JUNIPER NETWORKS

ASICS: THE HEART OF MODERN ROUTERS

Chang-Hong Wu
Distinguished Engineer, Juniper Networks

THE INTERNET EXPLOSION

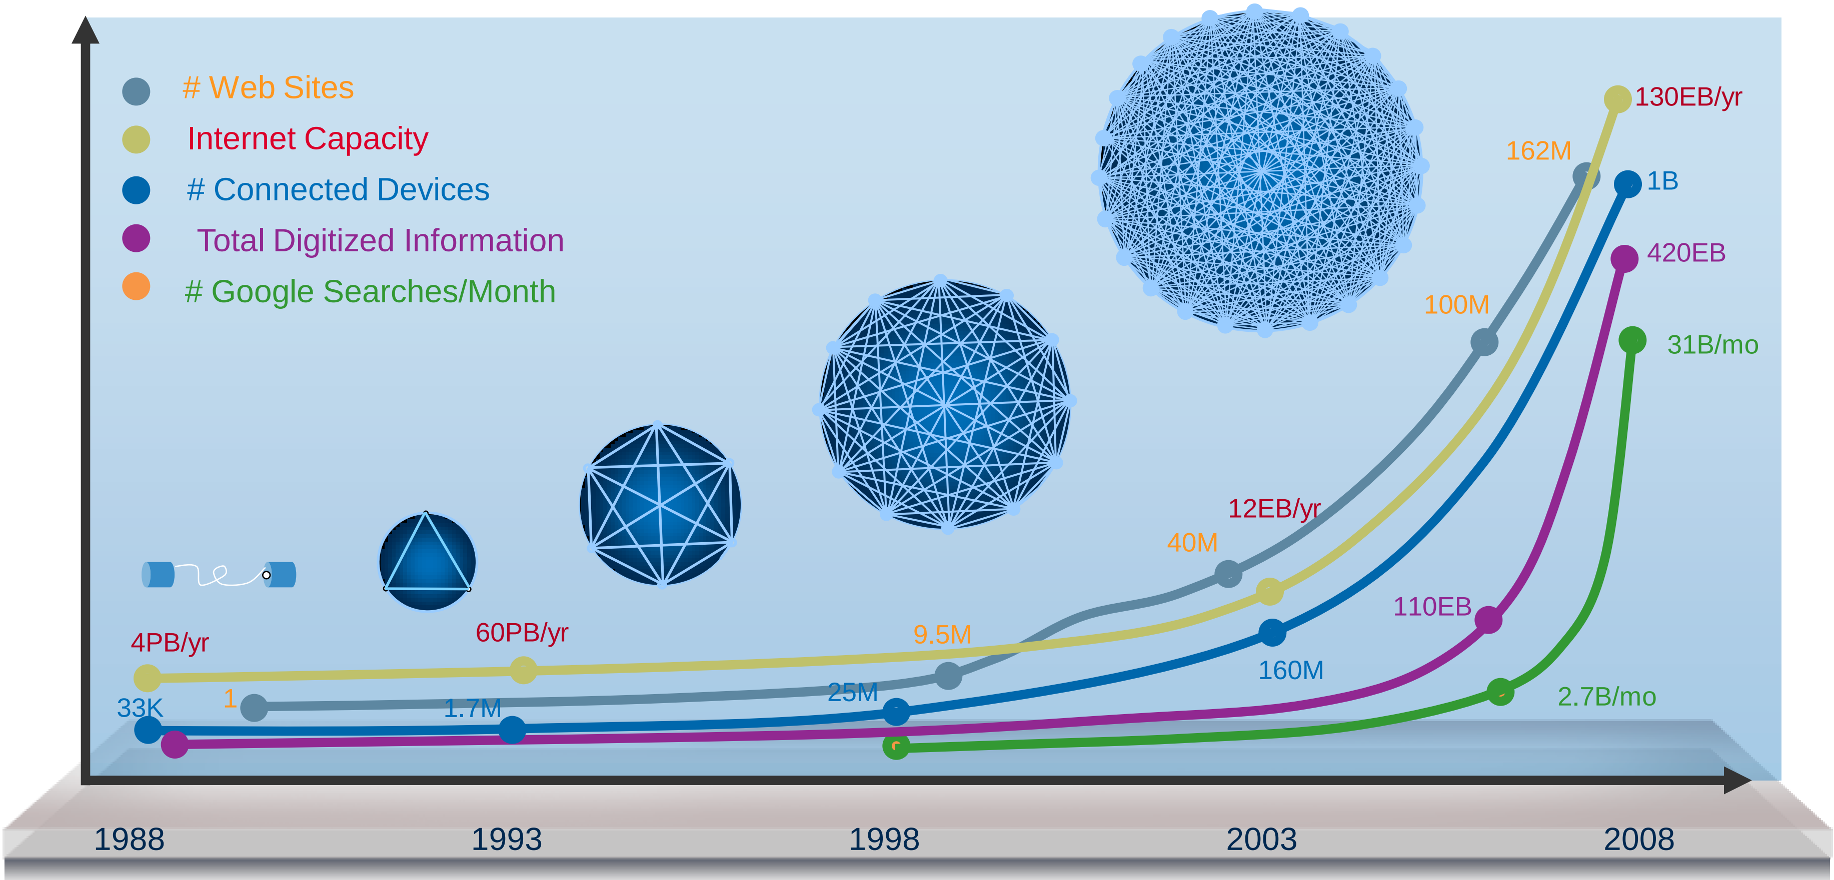

Exponential growth, no matter how you measure it!

The clearest indication of value delivered to end-users

DRIVING FORCE BEHIND EXPONENTIAL GROWTH

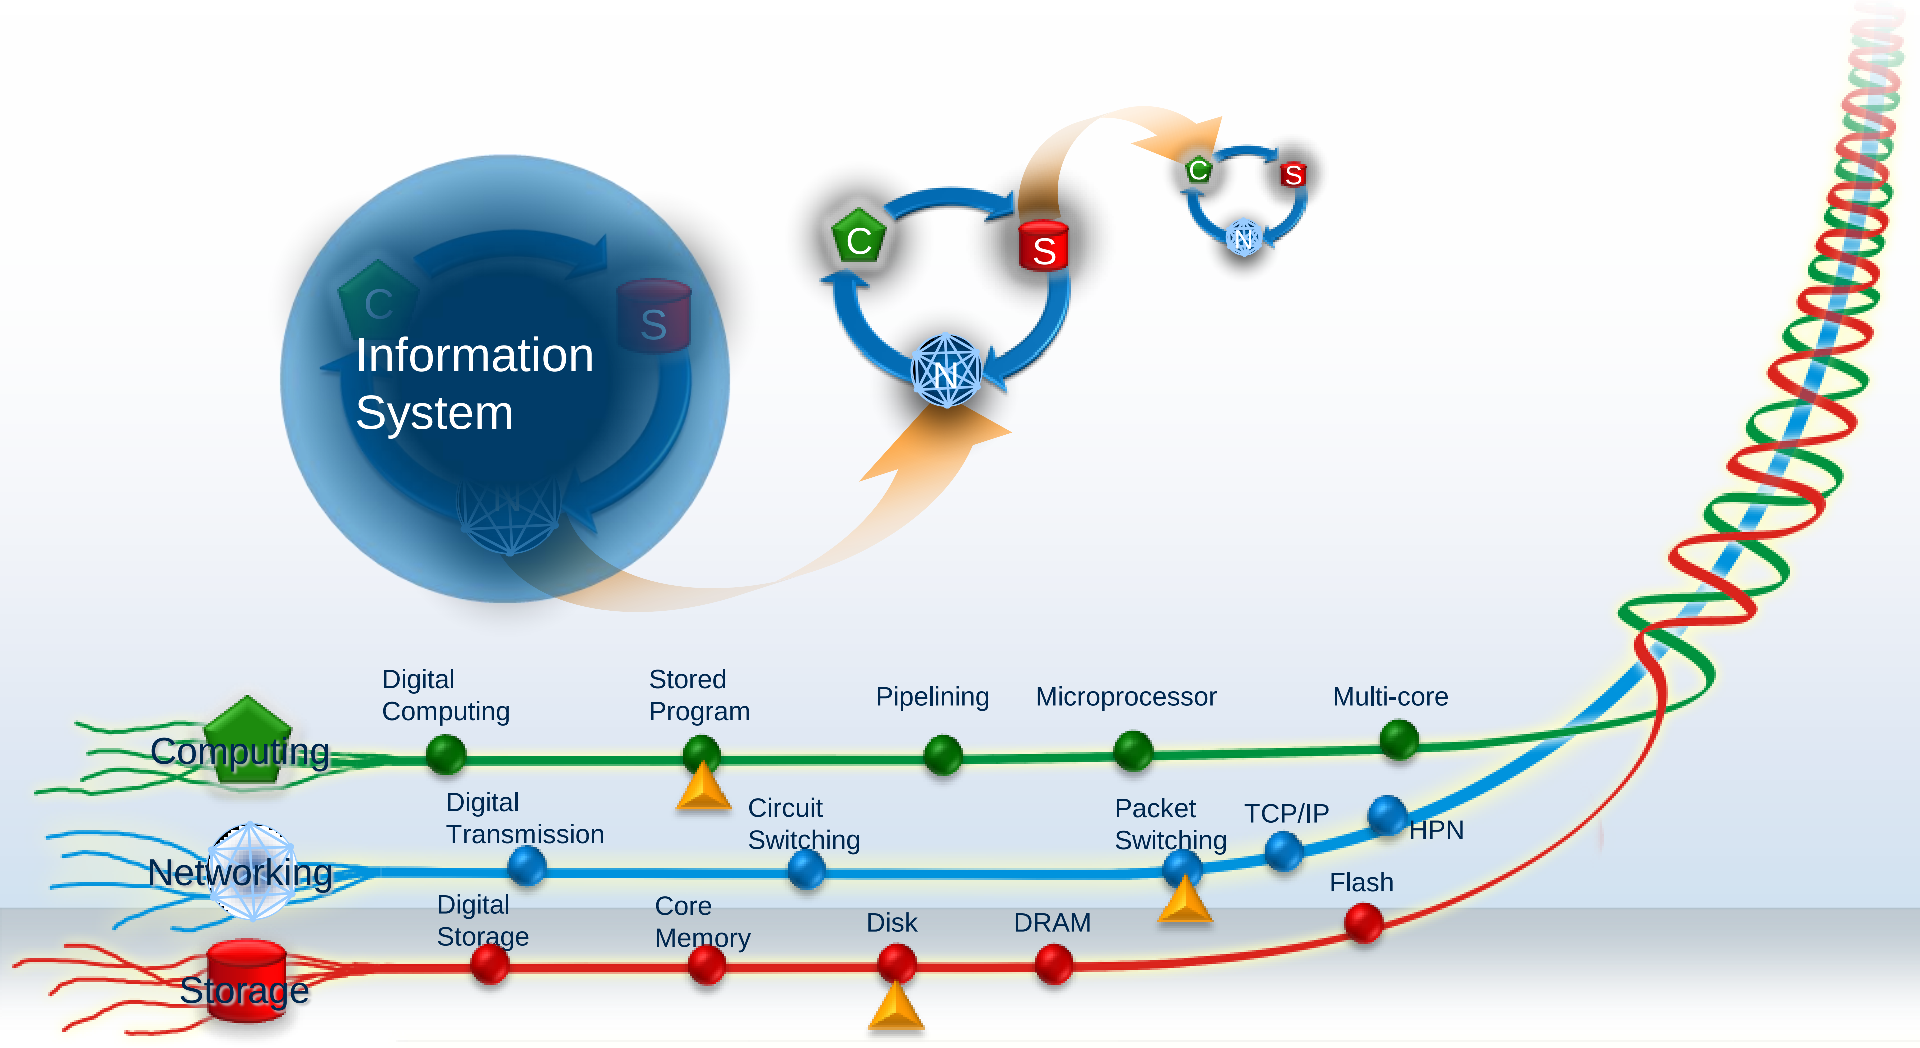

COMPUTER PERFORMANCE: 1988-2008

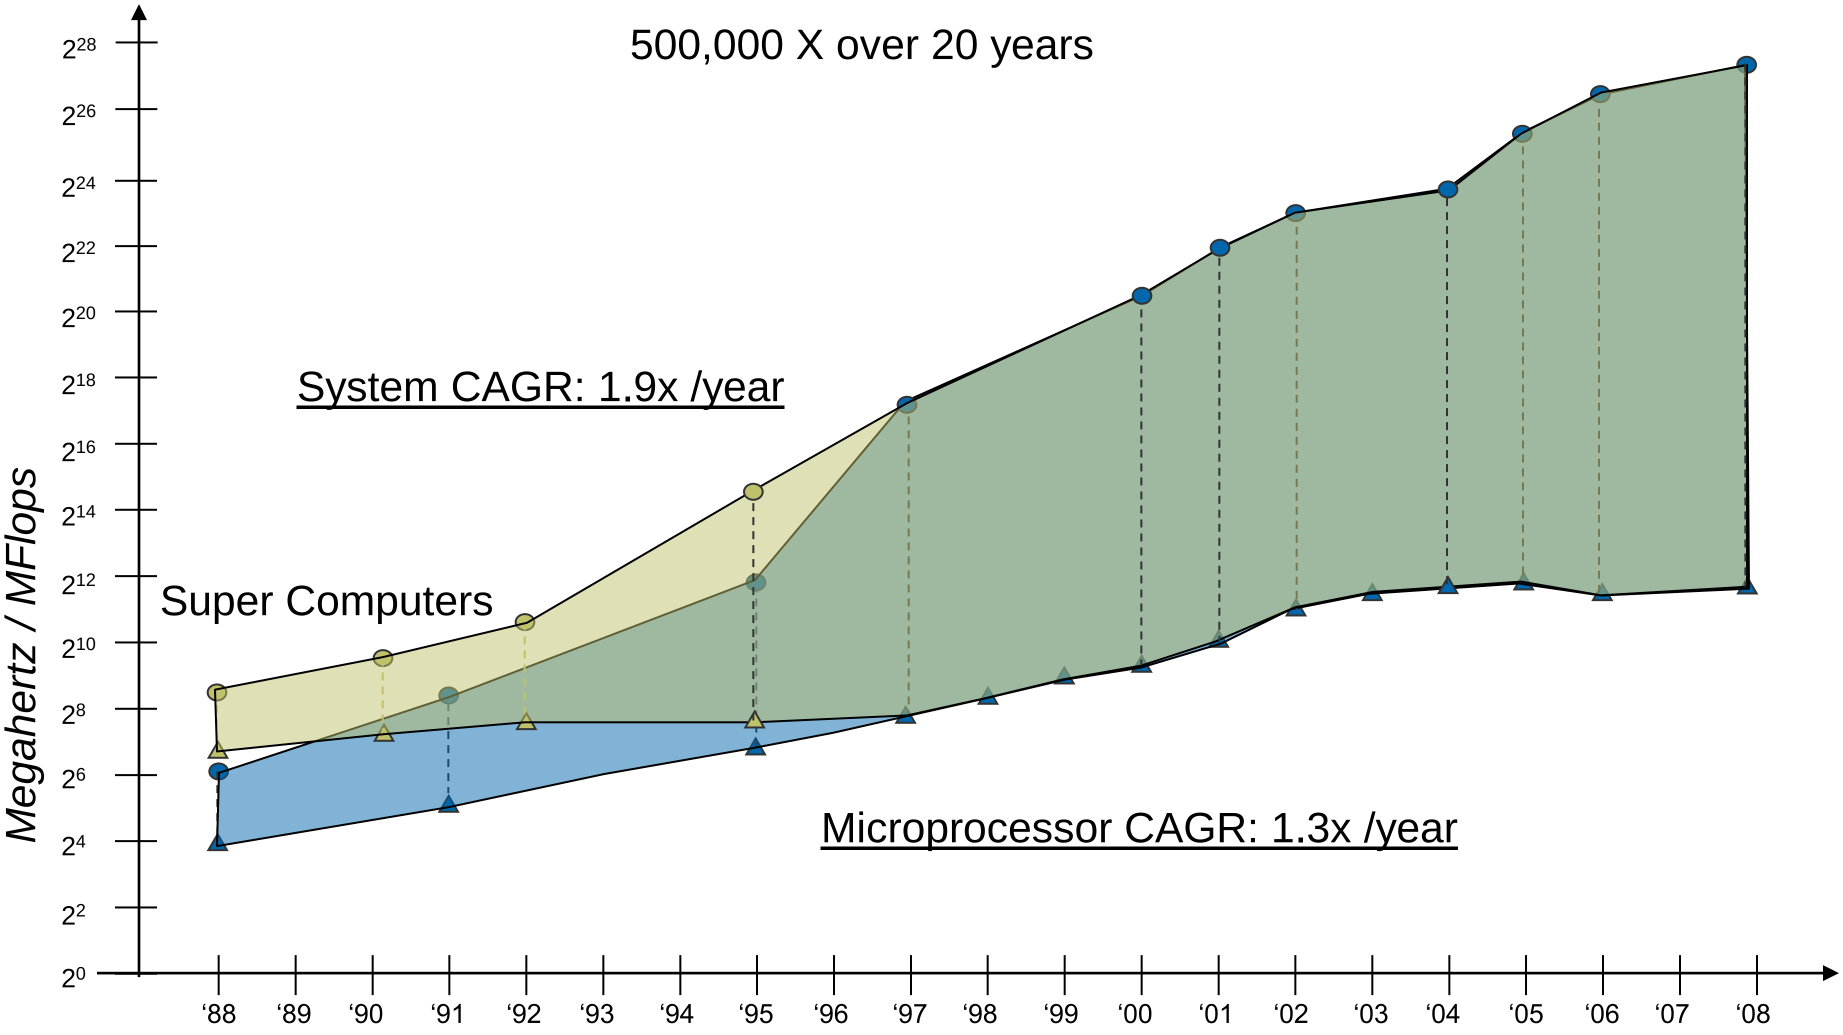

ROUTER PERFORMANCE 1988 – 2008

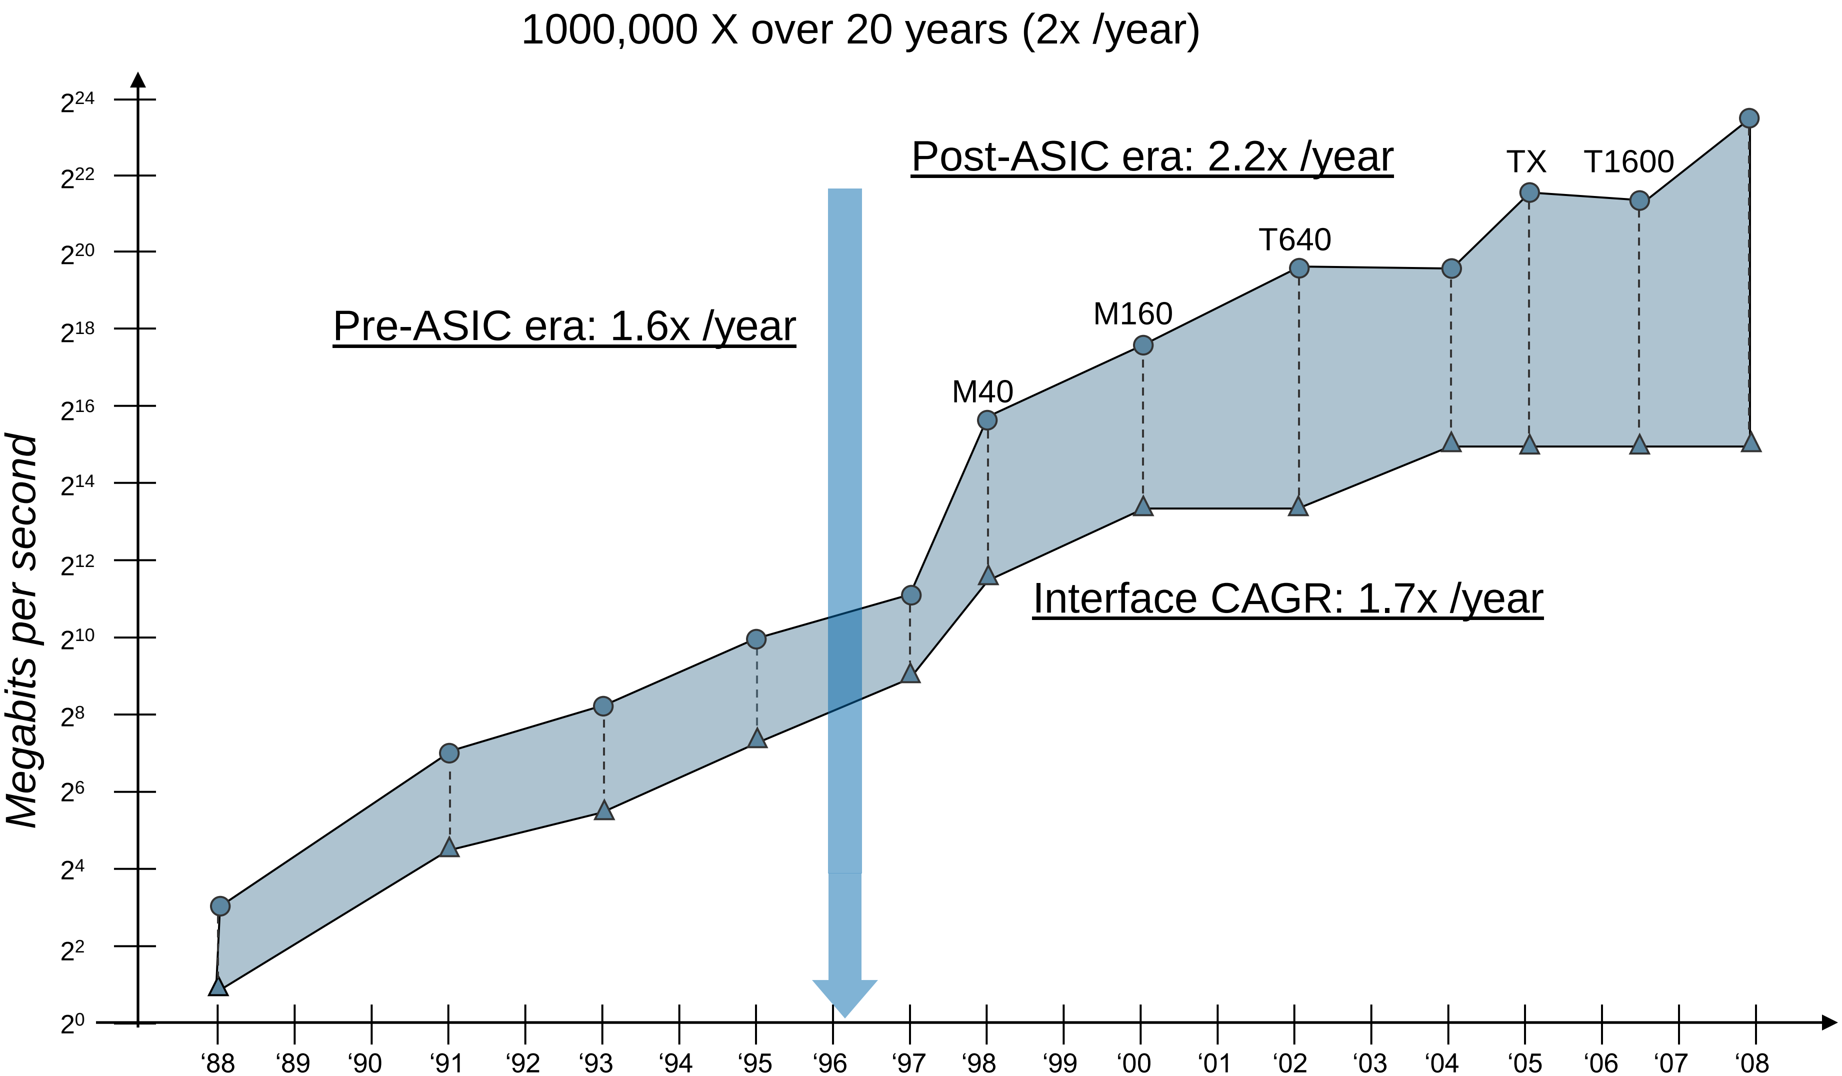

SILICON
THE FOUNDATION OF PERFORMANCE

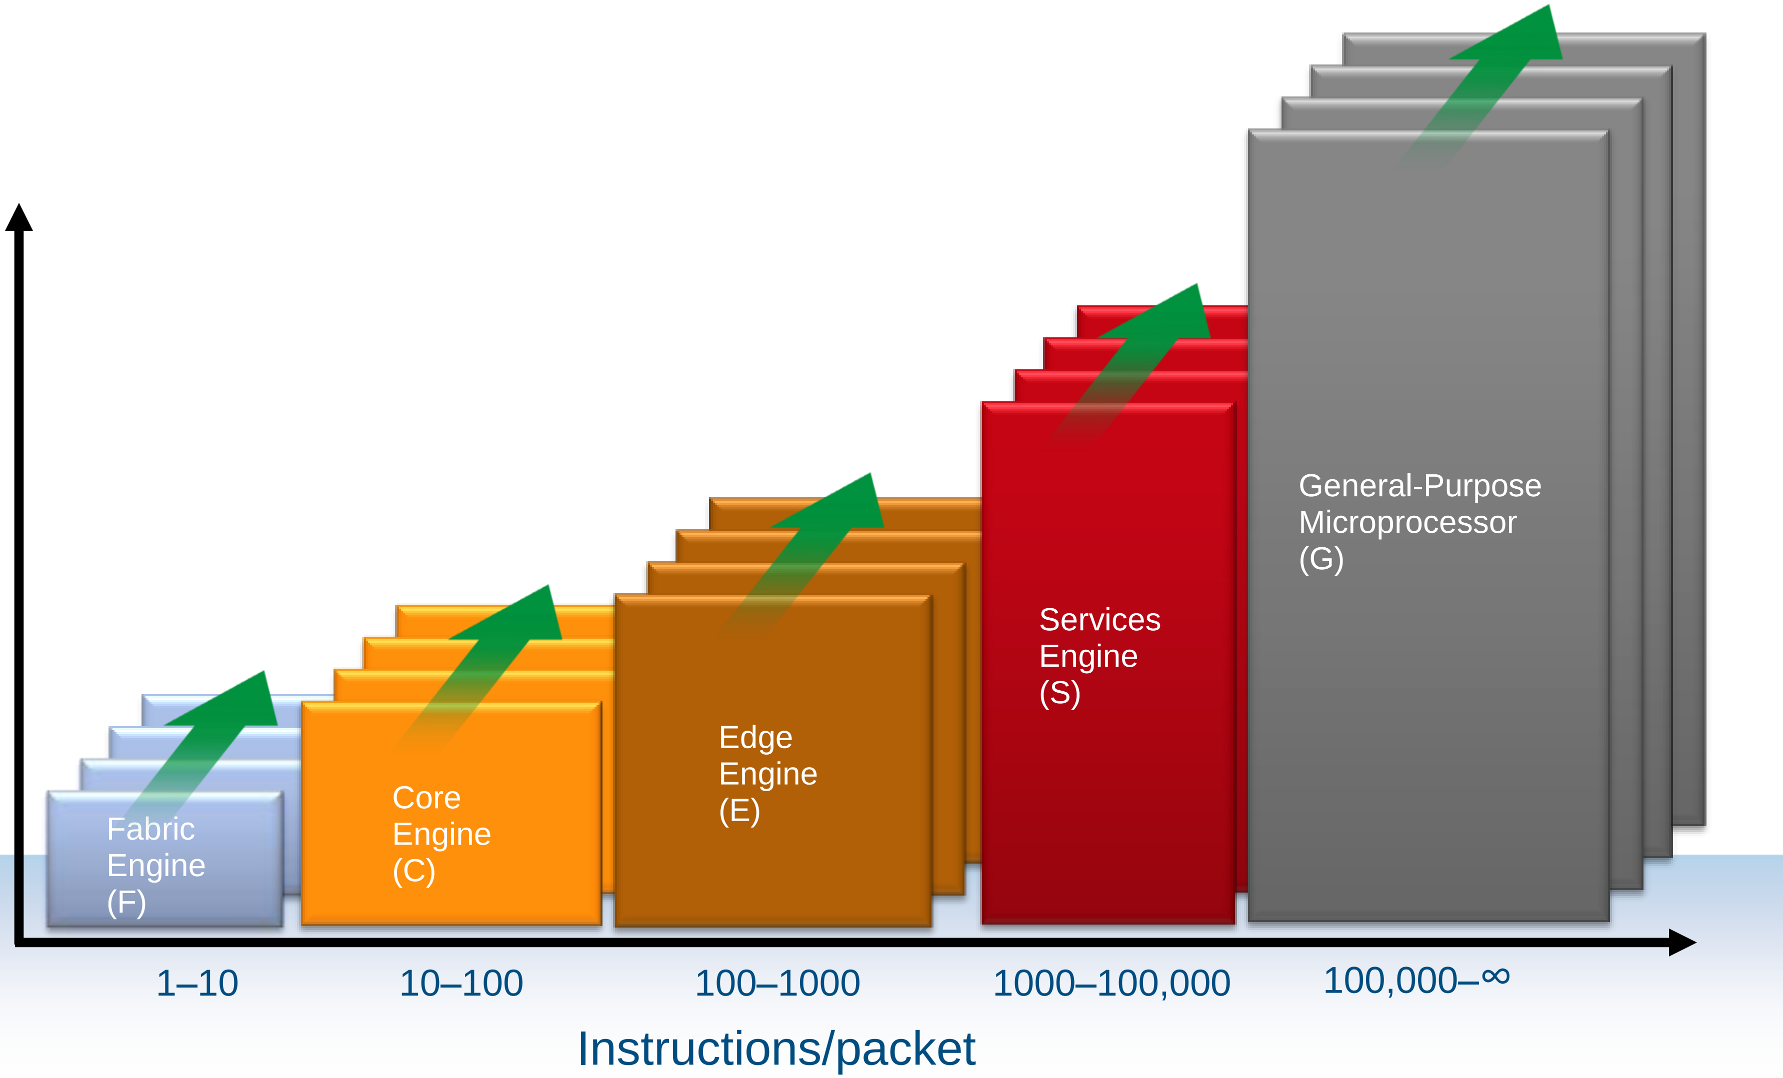

COMPARISON OF SILICON TECHNOLOGIES

Technology	Advantages	Disadvantages	Use Cases
General Purpose CPUs	Very flexible	Poor performance, density, and power	Flexibility is more important than performance
Field Programmable Gate Arrays	Smaller up-front development cost; Field upgrades	Lower performance, density, and power; High per part price	Volume is low; Changes are expected
Off-the-shelf Network Processors	Flexible. Jump straight into software design	Can fall short of performance, power, and functionality targets	Differentiation is not important
ASICs	Tailor to your specification	High upfront cost; Long development cycle	High performance; Low production cost

SYSTEM ARCHITECTURE

Market requirements

- Performance, density, feature, cost targets

Software/hardware interactions

- Functional partitioning

Silicon process technology evaluation

- Cost/performance tradeoffs

Memory choices

- Stores configuration, FIB tables, etc.
- Temporary working buffers

Chip partitioning

- IO and logic ratio, die size, interface simplicity

ASIC PROCESS TECHNOLOGY

- Greater density allows more features/functionality for the same price
- Moore’s Law: Transistor density doubles every 18 months
  - Holding up remarkably well. But how much longer?
- While density is increasing, performance is starting to level off
- The decrease in operating voltage, hence dynamic power, also slowed
- Static power is becoming an issue
- NRE costs associated with newer processes increasing dramatically
- Architectural innovations are needed to continue to provide value to customers

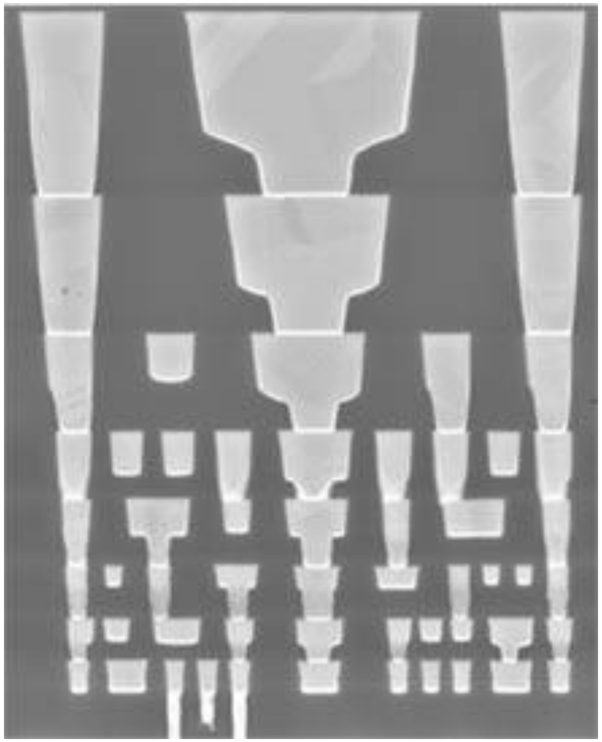

NETWORKING ASICS AND MEMORIES

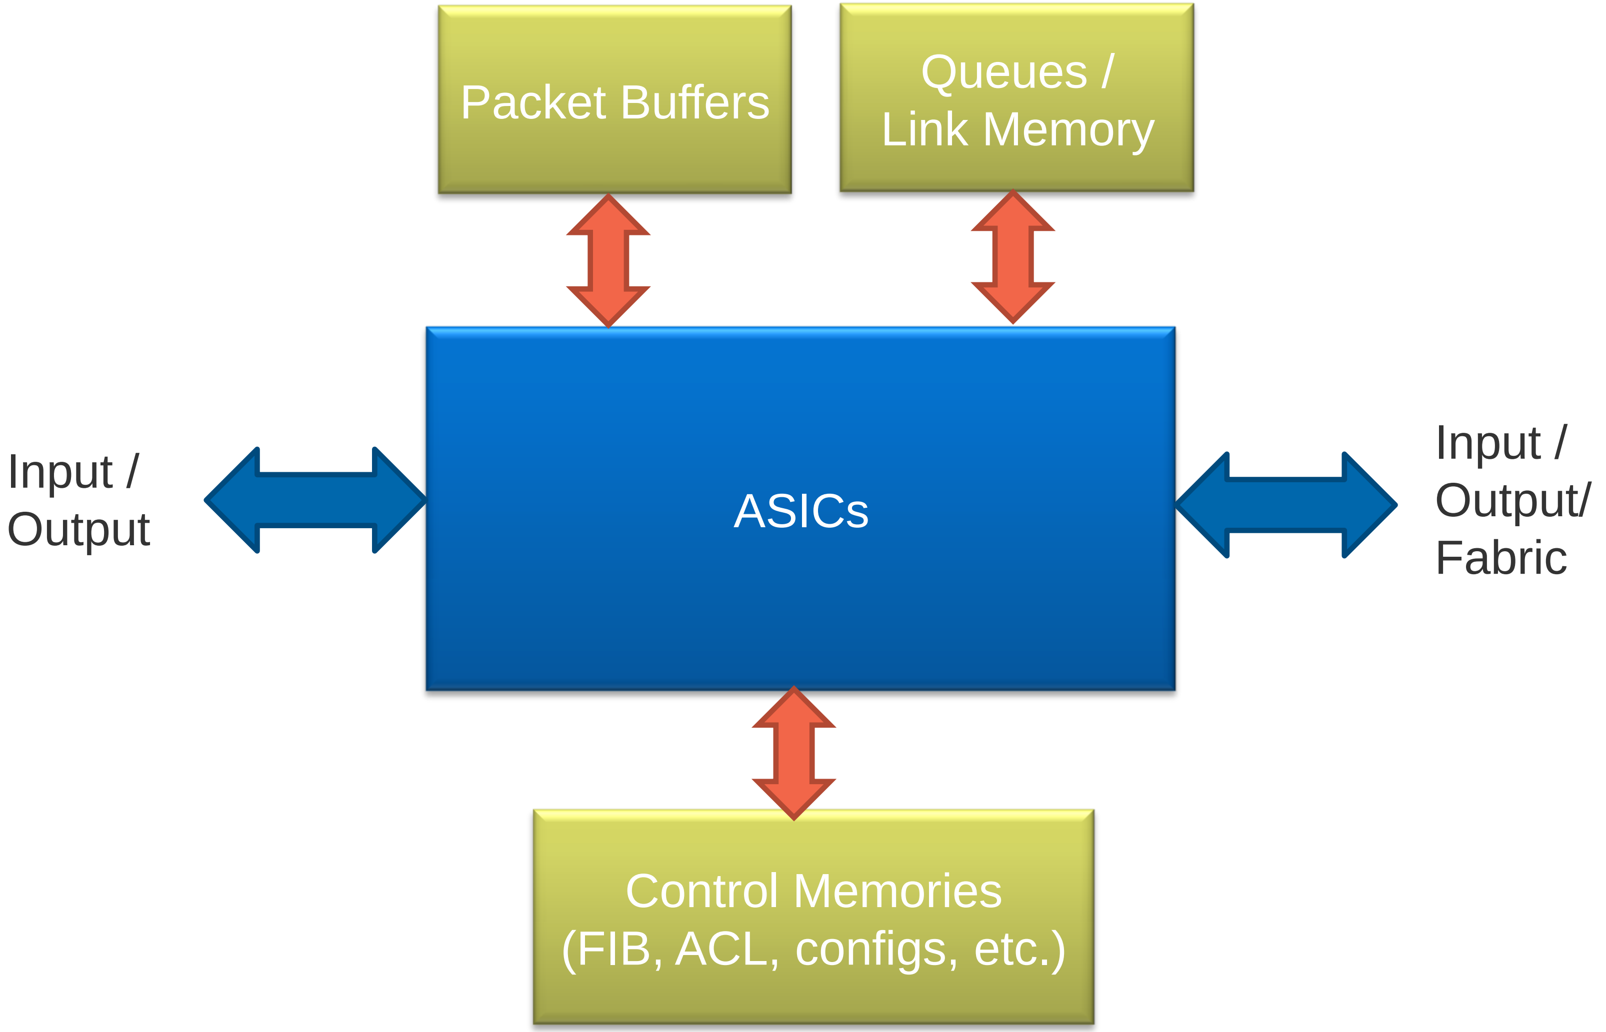

MEMORY TECHNOLOGY CHARACTERISTICS

Technology	Capacity	Frequency	Latency	Power	Cost
Embedded SRAM	L	H	L	M	H
Embedded DRAM	M	M	L+	L	M
Embedded TCAM	L	M	L+	H	H
External SRAM	M	L	M	H	H
External RLDRAM	H	M	M	L	H
External SDRAM	H+	M	H	L	L-
External TCAM	L	L	H	H	H

MEMORY CHOICES WITH NETWORKING ASICS

Packet buffering

- Need high throughput, high density
- Long bursts ok
- SDRAM or RLDRAM (Reduced Latency DRAM)

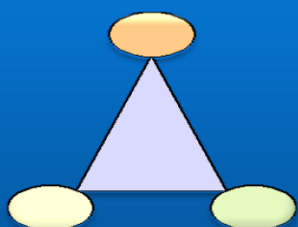

Queuing/Link memory

- Need high throughput, low latency
- Shorter bursts
- SRAM, RLDRAM, or SDRAM

Control memory

- Need high throughput, low latency
- Even smaller access quantum
- SRAM, TCAM, or RLDRAM

ARCHITECTURE – CHIP PARTITIONING

- Fewer chips does not necessarily mean less overall cost
  - Chips get very expensive once they cross a certain die size
  - Economics of silicon is all about fabrication yield
- Goals
  - Balance size of each chip within packet forwarding engine
  - Minimize pin-count on each chip
  - Minimize overall component cost
  - Flexibility of support different configs with the same chipset

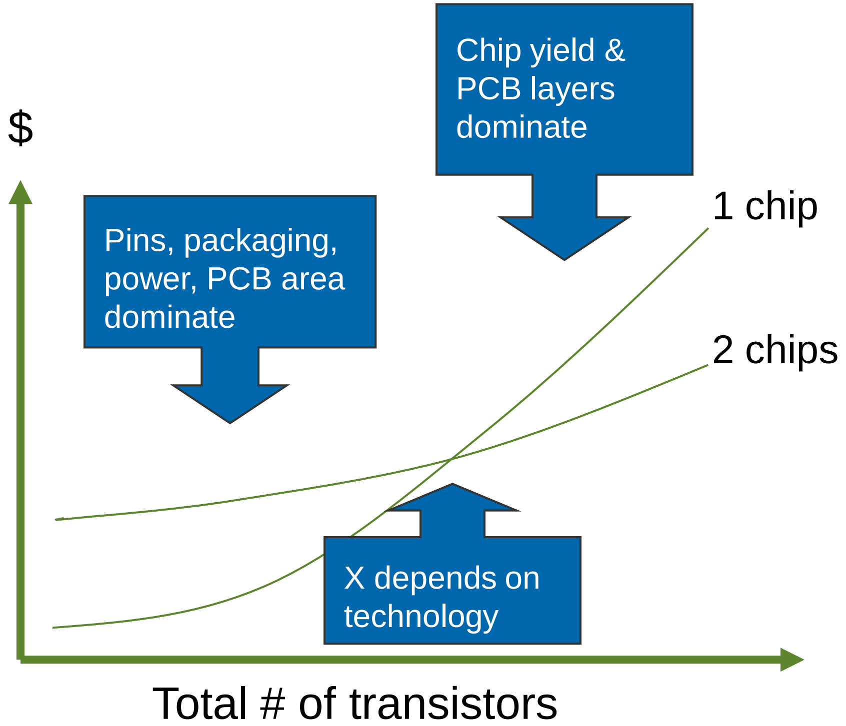

EXAMPLES OF SILICON PROCESS IMPROVEMENT, CHIP PARTITIONING, AND MEMORY USAGE

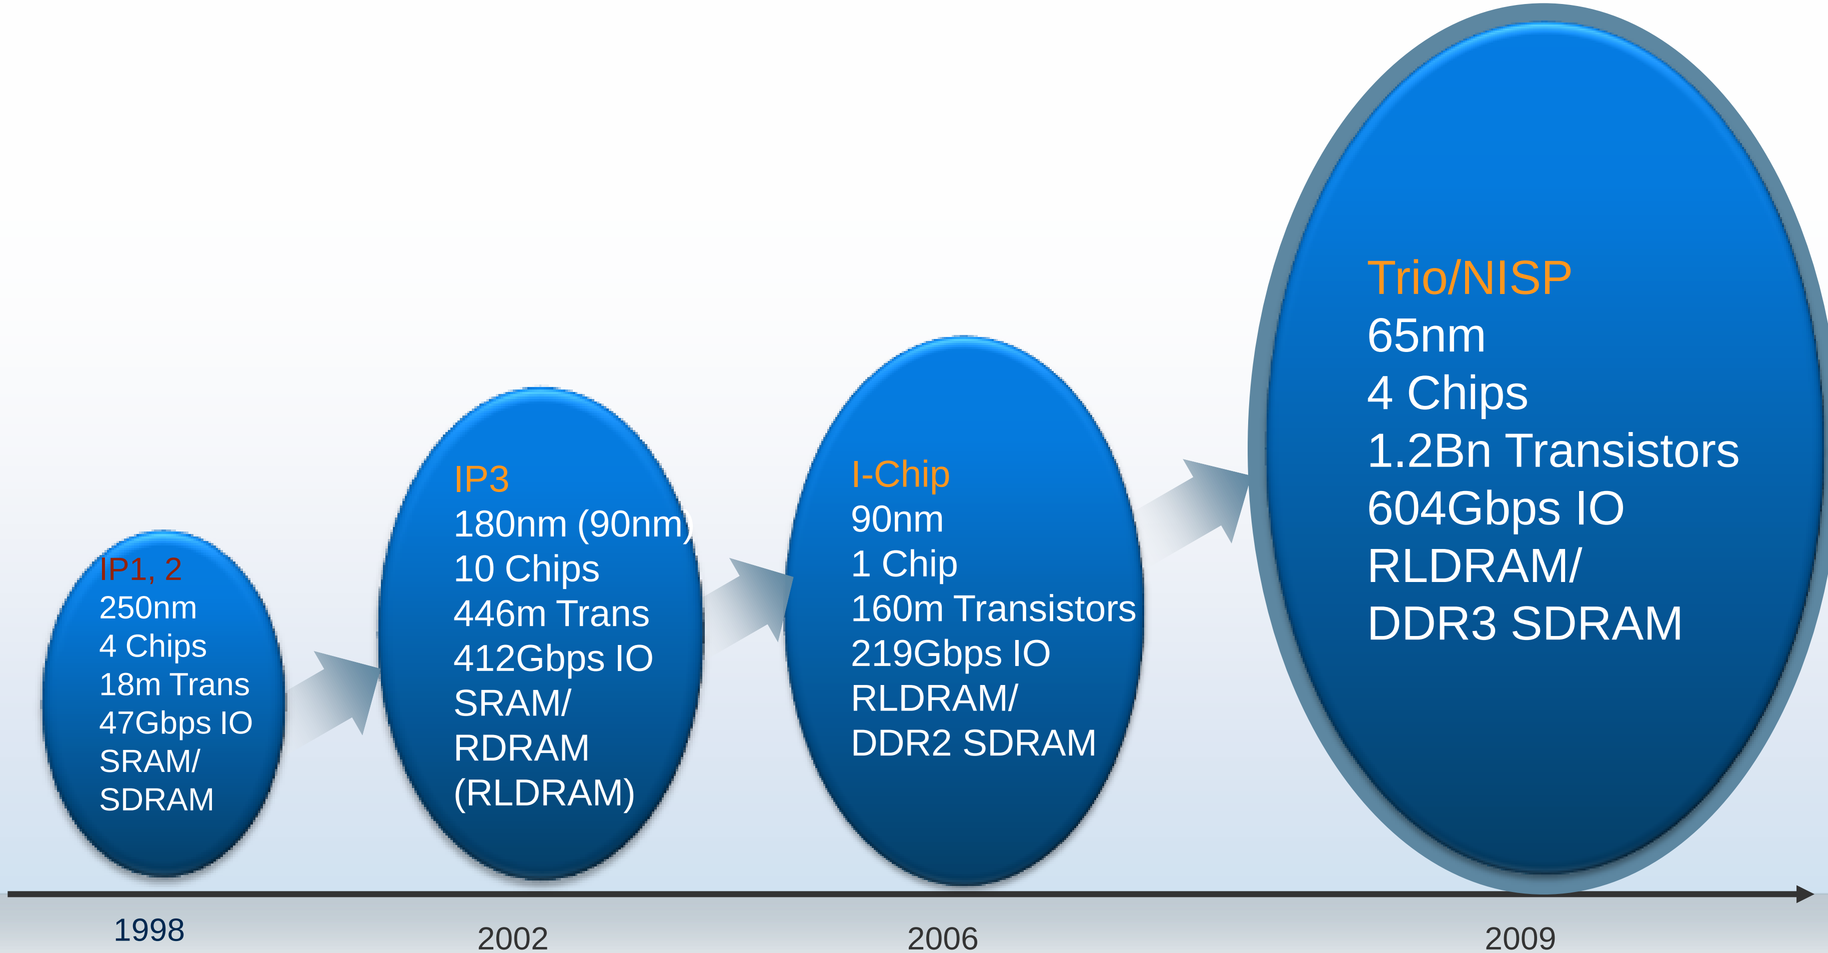

EXAMPLES: BENEFITS OF ASIC EVOLUTIONS

	M40	M160	T640	T1600
Slot Capacity, Gbps	3.0	10	40	100
System Capacity	40Gbps	160Gbps	640Gbps	1600Gbps
Max System Draw	1.5 KW	3.15 KW	4.52 KW	8.35 KW
EER (Gbps/KW)	13	25	71	96
FRS	1998	2000	2002	2007

MICRO ARCHITECTURE

Take each subsystem, divide into blocks, divide each block into sub-blocks, design down to the basic logic elements

Document both functionality and architecture

- Rigorous peer reviews of all documents

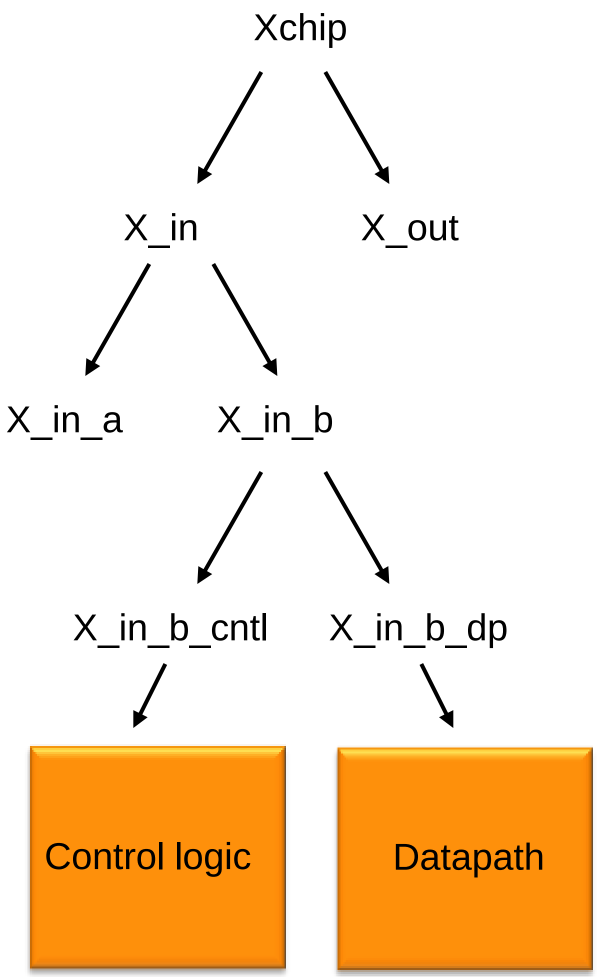

REGISTER TRANSFER LEVEL CODING

Translate micro architecture for all blocks to “Register Transfer Level” code.

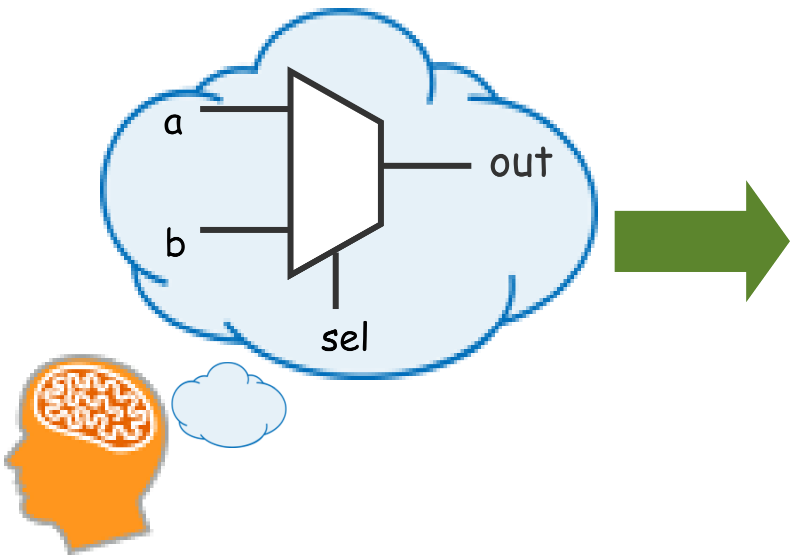


```
always @ (sel or a or b)
    begin
        if (sel == 1)
            out = a;
        else
            out = b;
end
                OR
assign out = sel ? a : b;
```

- A large chip will have hundreds of thousands of lines of RTL code
- Must always keep in mind physical placement and timing during the micro architecture phase
  - You pay now or you pay more later

SYNTHESIS & TIMING

Synthesis is the exercise of mapping RTL to GATES in the technology of choice

INPUT

- RTL code
- Specification of clocks and cycle-time (frequency)
- Input and output constraints for module being synthesized
- Wire-load models as basis to model interconnect effects on gates
- Recent trends: physical synthesis

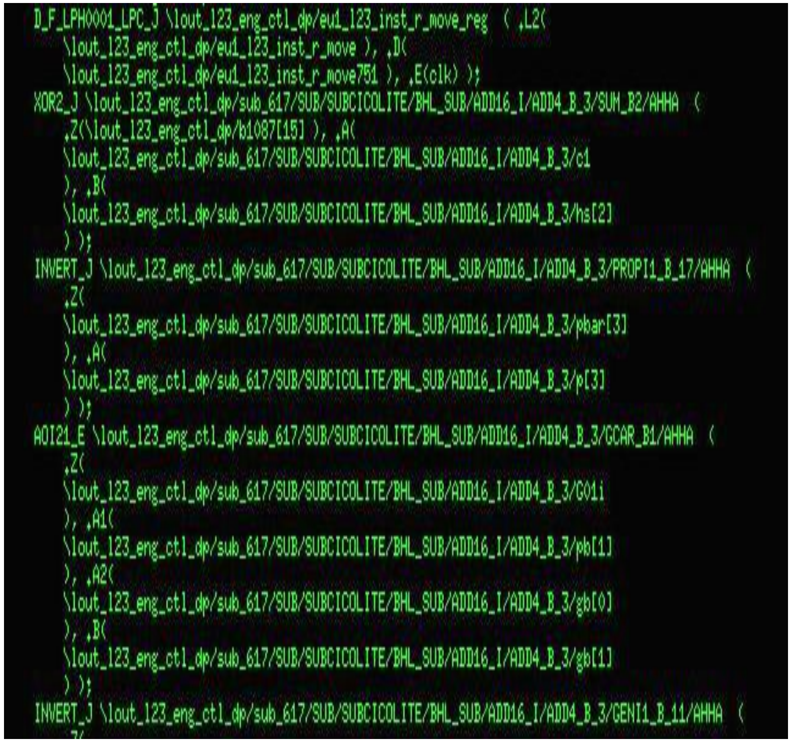

VERIFICATION

Goal: First-time-right silicon

- Avoid expensive ASIC respins
- Simulations are far easier to debug than real chips

Recipe: At least as many verification engineers as design engineers per chip

Performed at multiple levels

- Block level
- Chip level
- Sub-system level
- System level
- Software/hardware co-simulation

TOOLS

Test-bench tool
SystemVerilog
C/C++, Verilog
Coverage tools
Equivalency checkers
Simulators
Waveform viewers

PHYSICAL DESIGN

Power and clock planning

Perform high-level floor-planning

Place I/O, SRAMs, & Register Arrays

Random logic placements

Perform congestion analysis

Wire up all the logic and IOs

Run timing with physical placement

Many iterations of all of the above

PHYSICAL DESIGN EXAMPLE

1) Memory placement
2) Logic placement & clocks
3) M1 routing
4) M2 routing
5) M3 routing
6) M4 routing
7) M5 routing
8) M6 routing
9) M2/M4/M6 routing
10) M1/M3/M5 routing

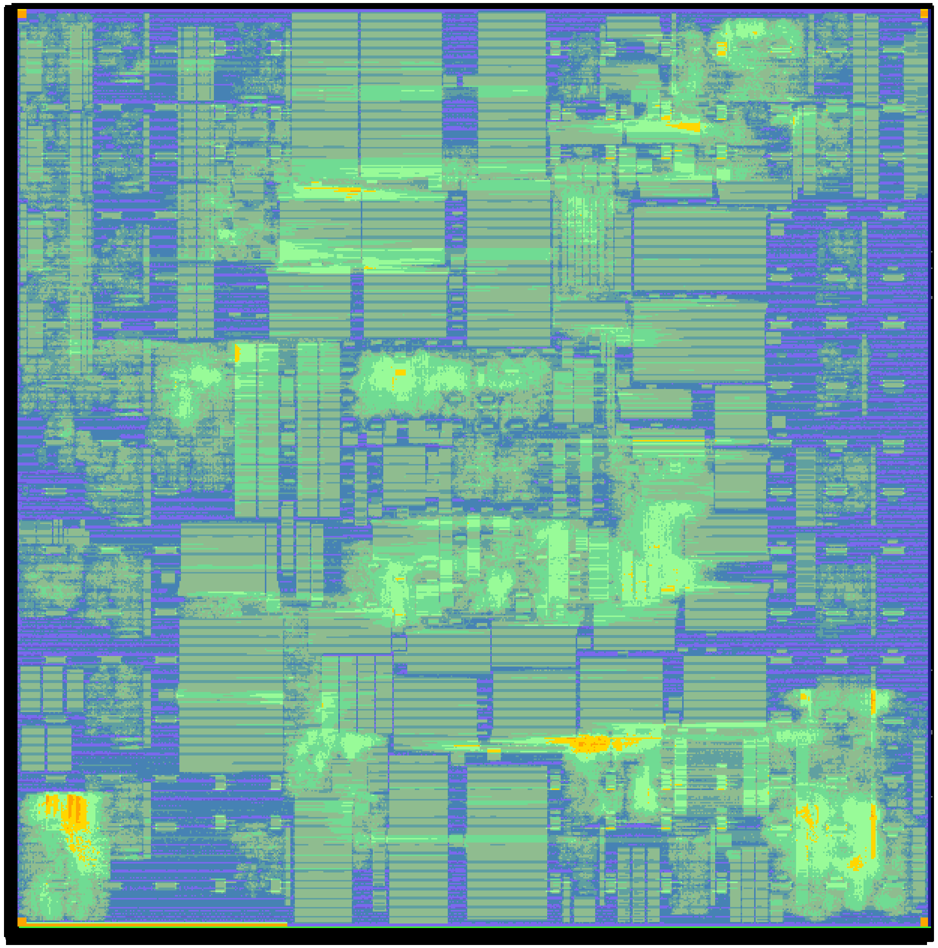

ASIC TAPEOUT

Criteria for ASIC Tapeout

- All functionality complete
- All verification complete
- Performance simulations meet goals
- Chip is error free from a testability perspective
- Chip meets timing under all process, temperature and voltage conditions
- Design and verification database is archived

MANUFACTURING

After the ASIC is taped out

- Masks are generated for photolithography
- ASICs are then built layer-by-layer on a silicon substrate wafer

Once the ASIC wafer is complete

- Each die is tested in wafer test
- Only good die are laser cut for packaging

Once cut die are available

- They are put in a package
- The packaged devices are then tested again

Tested packaged parts are put on system boards

- Test with other hardware and software

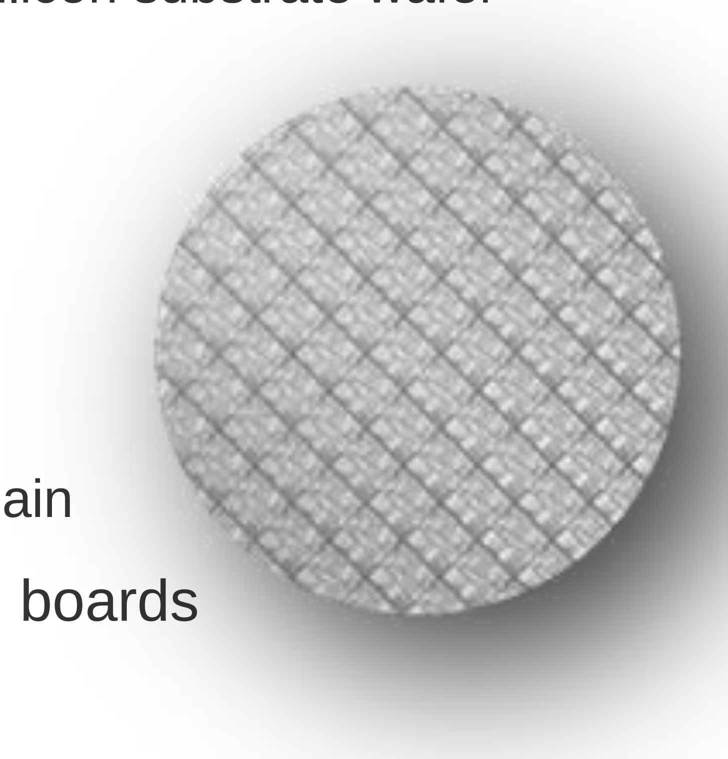

MANUFACTURING – CONTINUED

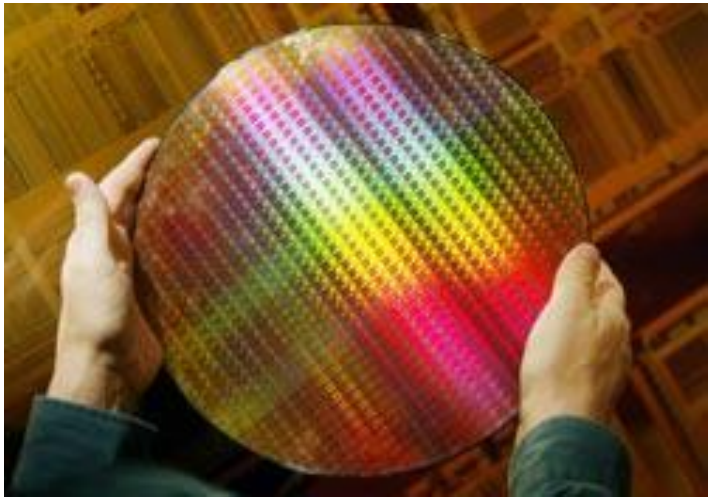

300mm wafer

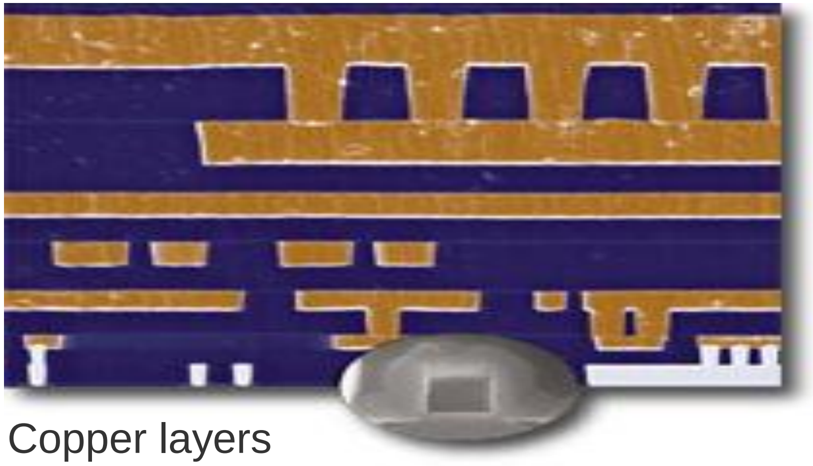

Copper layers

300mm wafer fab

Packaging

SUMMARY

ASIC technology has transformed the network industry

Silicon process technology is evolving at an impressive pace but architectural innovations are required to keep up with the demand for increasing performance at lower power

A rigorous architecture, design, and verification process is required to implement complex networking ASICs

There are a vast amount of architectural and design tradeoffs to be made so user community should provide feedbacks early and often everywhere