

Computer Architecture

计算机体系结构

Lecture 1. Overview: From IC to IDC

第一讲、概述：从集成电路到数据中心


Chao Li, PhD.

李超 博士

SJTU-SE346, Spring 2019

What is Computer Architecture

The science and art of designing, analyzing, selecting and interconnecting hardware components to create computers that meet functional, performance and cost goals



The fabulous processor chip



The fabulous Las Vegas

What is Computer Architecture

- **How is computer different?**

Type I objects: all things not requiring instructions

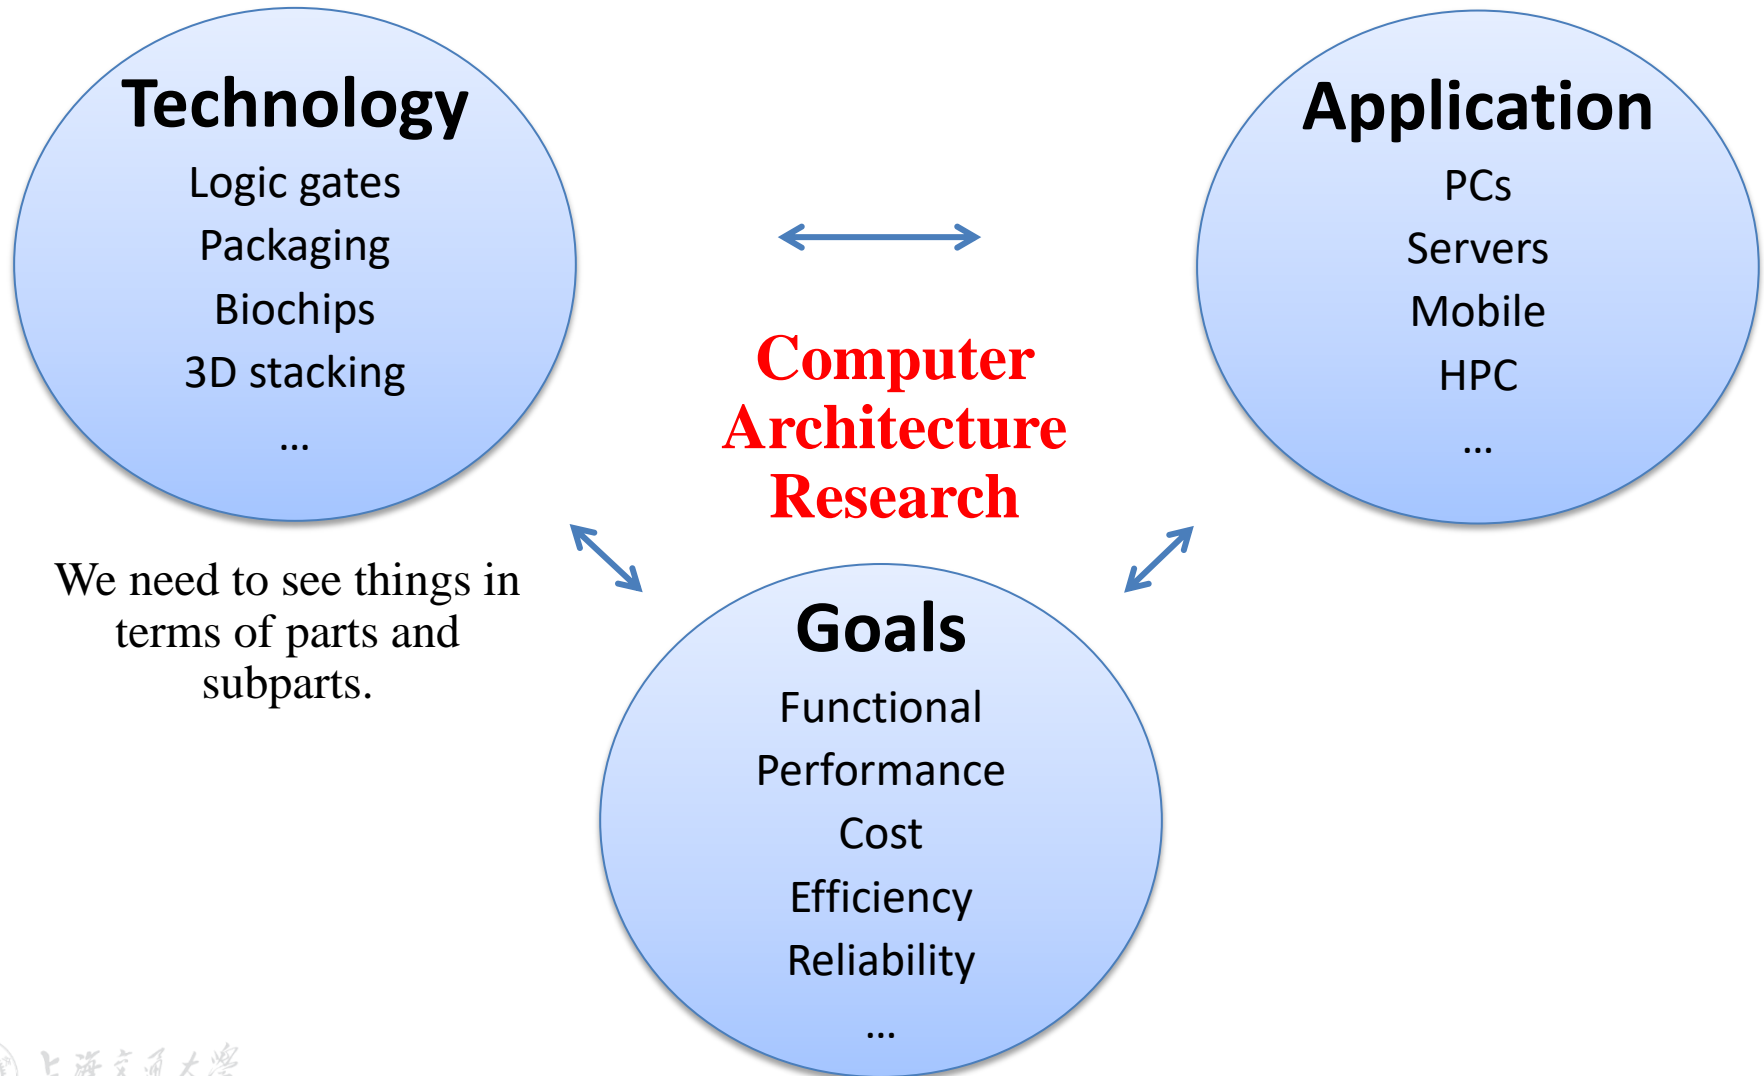
Type II objects: all things require instructions for their formation

anything permitted by chemistry and physics is attainable through the use of appropriate instructions—quite a remarkable statement.


information when properly organized and employed can instruct the formation of very specific and otherwise highly improbable structures

J. E. Mayfield, The Engine of Complexity

What is Computer Architecture



Exponential Growth of Computing




Outlines

- See “Small” in the 20th Century
 - A short history of the IC industry
 - A brief introduction to VLSI
- Think “Big” in the 21st Century
 - The server and data center industry
 - Why energy is a big Issue

The Earliest Electronic General-Purpose Computers

The 1930s and 1940s are considered the beginning of the modern computer era




ENIAC (Electronic Numerical Integrator and Computer)

- Managed by Univ. of Penn
- Completed in 1946
- Operated until 1955
- Performed decimal arithmetic
- Human 20h => ENIAC 30s

Power	Size	Speed	Reliability	Components
150KW	167 m^2	5 KHz	fails every 1~2 days	17468 vacuum tubes, etc.

The Invention of Transistor



The First Point Contact Transistor
(Dec 23, 1947, New Jersey, USA)
Two gold contacts lightly touching a germanium crystal that was on a metal plate connected to a voltage source.



William Bradford Shockley

Prize share: 1/3



John Bardeen

Prize share: 1/3



Walter Houser Brattain

Prize share: 1/3

The Nobel Prize in Physics 1956 was awarded jointly to William Bradford Shockley, John Bardeen and Walter Houser Brattain *"for their researches on semiconductors and their discovery of the transistor effect"*.

- The transistor went on to replace bulky vacuum tubes and mechanical relays.
- It becomes the basic building block upon which all modern computer rests.

The “Traitorous Eight”



The “Traitorous Eight” (1960)



Robert Noyce and Gordon Moore (1970)

1957: Fairchild Semiconductor

- Directly or indirectly involved in the creation of dozens of corporations


1968: Intel Corporation

- Robert Noyce and Gordon Moore


Integrated Circuit

“A body of semiconductor material ... wherein all the components of the electronic circuit are completely integrated”

- Geoffrey Dummer (UK)
 - First conceptualize the idea
- Jack Kilby and Robert Noyce (US)
 - First independently invented IC
 - “Semiconductor device-and-lead structure”
 - US Patent 2,981,877
 - (Noyce filed in 1959, granted in 1961)
 - “Miniaturized Electronic Circuits”
 - US Patent 3,138,743
 - (Kilby filed in 1959, granted in 1964)



Evolution of Processors



Intel 4004 Processor

Introduced 1971

Initial clock speed

108KHz

Number of transistors


2,300

Manufacturing technology

10 μ

4-bit processors. The same computing power as ENIAC. The following 8008 design (8-bit) doubles the computation capability

Evolution of Processors



Intel 8086 Processor

Introduced 1978

Initial clock speed

5MHz

Number of transistors


29,000

Manufacturing technology

3 μ

**The first 16-bit processors. The first x86 CPU.
Up to 10x the performance of 8080**

Evolution of Processors



Intel 486 Processor

Introduced 1989

Initial clock speed

25MHz

Number of transistors


1,200,000

Manufacturing technology

1 μ

32-bit; first tightly pipelined x86; have over 1 million transistors; L1 cache integrated

Evolution of Processors



Intel Pentium Processor

Introduced 1993

Initial clock speed

66MHz

Number of transistors

3,100,000

Manufacturing technology

0.8 μ


The first superscalar IA-32 processor

Our World
in Data

The complexity for minimum component costs has increased at a rate of roughly a factor of two per year

Licensed under CC-BY-SA by the author Max Roser.


Processor Clock Rate Trend



Outlines

- See “Small” in the 20th Century
 - A short history of the IC industry
 - A brief introduction to VLSI
- Think “Big” in the 21st Century
 - The server and data center industry
 - Why energy is a big Issue


Abstraction Layer




From Silicon to Chip



Silicon Ingot



Wafer




“Naked” Die





Packaged Die

Silicon Ingot




Doped Semiconductors







- The advantages of semiconductors emerge when impurities are added (called doping)
 - N-Type: electrons are majority carrier
 - P-Type: hole are majority carrier

Transistor Structure (NMOS)




Equivalent Ideal Switches of Transistors

- Transistors can be thought as a switch controlled by its gate signal




*n*MOS

*p*MOS




*n*MOS ($V_{in} > 0$, $V_{DS} \geq 0$, $V_{GS} \geq 0$)




*p*MOS ($V_{tp} < 0$, $V_{DS} \leq 0$, $V_{GS} \leq 0$)

Transistors in Series/Parallel Connection


- NMOS switch closes when switch control input is high
 - NMOS Transistors pass a “strong” 0 but a “weak” 1 *Q: why?*
- PMOS switch closes when switch control input is low
 - PMOS Transistors pass a “strong” 1 but a “weak” 0 *Q: why?*




$Y = X$ if A and B



$Y = X$ if \bar{A} AND \bar{B}




$Y = X$ if A OR B

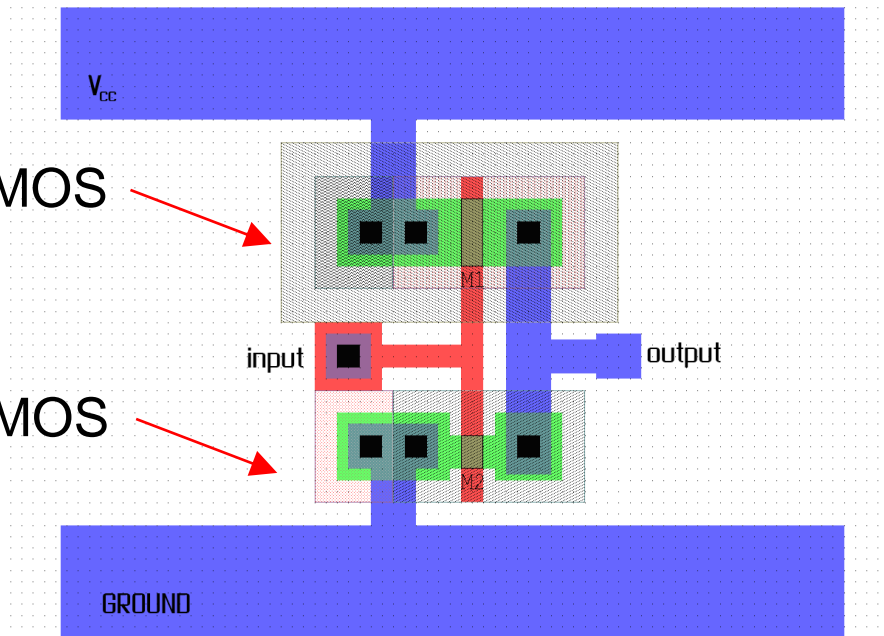


$Y = X$ if \bar{A} OR \bar{B}

From Circuit to Layout



CMOS Inverter



Active areas for thin oxide region



Polysilicon for the gate



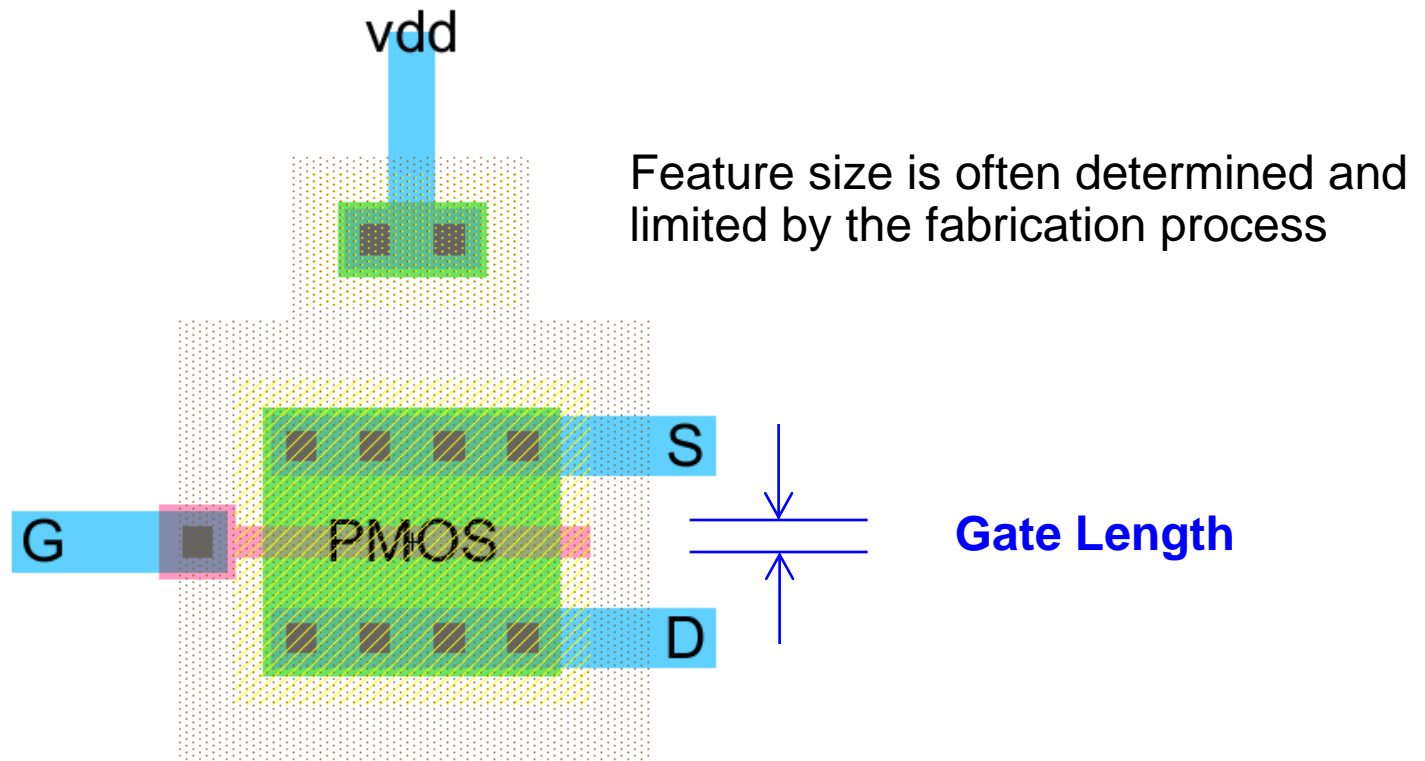
Metal for interconnection







Contact (Via) for inter-layer connection

Transistor Gate Length


- The minimum process dimension is feature size = 2λ
 - Feature size reflect the typical length of a transistor channel (**gate**)
 - In 1978, $\lambda = 1.5 \mu\text{m}$ (a.k.a. 3 micrometer technology)
 - In 2004, $\lambda = 0.045 \mu\text{m}$ (a.k.a. 90 nanometer technology)



Integrated Circuit Fabrication

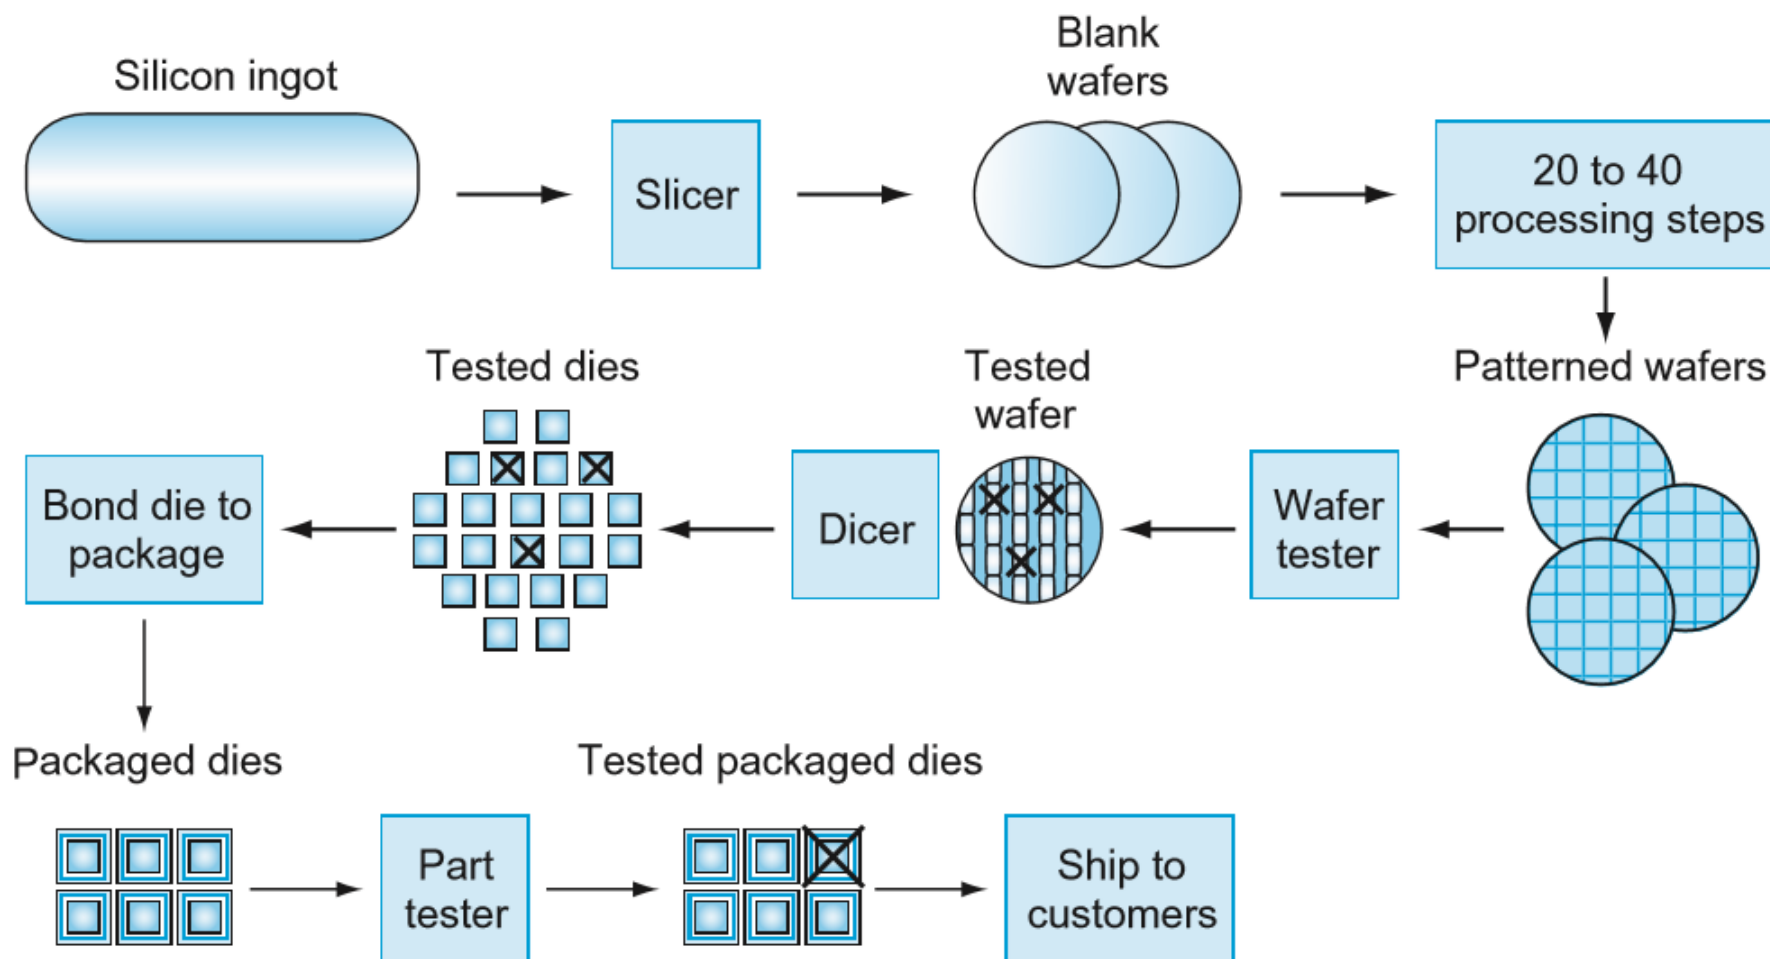


Moore's Law



Feature size versus calendar year

Manufacturing Process




Cost per Die

$$\text{IC cost} = \frac{\text{Die cost} + \text{Testing cost} + \text{Packaging cost}}{\text{Final test yield}}$$

$$\text{Die cost} = \frac{\text{Wafer cost}}{\text{Dies per Wafer} * \text{Die yield}}$$

$$\text{Dies per wafer} = \frac{\pi * (\text{Wafer_diam} / 2)^2}{\text{Die Area}} - \frac{\pi * \text{Wafer_diam}}{\sqrt{2 * \text{Die Area}}} - \text{Test dies}$$



$$\text{Die Yield} = \text{Wafer yield} * \left\{ 1 + \frac{\text{Defects_per_unit_area} * \text{Die_Area}}{\alpha} \right\}^{-\alpha}$$

- $\text{Die Cost} \propto \text{die area}^4$

Die Cost: Some Example (1994)

Chip	Metal layers	Line width	Wafer cost	Def./cm ²	Area mm ²	Dies/wafer	Yield	Die cost
386DX	2	0.90	\$900	1.0	43	360	71%	\$4
486 DX2	3	0.80	\$1200	1.0	81	181	54%	\$12
Power PC 601	4	0.80	\$1700	1.3	121	115	28%	\$53
HP PA 7100	3	0.80	\$1300	1.0	196	66	27%	\$73
DEC Alpha	3	0.70	\$1500	1.2	234	53	19%	\$149
Super Sparc	3	0.70	\$1700	1.6	256	48	13%	\$272
Pentium	3	0.80	\$1500	1.5	296	40	9%	\$417

Outlines

- See “Small” in the 20th Century
 - A short history of the IC industry
 - A brief introduction to VLSI
- Think “Big” in the 21st Century
 - The server and data center industry
 - Why energy is a big Issue

Scalability

The capability of a system, network, or process to handle a growing amount of work, or its potential to be enlarged in order to accommodate that growth


- **Scale Out** (Horizontal Scaling)
 - Add more components to a system
 - e.g. double the nodes in a cluster
- **Scale Up** (Vertical Scaling)
 - Add resources to a single component in a system
 - e.g. upgrade your memory

Servers

- Tower/Rack-Mounted/Blade/Mainframe
- Usually accessed only via a network
- Engineering/Scientific/Business application




From Server to Large-Scale Systems



1U server

A rack unit (abbreviated **U** or **RU**) is a unit of measure defined as 1.75 inches



Supercomputer (HPC Center)

- Scientific computing
- Engineering tasks
- Military projects
- High-quality components
- Throughput matters



Sunway TaihuLight
15.4MW, 6Gflops/W




TIANHE 2
17.8 MW



TITAN
8.2 MW



Sequoia
7.9 MW



The K
12.7 MW



MIRA
3.9 MW

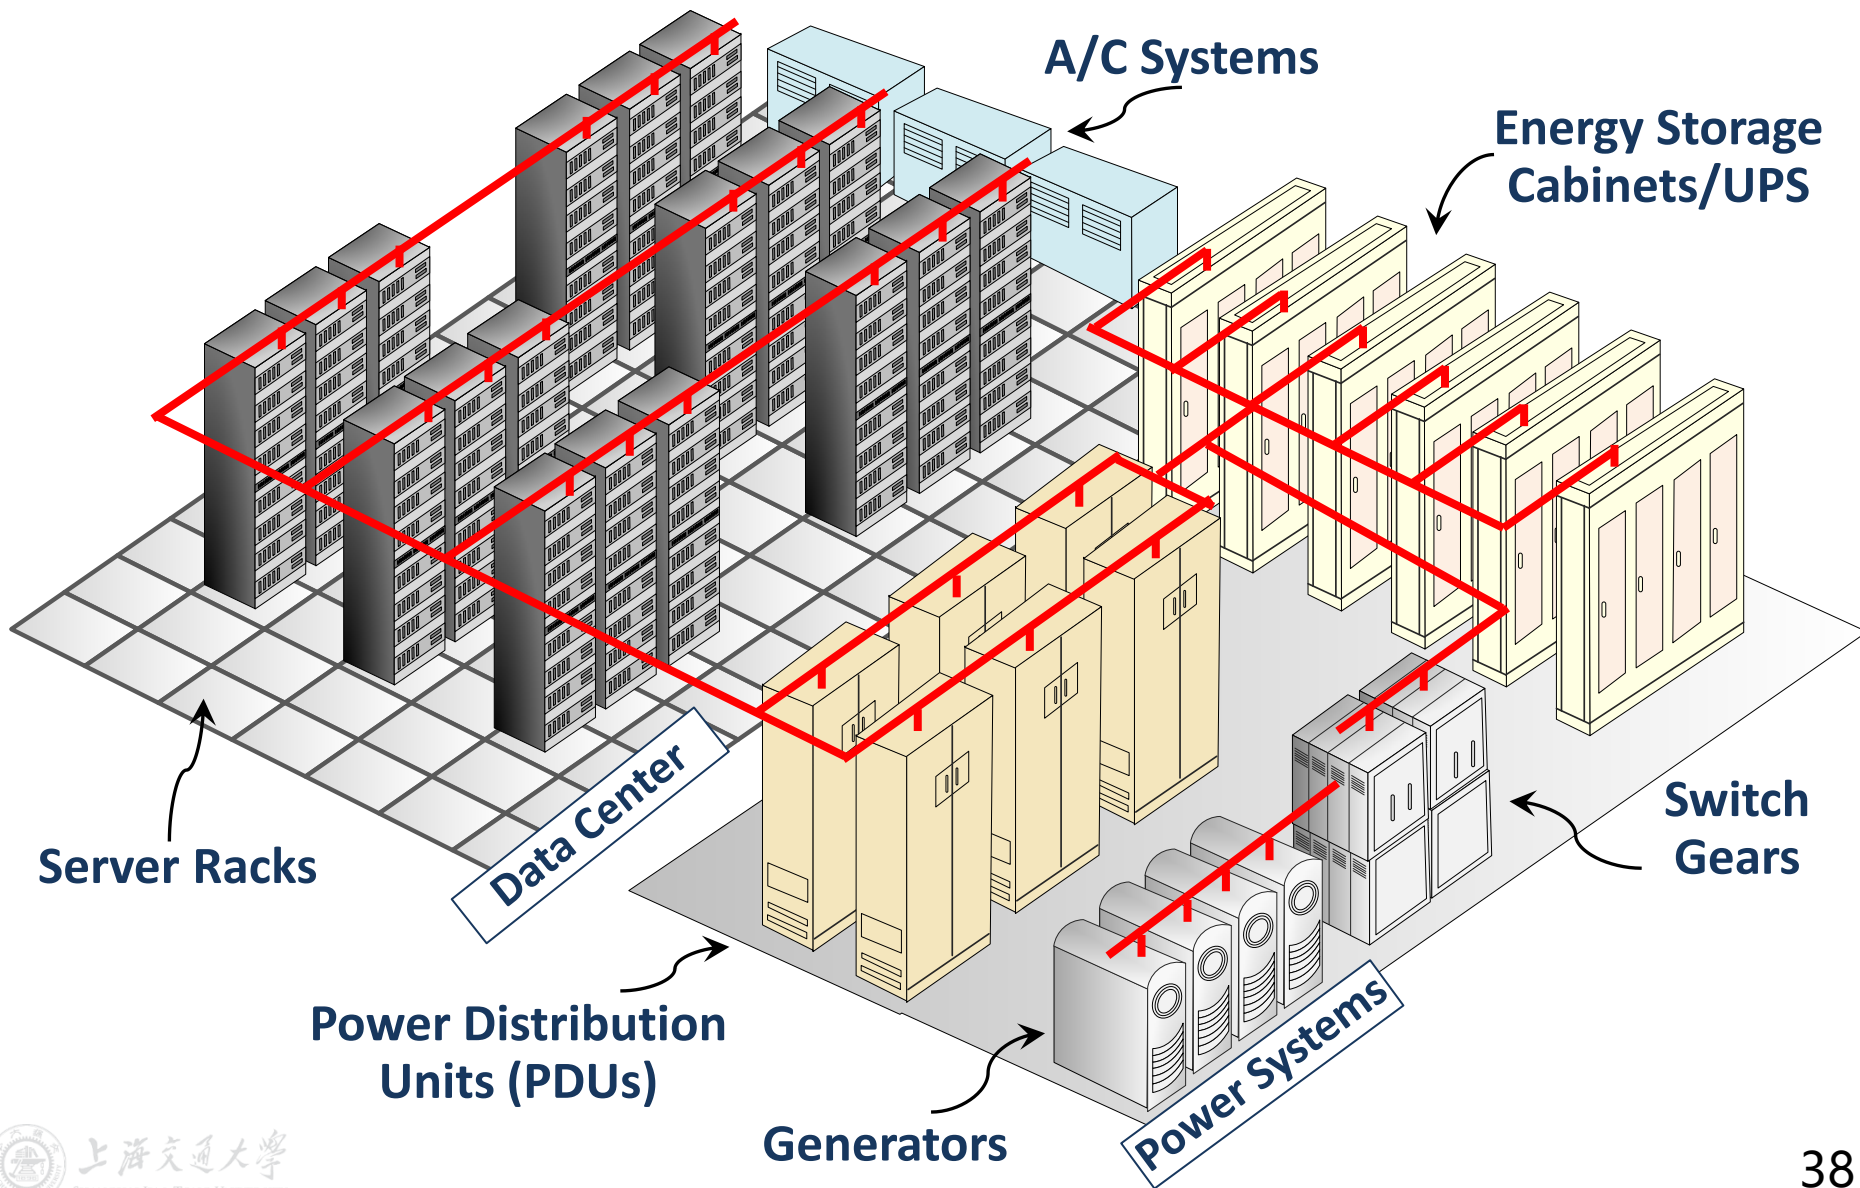
Internet Data Centers (IDC)

Email Services
Search
Social networking
Online maps
Online gaming
Video sharing
File sharing
E-Business
Cloud computing
...




- Inexpensive, commodity components
- Quality of service (latency) matters


A Typical Data Center



The Three Pillars of a Data Center



Data Center TCO




- TCO: Total Cost of Ownership
 - CapEX (Capital Expenditure) + OpEx (Operational Expenditure)

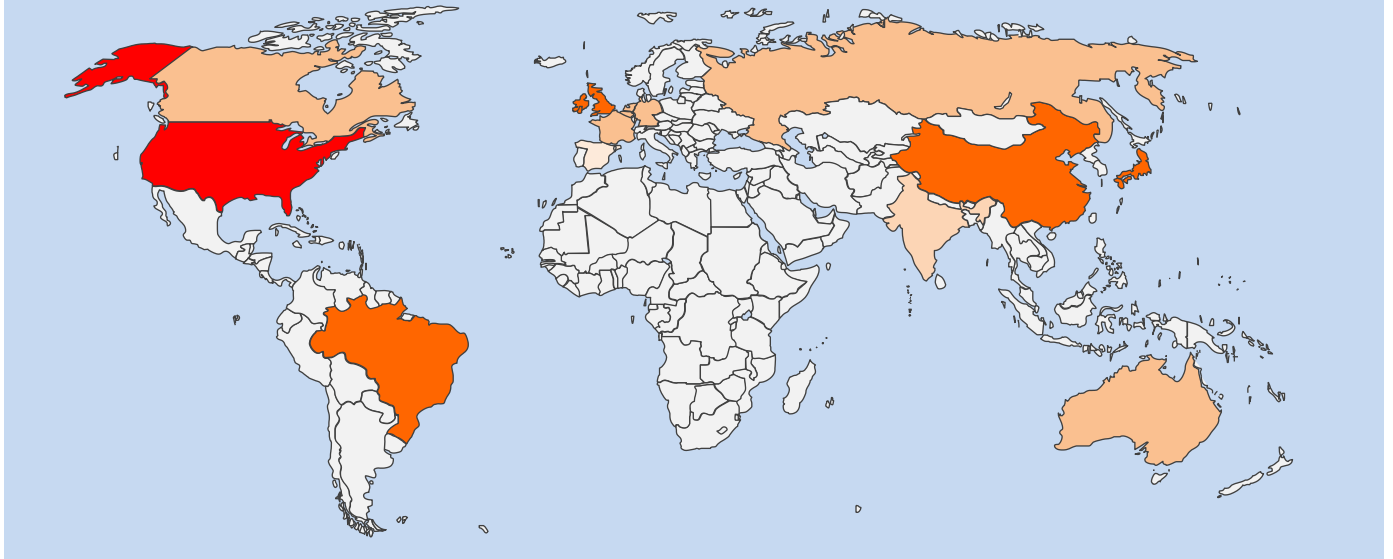
Outlines

- See “Small” in the 20th Century
 - A short history of the IC industry
 - A brief introduction to VLSI
- Think “Big” in the 21st Century
 - The server and data center industry
 - Why energy is a big Issue

Energy Consumption Issue



The increase in server energy demand (2012-2013)^[2]



- The global data center electricity usage in 2012: 300 ~ 400 TWh
 - 2% of global electricity usage
 - Expected to triple by 2020 ^[1]



302 TWh
TOTAL ENERGY
was consumed
in **CA** in 2012 ^[1]


[1] C. Belady, Projecting Annual New Datacenter Construction Market Size, Global Foundation Services, 2011

[2] DCD Industry Census 2012: Energy, <http://www.dcd-intelligence.com/>


[3] http://energyalmanac.ca.gov/electricity/total_system_power.html

Energy Consumption Issue: Cost

The 3-Year Energy Expenditure
(% of Total IT Equipment Cost)



Historical Electricity Prices in UK



- Escalating energy consumption drives data center cost up
 - Need to think alternative power provisioning solutions

[1] Conference report: The Future of the Data Centre, <http://www.information-age.com>


[2] Ken Brill, The Economic Meltdown of Moore's Law and the Green Data Center

[3] <https://www.gov.uk/government/organisations/department-of-energy-climate-change>

Energy Consumption Issue: Environmental Impact




Hurricane Sandy, 2012
(Northeastern US)




Typhoon Haiyan, 2013
(Southeast Asia)

- The greenhouse effect & climate change
- 1MW data center → 10~15 Kt CO₂ yearly
- Data centers are carbon-constrained:
 - They must cap carbon emissions



Power Capacity Issue

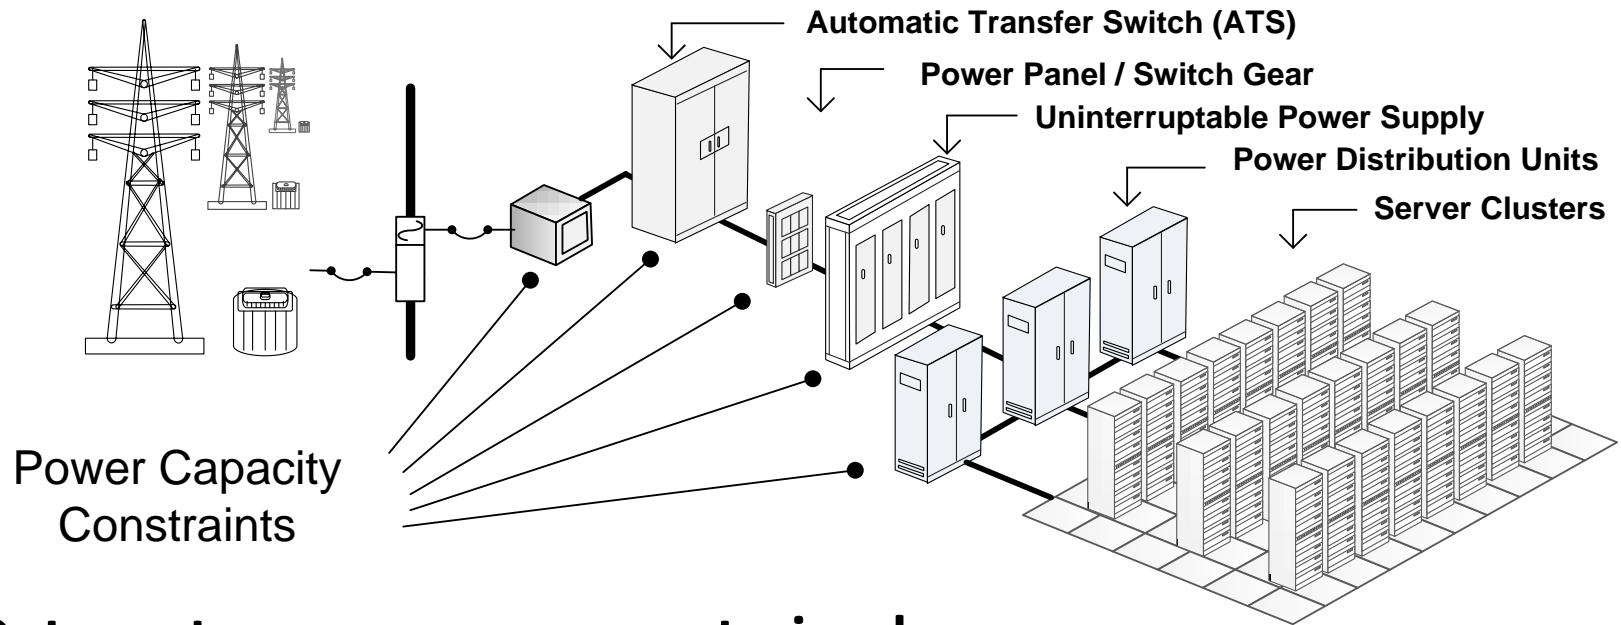


“Exascale computers ... need a dozen nuclear power stations to run it.”

- E&T: the Engineering & Technology Magazine




Power Capacity Issue: Scalability




- **Datacenters are power-constrained:**
 - Limited power capacity headroom *Q: how do you understand it?*


Run out of power capacity in 2012 ?



Capacity expanded in the last 5 years?



How To Scale Power Capacity?



[1] the Uptime Institute 2012 Data Center Industry Survey, 2012

Summary

- What is computer architecture
- History of IC
- Transistor basics
- Feature length
- HPC vs IDC
- Scale up/out
- Energy/power issues
- The trend of computer architecture research