

Ultraperformance Nanophotonic Intrachip Communications: UNIC



Jagdeep Shah
DARPA/MTO

Frontiers of Extreme Computing 2007
Santa Cruz, CA
October 21-24, 2007



UNIC: Outline



- **Introduction: Vision and Challenges**
 - **Microprocessor Challenges**
 - **Photonics/EPIC**
 - **UNIC**
 - **Summary**






CHALLENGE AND VISION



DEVICE SCALING MICROPROCESSOR SCALING

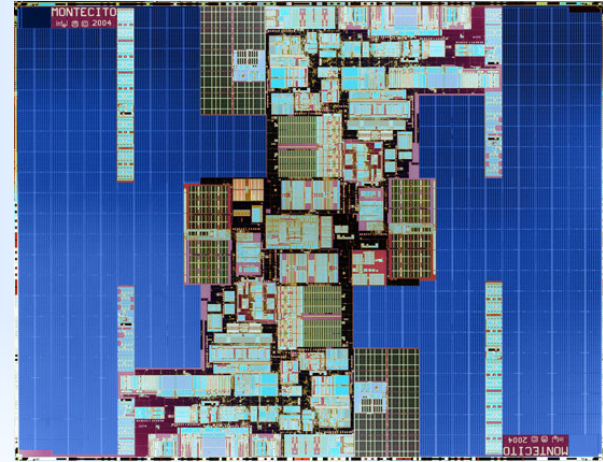
Intel 4004 (1971)

2312 transistors, 11 mm²
~20,000 transistors/cm²



Intel Itanium 2 (2006)

1.72 **billion** transistors, 596 mm²
~300,000,000 transistors/cm²




15,000 X Increase in Device Density

The Impact of Moore's Law:

- Device scaling most obvious in microprocessors
- On the path to >> 1 billion devices/ cm²
- 3D layer stacking on all foundry roadmaps

**Microprocessors are becoming
Ultradense Systems**



DARPA/MTO VISA Program






Other Ultradense systems



DARPA MoleApps – Aim: 10^{15} devices/ cm^3

Courtesy HP



**17 nm half-pitch,
 $3.5 \cdot 10^{11}$ / cm^2
demonstrated**



CHALLENGE AND VISION



CHALLENGE

- How will these ultradense functional units communicate
 - With each other?
 - With the external world?

VISION

**DEVELOP A PATHWAY FOR
SUCH COMMUNICATIONS**





UNIC: Ultrapformance Nanophotonic Intrachip Communications



- DARPA is about to launch a program to develop such a pathway:
 - **for communications between ultrahigh functional units on a chip and from the chip to the outside world**
- **Enormous challenges: Focus on a specific challenge of microprocessors**
- **Use microprocessor circa 2017 as a design driver**



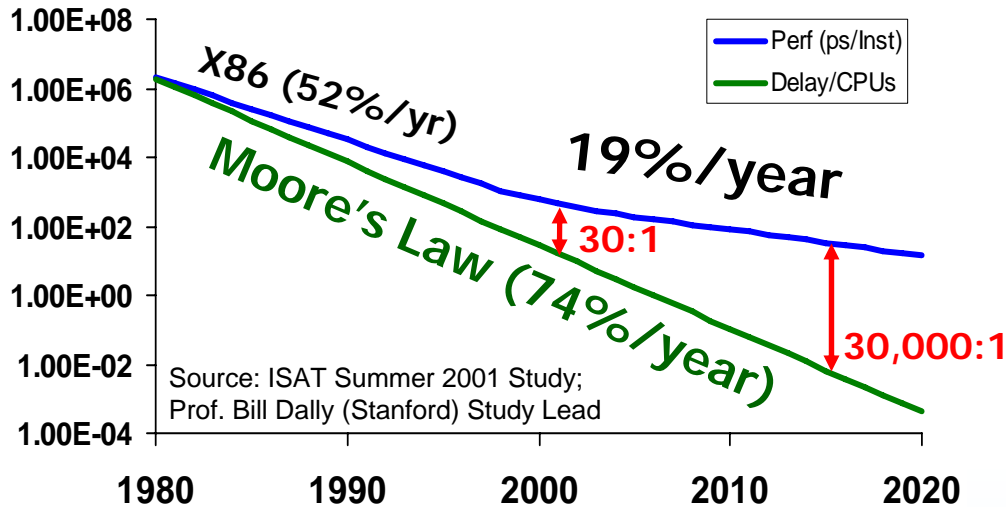


MICROPROCESSOR CHALLENGES





Diminishing Performance Returns



We're already seeing this effect:
Microprocessor performance is not keeping pace with device scaling due to limitations at the circuit level

- Diminishing returns of ILP
- Thermal constraints
- Increasing complexity
- **Growing communications gap**

Supercomputer Comms Gap

- Memory and bisection bandwidth imbalance a growing problem
- Limited by power consumption

Example: Cray XT4, 380 TF
Bytes/Flop ~0.06

System-level balance

- Memory/ Bisection-Bandwidth balanced against peak performance
Metric: Bytes/FLOP ($[B/s]/[FLOPS]$) ~1
- On-chip/ off-chip bandwidth balanced
- Power Consumption
Uniform Distribution: 1/3 Processing, 1/3 Communications, 1/3 Memory






Communications challenges prevent actual system performance from meeting theoretical peak performance



Multicore Processors Today and Tomorrow



Multicore architectures designed to deliver increased performance at the circuit level



1 TFLOPS on-chip reported

Challenges

- Reducing system-level power consumption
- High-bandwidth, low-power access to ultradense devices
- The communications gap:
 - Latency/Bandwidth limits
 - Power dissipation as wires shrink
 - Power-hungry off-chip communications to memory

- Processor/On-chip Memory
 - 100-1000 compute cores in ~6 cm²
 - ~10 TFLOPS peak performance
 - 3D-integration
 - >100 billion active devices!
- Required Communications Network
 - ~80 Tb/s on-chip bandwidth
 - ~80 Tb/s off-chip memory BW
- Total System Power ~200 W

Global Information
Grid: 5 Tb/s

Electronics cannot meet communication requirements within power budget. Actual performance cannot meet theoretical peak performance.

A new communications strategy is required



PHOTONICS






Optical Communications




The high-bandwidth, low-latency properties of optical interconnects are continually meeting our increasing communication needs at smaller scales.


Increased Bandwidth and Device Density



Board to Board



Chip-to-Chip



On Chip

Savage, *IEEE Spectrum*, 39 (8), 2002.



Is photonic signaling a potential solution to the chip-scale communications problem?



PHOTONICS



BENEFITS

- **High Bandwidth/Capacity**
 - **Wavelength Division Multiplexing (WDM)**
 - **Time Division Multiplexing (TDM)**
 - **Space Division Multiplexing (SDM)**
- **No need for power hungry repeaters**
- **Power Independent Of Distance (~ 10 cm)**
- **Seamless I/O**

CHALLENGES

- **Compatibility with Silicon**
- **Current devices too large: can they be scaled down in size?**
- **Current devices are power hungry: can the power be reduced?**
- **Spectral bandwidth vs. thermal stability**

Devices required to meet these challenges will be very different from today's devices performing the same functions

EPIC

ELECTRONIC PHOTONIC INTEGRATED CIRCUITS



**Success of EPIC gives
confidence
that this is possible**



- Legacy microphotonic devices were discrete components using a variety of materials
- Moore's law has opened the way for integrated photonics in Si
- High index-contrast of Si/SiO₂ and smooth features below 90nm node enable nanophotonic devices
- DARPA EPIC Program has demonstrated high performance, monolithically integrated photonic and electronic devices using a standard CMOS foundry
- Application-specific EPIC chips demonstrated





UNIC

**PHOTONICALLY-ENABLED
MICROPROCESSOR**





UNIC APPROACH



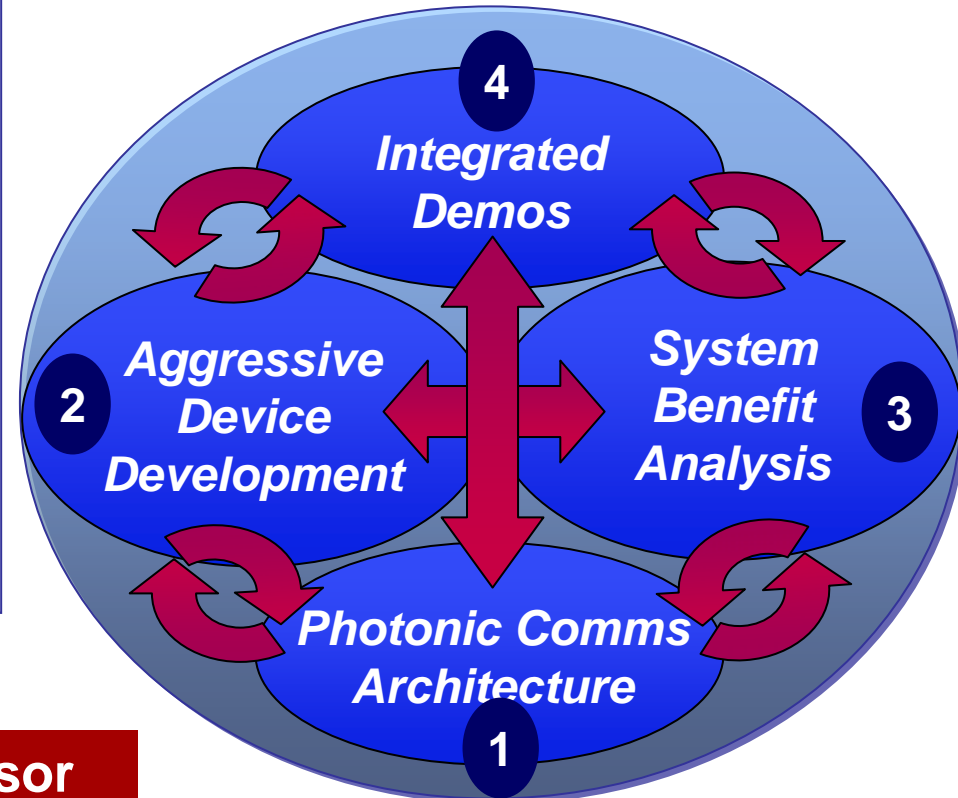
- **Select a Design Driver (with ultradense functional units: microprocessor circa 2017)**
- **Design a photonic communication network within the constraints imposed by the design driver**
- **Quantify the system benefits of the approach**
- **Quantify device requirements to enable such communications**
- **Demonstrate the required device performance**
- **Demonstrate on-chip functional communication links with all essential components working in unison (sufficiently aggressive to convince the skeptics)**
- **Demonstrate multiple high performance microprocessors communicating via on-chip optical links**



Demonstrate to the microprocessor community (and others) that photonic intra-chip – and seamless off-chip – communication is a credible technology that will allow actual system performance to scale to a level not possible with electronic communications

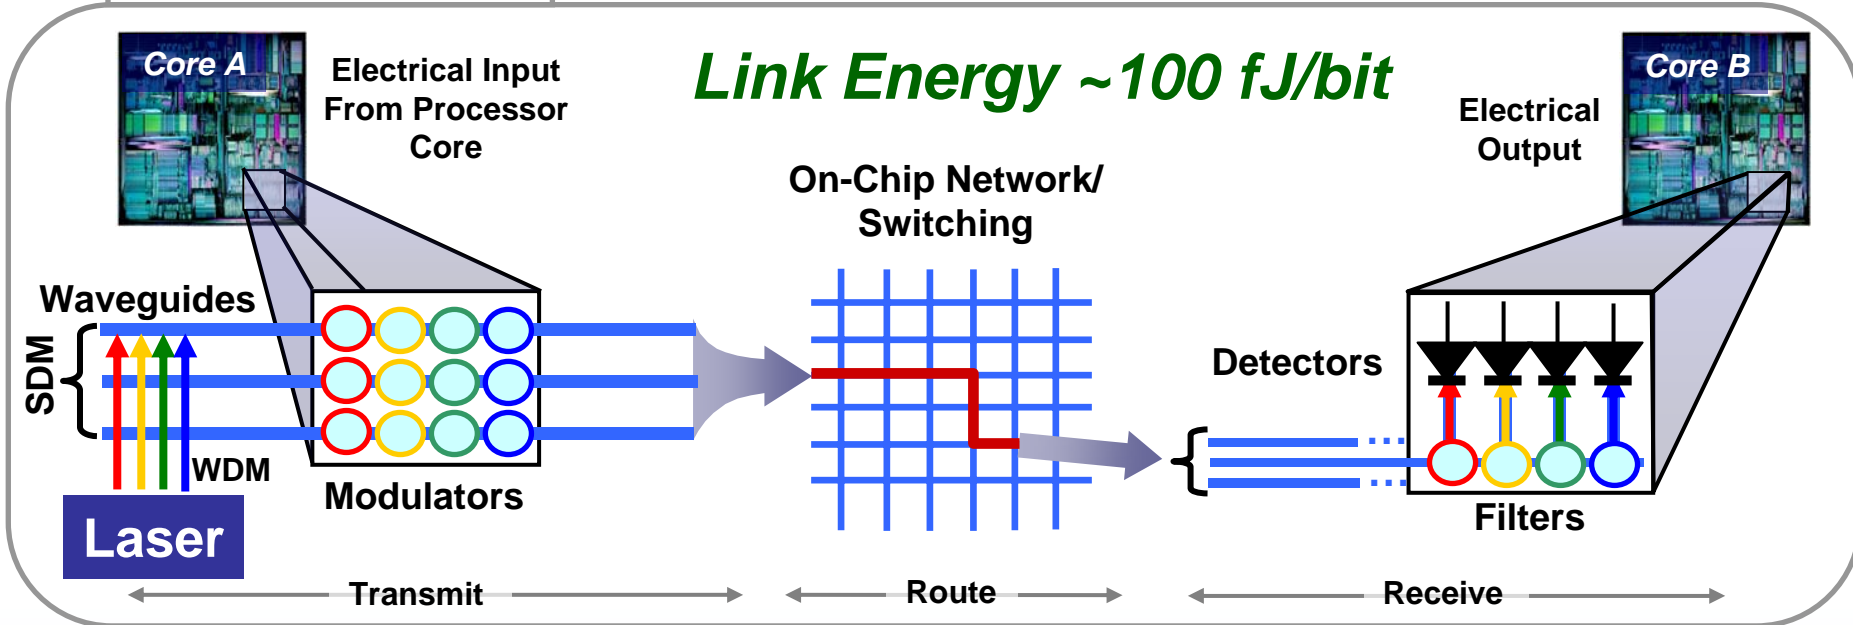
Comprehensive, team-based efforts encompassing

1. Photonic Communication Architecture (on/off chip)
2. Device demonstrations (compatible with CMOS fabs) far beyond EPIC
3. System-level performance-benefit analysis
4. Full-link demonstrations of all critical technologies working together; application emulation
5. Microprocessors communicating via on-chip optical links



No effort on processor design, architecture...

Optical Link Example



How can 80 Tb/s be achieved?


- **Transmit** – Spatial Division Multiplexing (SDM) and Wave Division Multiplexing (WDM) of optical signals with multiple, low-insertion-loss, high-BW modulators
- **Route** – Spectrally filtered passive networks or active, arbitrated spatial switch networks using low loss, high power handling waveguides. Power consumption due to tuning must be kept at a minimum by using thermally tolerant designs.
- **Receive** – High-BW, high-responsivity, low-power (no TIA) photo detectors

System-level constraints mandate stringent performance requirements

Device Example (w/drivers)		EPIC Demonstrated	UNIC Requirements	Required Improvement
10+ Gb/s Modulator	Area	700 μm^2	28 μm^2	25 X
	Power	330 mW	0.8 mW	825 X
WDM Filter	Area	0.6 mm^2	0.005 mm^2	120 X
	Power	73 mW (tuning)	0 mW ?	???
10+ Gb/s Detector	Area	0.16 mm^2	0.01 mm^2	16X
	Power	36 mW	1 mW	36 X



EPIC



UNIC

UNIC will consist of many EPIC-like circuits with dramatically reduced power consumption and dimensions

Next Generation Intra-chip Communication Devices Require Dramatic Size and Power Reductions

PHOTONIC TECHNOLOGY BENEFITS

- High Bandwidth: Wavelength, Time, and Space Division Multiplexing
- Low Power: No repeaters, buffers, regenerators; independent of dist.
- Seamless I/O: No need for power-hungry off-chip communication
- Enables high performance for low system power

System-level benefits


- Restores B/F system balance to maximize actual performance
- Facilitates architectures which reduce programming complexity (e.g. shared memory)
- Reduces chip and system power
- Enables deployable, chip-scale supercomputers



Real-Time, High-Performance Embedded Processing



Image Processing



SAR Processing



Autonomous Ops

Supercomputers



Cryptanalysis

Frontiers of Extreme Computing 2007: Santa Cruz, CA



SUMMARY



- **Device scaling is producing ultradense electrical microsystems**
- **Microprocessors are becoming ultradense supercomputers on-chip**
- **High-performance computer systems are becoming unbalanced due to communications bottleneck**
- **Photonic communications may provide a novel solution to high BW on- and off-chip communications challenges**
- **UNIC addresses this problem head-on**

