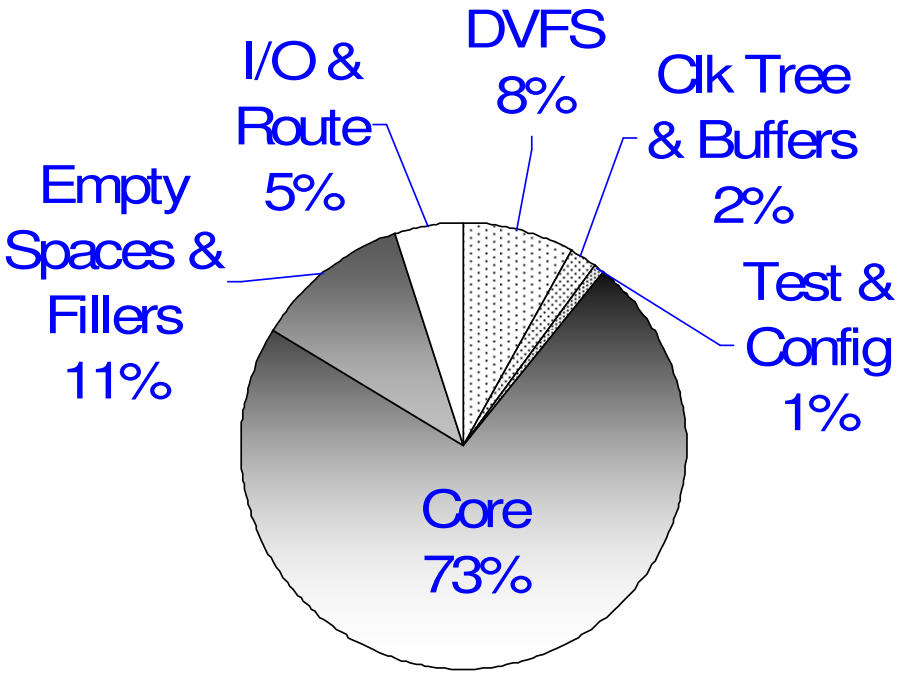

Tile Area

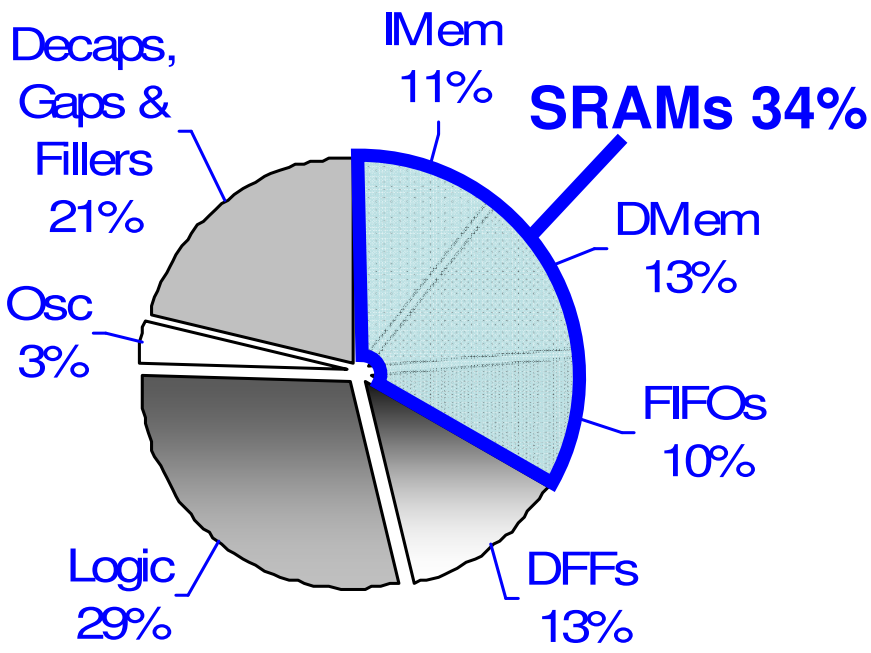

Core Area

Die Micrograph and Key Data

- 65 nm STMicroelectronics low-leakage CMOS

Single Tile	
Transistors	325,000
Area	0.17 mm^2
Max. frequency	1.19 GHz @ 1.3 V
Power (100% active)	59 mW @ 1.19 GHz, 1.3 V
	608 µW @ 66 MHz, 0.675 V
Power (JPEG)	6.3 mW @ 1.095 GHz, 1.3V

55 million transistors, 39.4 mm^2

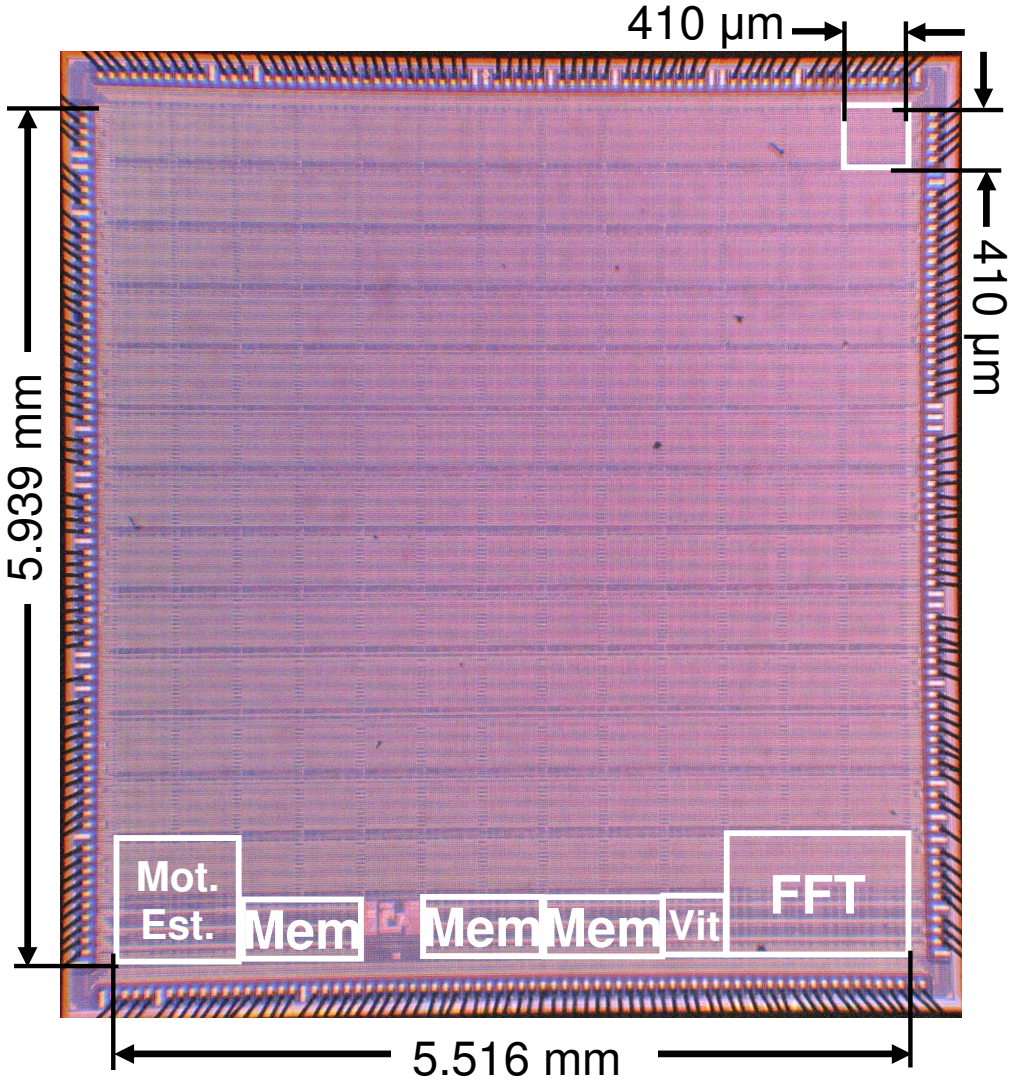

Complete 802.11a Baseband Receiver

- 22 processors
  - 32 processors using only nearest-neighbor connections (46% increase)
- 54 Mbps throughput
  75 mW @ 610 MHz, 1.3 V
- 6x faster than TI C62x,
  2x faster than SODA,
  4x faster than LART
  (all scaled to 65 nm technology, 1.3 V and 610 MHz)

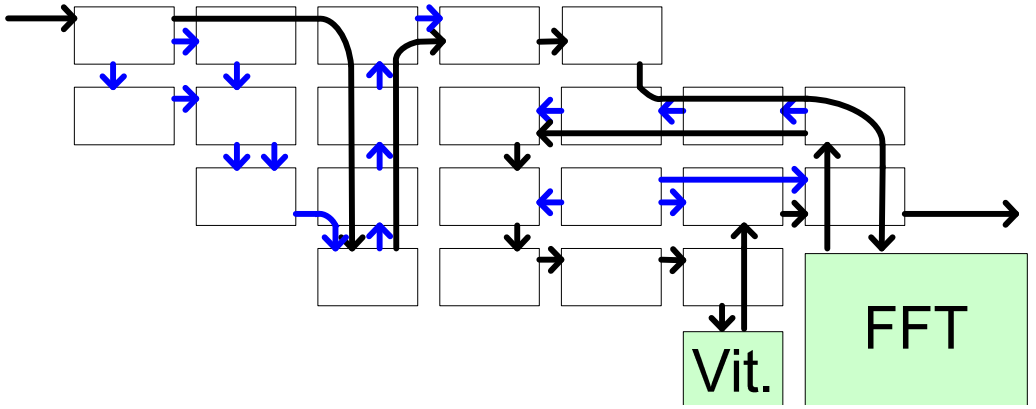

802.11a with long-distance

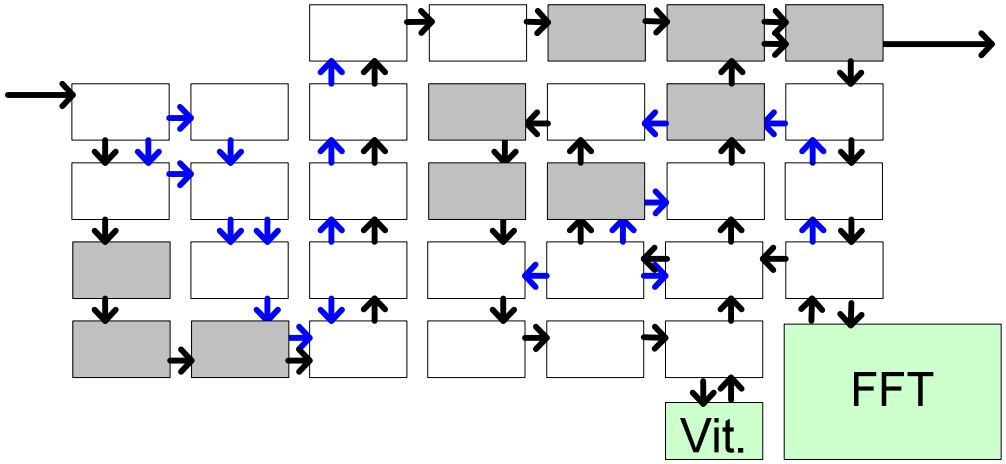

802.11a without long-distance

Summary

- 65 nm low power ST Microelectronics process
- Maintains the basic GALS architecture of AsAP
- 164 homogenous processors
  - 1.2 GHz, 59 mW, 100% active @ 1.3 V
- Three 16 KB shared memories
- Three dedicated-purpose processors
- Long-distance circuit-switched communication increases mapping efficiency without overhead
- DVFS nets a 48% reduction in energy for JPEG with only 8% performance loss

Acknowledgements

- STMicroelectronics
- Intel
- UC Micro
- NSF Grant 430090 and CAREER award 546907
- SRC GRC
- Intellasys
- SEM
- J.-P. Schoellkopf
- K. Torki and S. Dumont
- R. Krishnamurthy and M. Anders