Network-on-Chip: Hype or new System-Paradigm?

Claas Cornelius
University of Rostock, Germany

GRK-Workshop
Schwarzenhof, November2005

Outline

- Classification and Trends
- Historical Origin
  - User-View
  - Problems of Development/Architecture
- Network-on-Chip (NoC)
  - Solution Attempts
  - Characterization
- Summary

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Classification

„Computer as the workhorse of the Semiconductor Industry“

- Motivation
  - Performance
  - Flexibility
  - Mobility

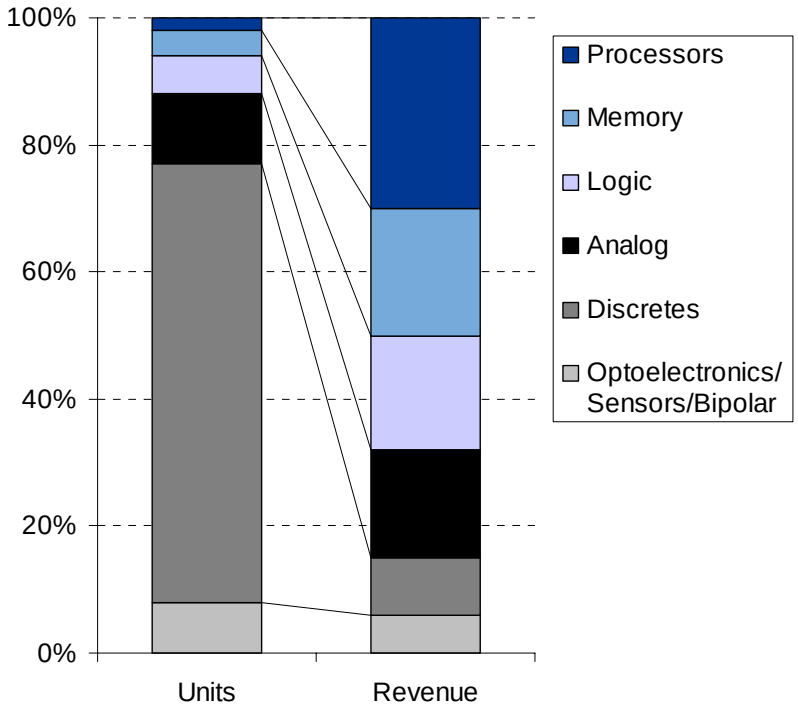

Source: WSTS ‘02

~ 0,04% of all sold units are General Purpose Processors (GPP)
~ 15% of the total revenue is based on GPPs
~ 40% of all semiconductors are used in the PC-area

Trends

- Mobile phones
- Paddles
- Computers

Applications:

- Accounting
- Text processing
- Games
- Multimedia
- Web and Communication
- ...

Algorithm-on-Chip

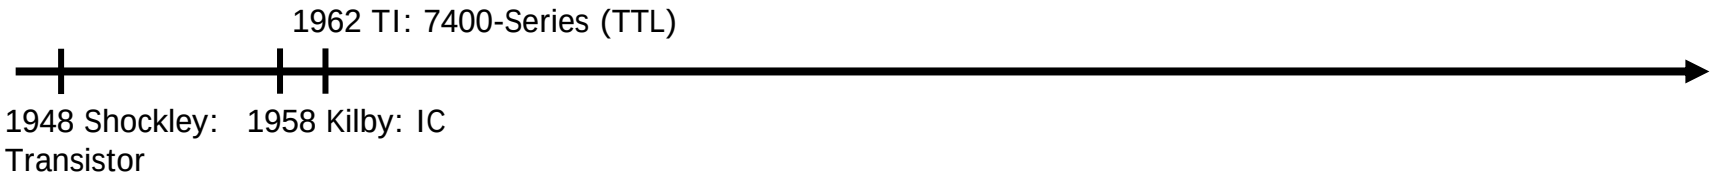

E.g.: Fulladder

$S = A \text{ XOR } B \text{ XOR } C_{in}$

$C_{out} = AB + (A \text{ XOR } B)\, C_{in}$

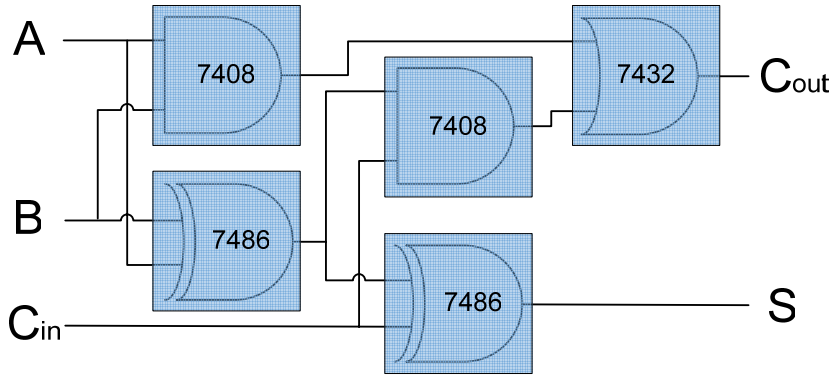

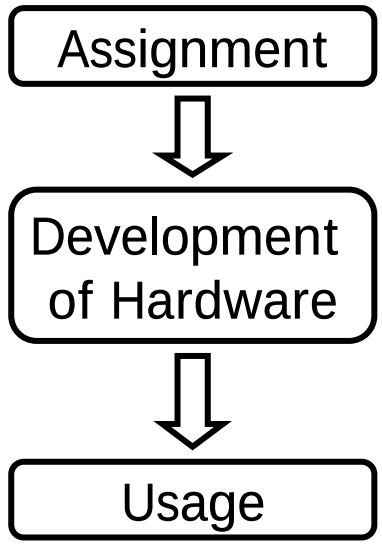

Algorithm-on-Chip

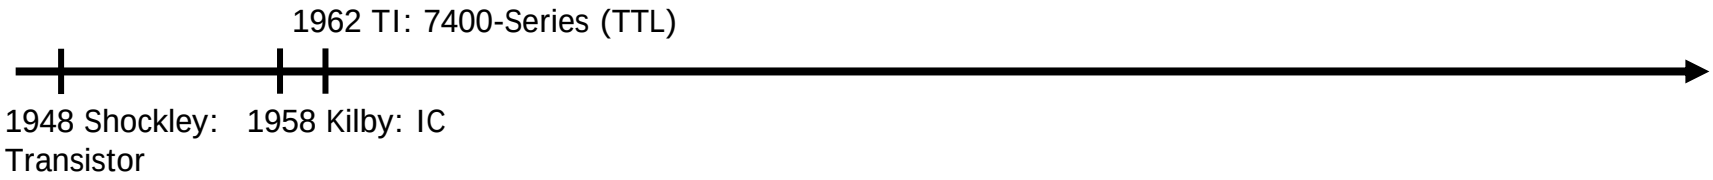

E.g.: Fulladder

$S = A \text{ XOR } B \text{ XOR } C_{in}$

$C_{out} = AB + (A \text{ XOR } B)\, C_{in}$

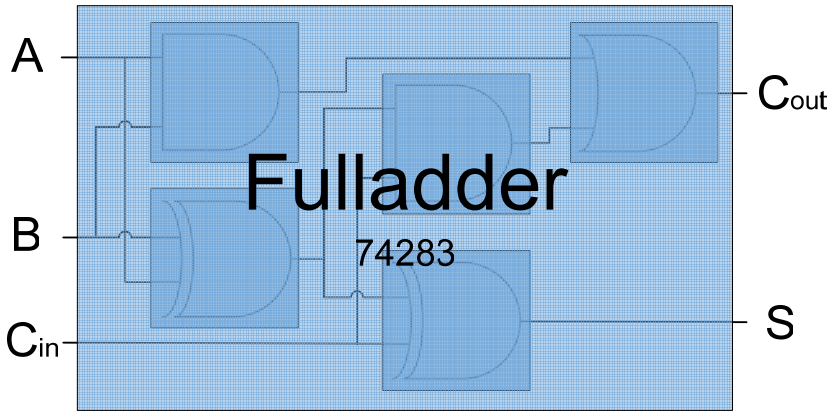

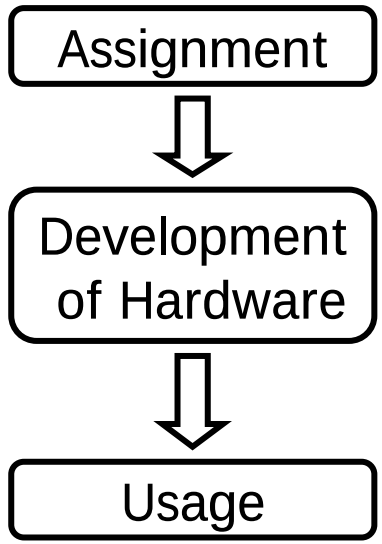

Algorithm-on-Chip

- ➔ Hardwired Communication
- ➔ Hardwired Computation

No flexibility

Mainframe

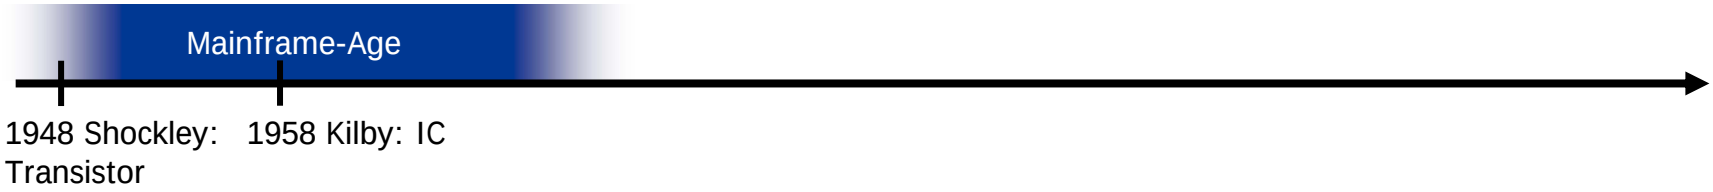

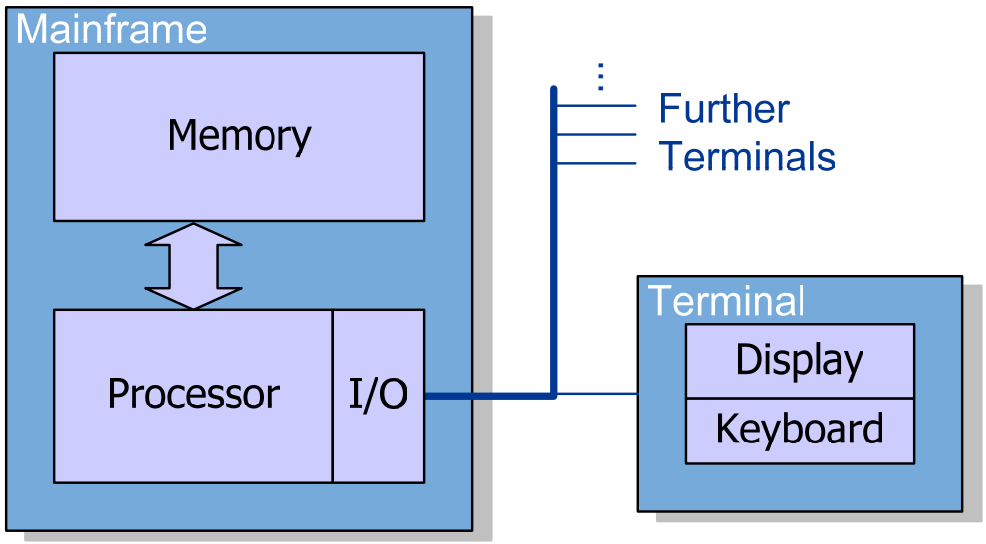

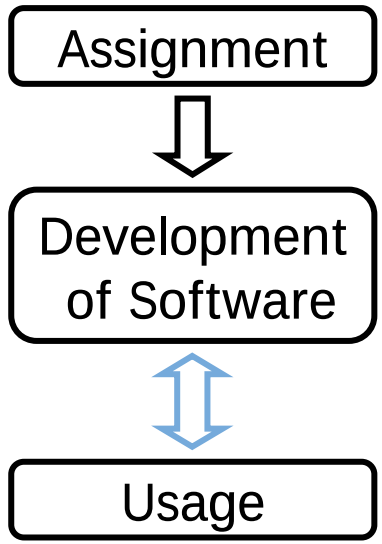

- Mainframes
  - Huge and expensive
  - Programmable

Flexible

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Mainframe

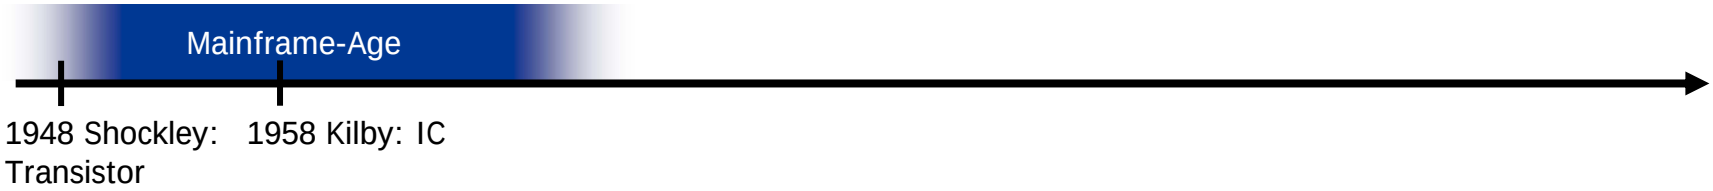

First commercial Mainframe:

- IBM 7030 (Stretch), 1961
  - Purchase price: $13,5 Mio.
  - Size: 232 m²
  - Power consumption: 21,6 kW
  - 3 MFlops
  - 169.100 Transistors

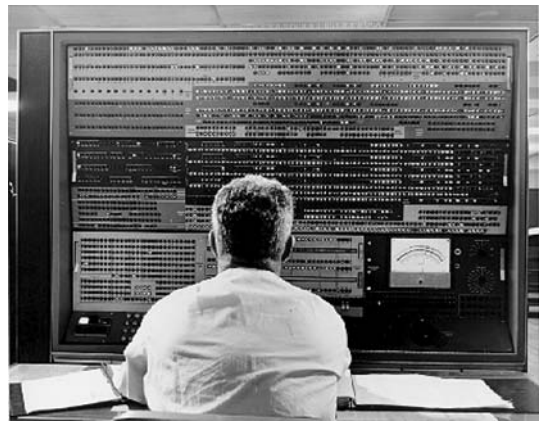

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Efficiency of Hardware

E.g.: Decoding/Playing of mp3-Files

	Frequency	Power
PC	80 MHz	20 W
mp3-Player	12 MHz	0,1 W

200 times

ASSP: Application Specific Standard Product 10-1.000 times

ASIC: Application Specific Integrated Circuit 1.000-10.000 times

Personal Computer

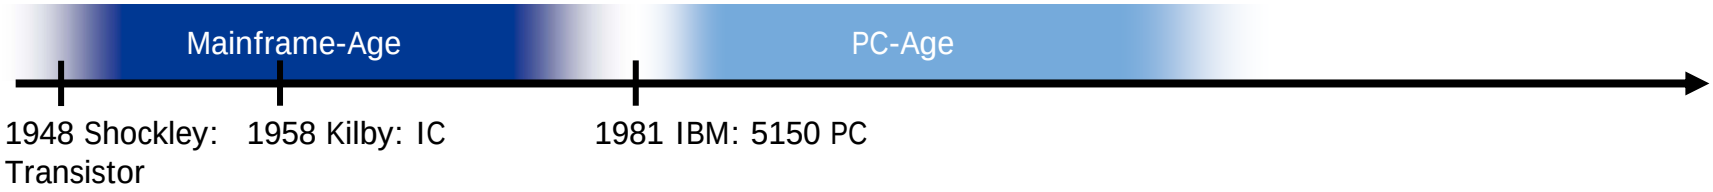

Structure of PCs

- Von Neumann-Architecture
- Hardware-Components
  - CPU
  - Graphic
  - Audio
  - I/O (USB, LAN ...)
- Bus-System
  - ISA, PCI
  - PCI-Express

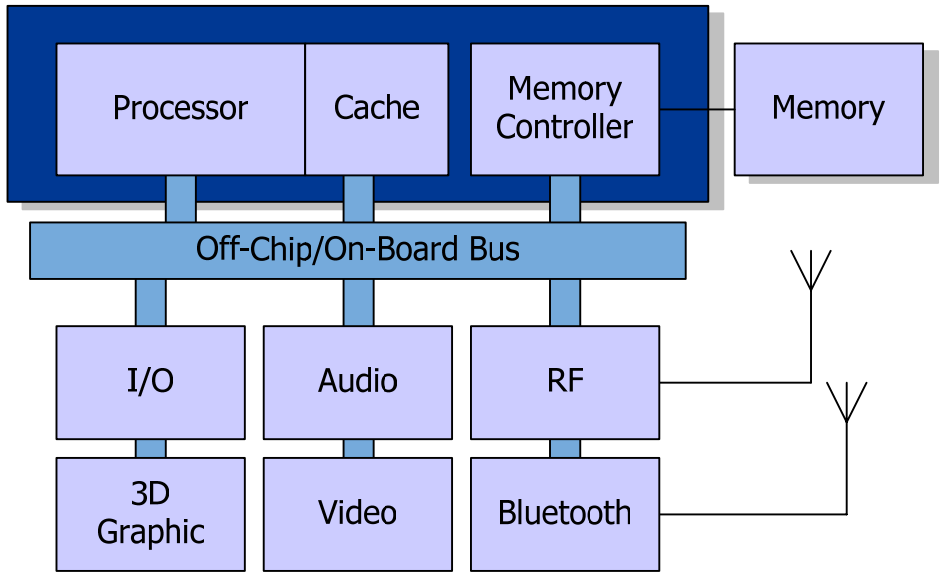

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Productiviy-Crisis

- More Transistors per Chip
- Growth in Productivity too small
  - ➔ Design-Productivity-Gap
    - Better Design-Tools ???
    - New System-Paradigm ???

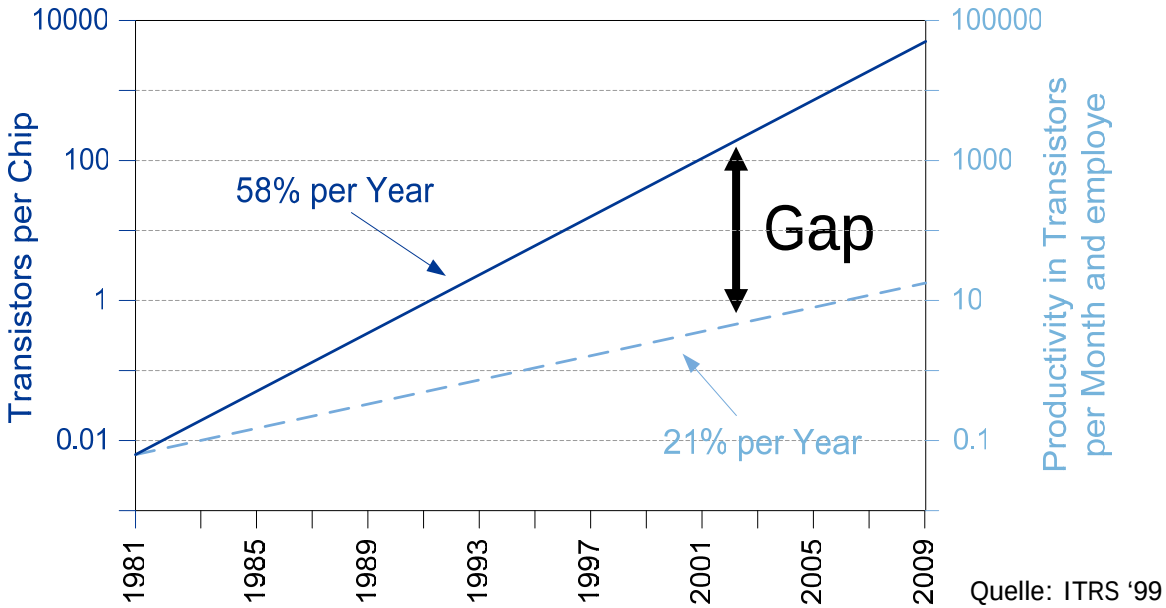

Quelle: ITRS '99

125 Mio. Transistors
Intel Prescott (2004)

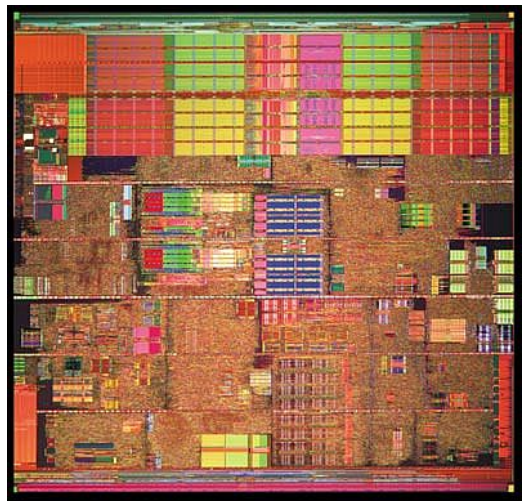

Scale: Number of Transistors

29k Transistors
Intel 8088 (1978)

System-on-Chip

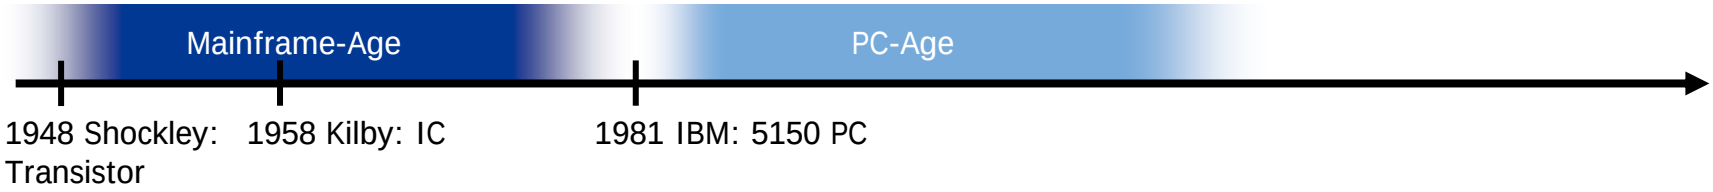

SoC-Example:

- Infineon ADM8668
  „ WildPass “
  - CPU, WLAN, Ethernet, USB, PCI, IDE

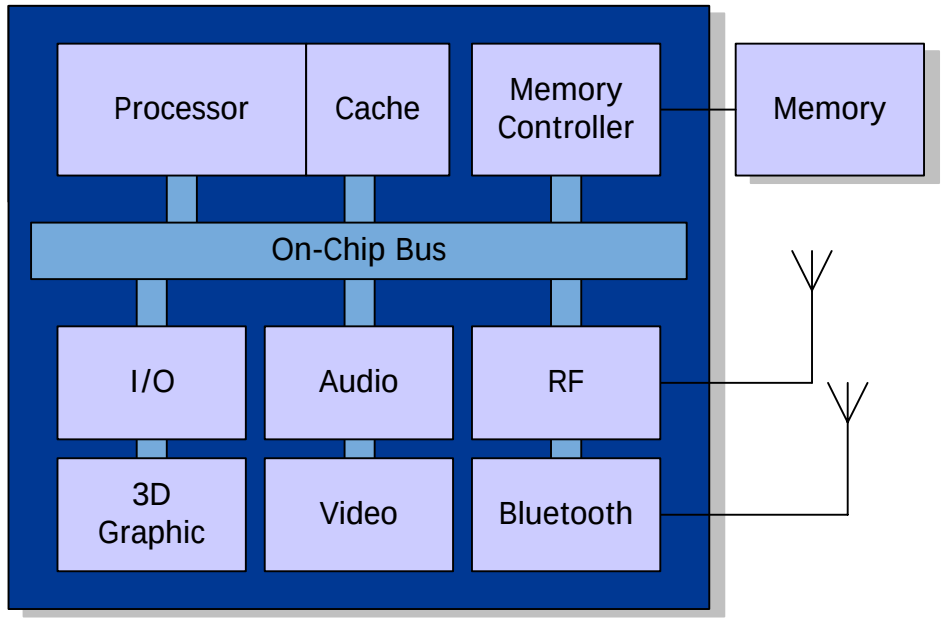

System-on-Chip

➔ Hardwired Communication

➔ Programmable Computation

Synchronization

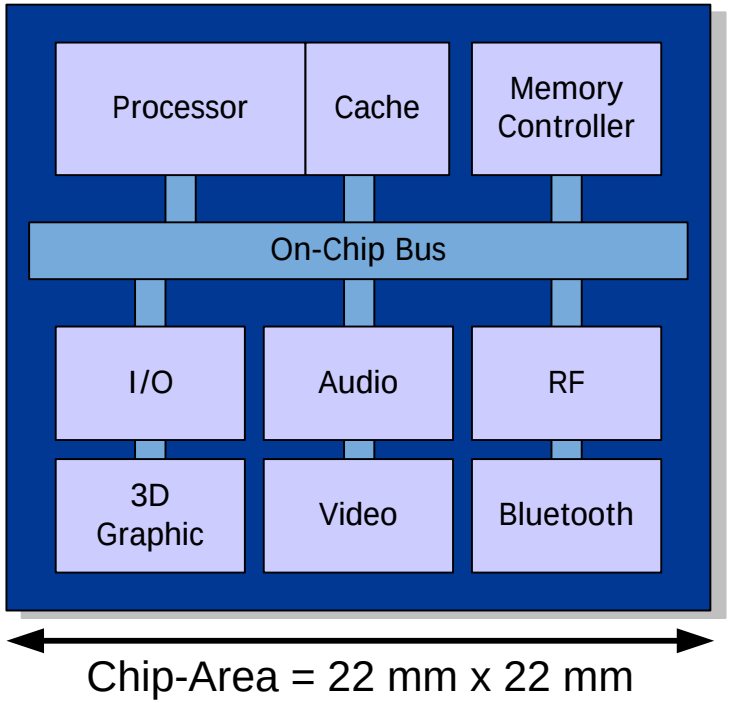

E.g.: Clock distribution (optimal)

$$f = 4\,\text{GHz} \rightarrow T = 0{,}25\,\text{ns}$$

Velocity of propagation

$$v = \frac{2}{3} c_0 \approx 200 \frac{\text{mm}}{\text{ns}} = 50 \frac{\text{mm}}{\text{T}}$$

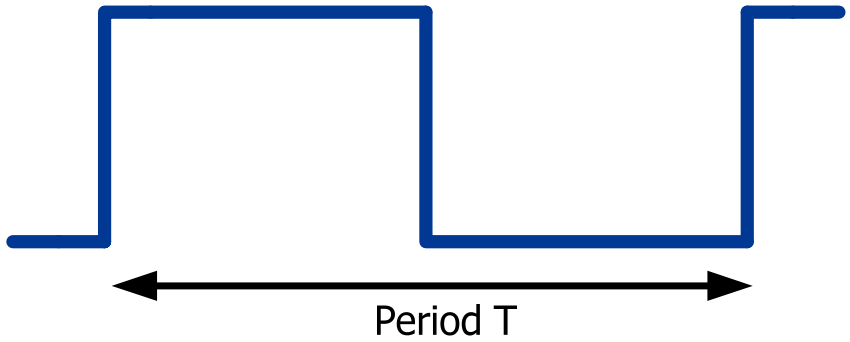

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Synchronization

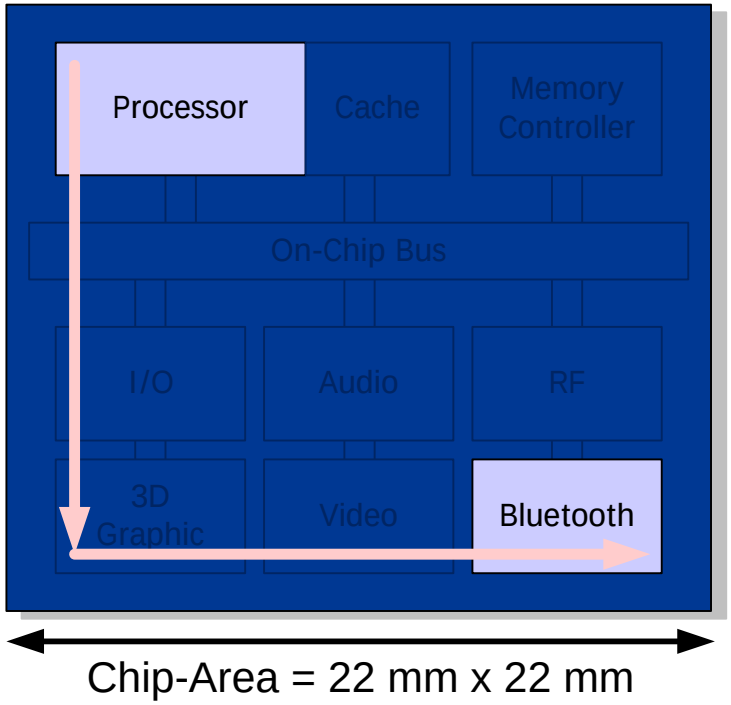

E.g.: Clock distribution (optimal)

$$f = 4\ \text{GHz} \rightarrow T = 0{,}25\ \text{ns}$$

Velocity of propagation

$$v = \frac{2}{3} c_0 \approx 200 \frac{\text{mm}}{\text{ns}} = 50 \frac{\text{mm}}{\text{T}}$$

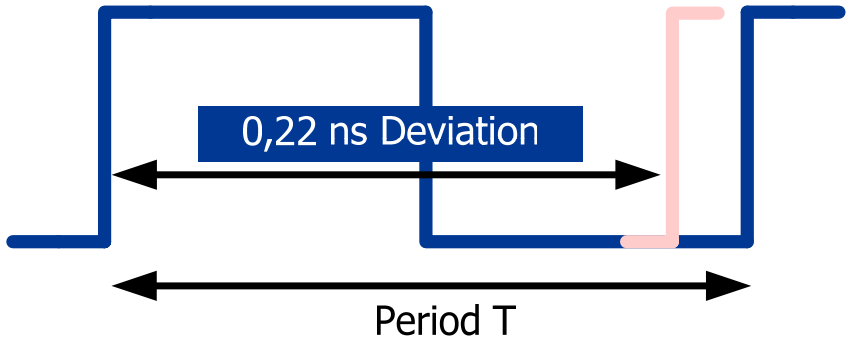

- Frequency rises
- Chip-Area increases
- Velocity of propagation is constant or decreases
  - ➔ Synchronization not/only with extensive effort obtainable
  - ➔ Clock distribution accounts for 20-25% of the total power consumption

Interconnectivity

- Wires between the logic dominate:
  - ➔ Latency
  - ➔ Power consumption (40-50%)

E.g.: 100 System-Components, 64bit-Bus

Wanted is minimum total Wire-Lengths:

- Point-to-Point-Connections ..... ~37500 m
- Multi-Bus-Systems..... ~3000 m
- Network ..... ~100 m

- Point-to-Point-Connections
- Bus-Systems (Shared Medium)

} ➔ Not scalable

Memory-Bottleneck

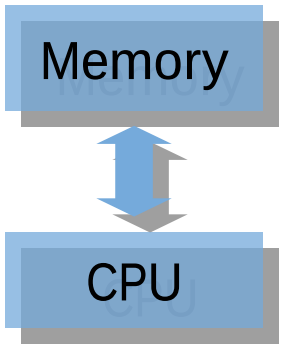

Von Neumann-Architecture

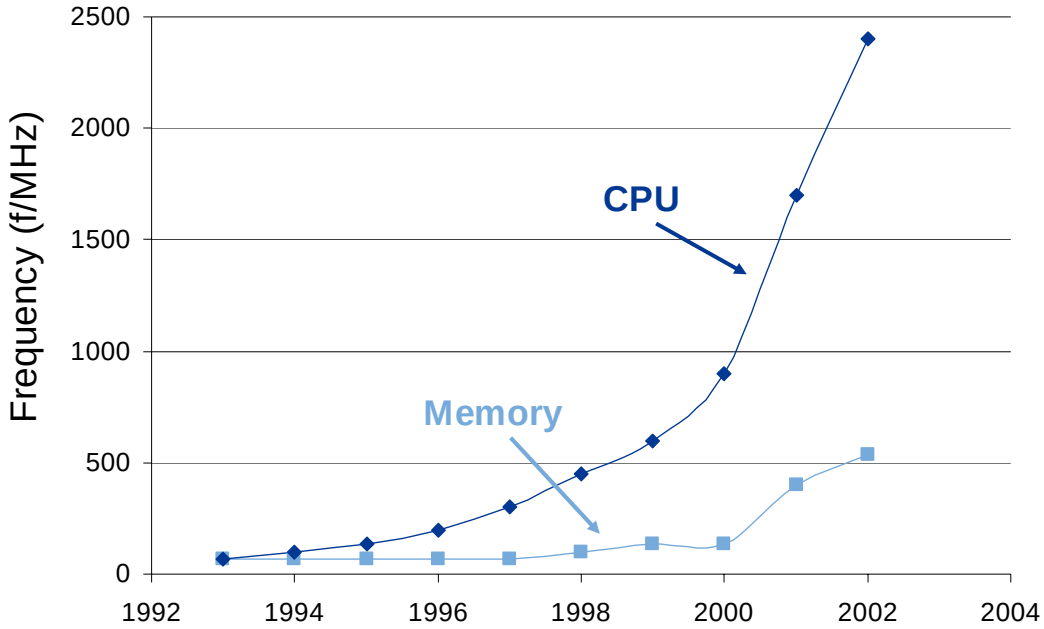

Memory-Bottleneck

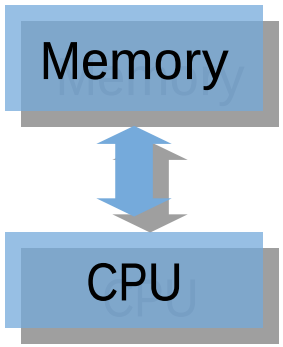

Von Neumann-Architecture

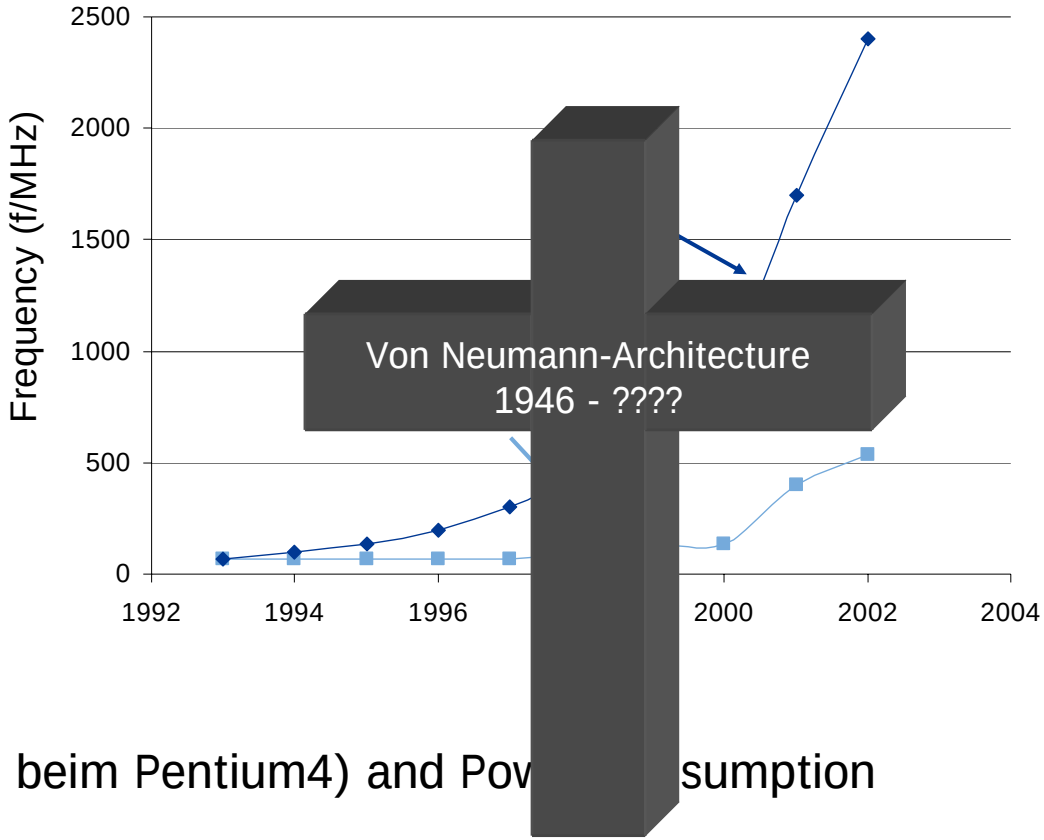

Memory is too slow

- Use of Multi-Level-Cache
  - ➔ Lots of Chip-Area (~50% beim Pentium4) and Pow sumption
- Multi-Threading
  - ➔ Task-Switch-Overhead
- „Real“ Parallelization
  - ➔ Not realizable in Single-Core

Pollack's law

Consequence of the development problems:

„Power outperforms Performance."

Growth rate when new technology is introduced:

	Chip-Area	Power consumption	Performance
General Purpose Processors (GPP)	2X	2-3X	~1.4X

Requirement:

- More MIPS/mm²
- More MIPS/Watt

SoC - Summary

- Hard- and Software-Components
- Computation-oriented
- Extensive traffic between the components
- High System-Complexity
- Large Chip-Area
  - ➔ Productivity (Development, Scalability, Verification)

- Von Neumann-Architecture
- Memory-Bottleneck
  - ➔ Memory/System-Bus limit further performance increase
  - ➔ Cache requires lots of Chip-Area

Has the „old“ System-Paradigm clapped-out?

Targets for a new System-Paradigm

- Von Neumann

- Performance ↑
- Functionality ↑
- System-Complexity ↓
- Power consumption ↓

- Productivity ↑
  - Development
  - Scalability
  - Verification
- Costs ↓

Claas Cornelius, November 2005
Institute of Applied Microelectronics and Computer Engineering, University of Rostock

Parallelism

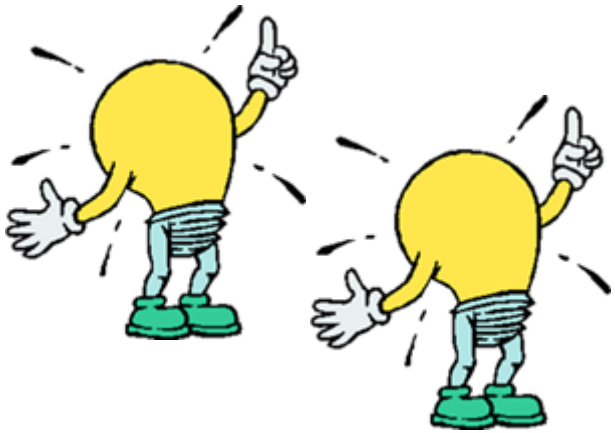

Definition:

What shines brighter than one bulb?

Two bulbs.

„Parallelism is a good possibility for performance increase beyond the Gigahertz“

Paul Otellini, Intels Vice President, 2004

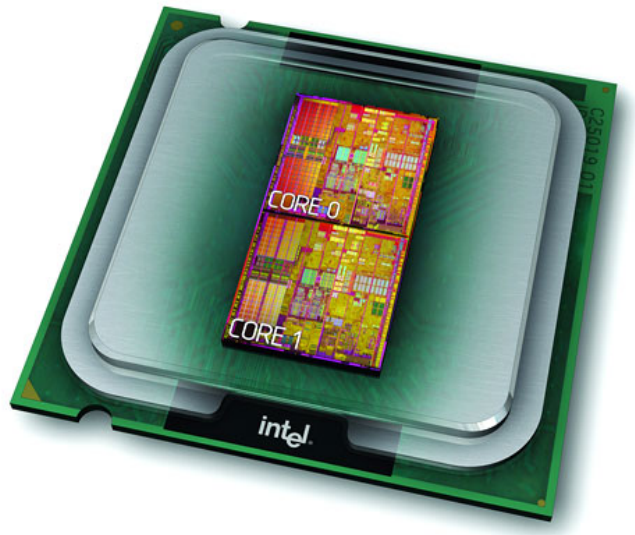

Montecito

- Dual-Core-Chip, 64-Bit Itanium CPU
- 1,72 Milliarden Transistors
- 24 MByte Cache

Modularity

- Independent development of System-Components
- Standardized interface for the connection of Ressources

  - Switches and Links form a Network
  - Encapsulation of the Network and the Ressources

  - Network-Topology
    - Regularity → 2D-Mesh

  - Level of abstraction for System-Developers:
    - Modules
    - Logic-Gates
    - Transistors

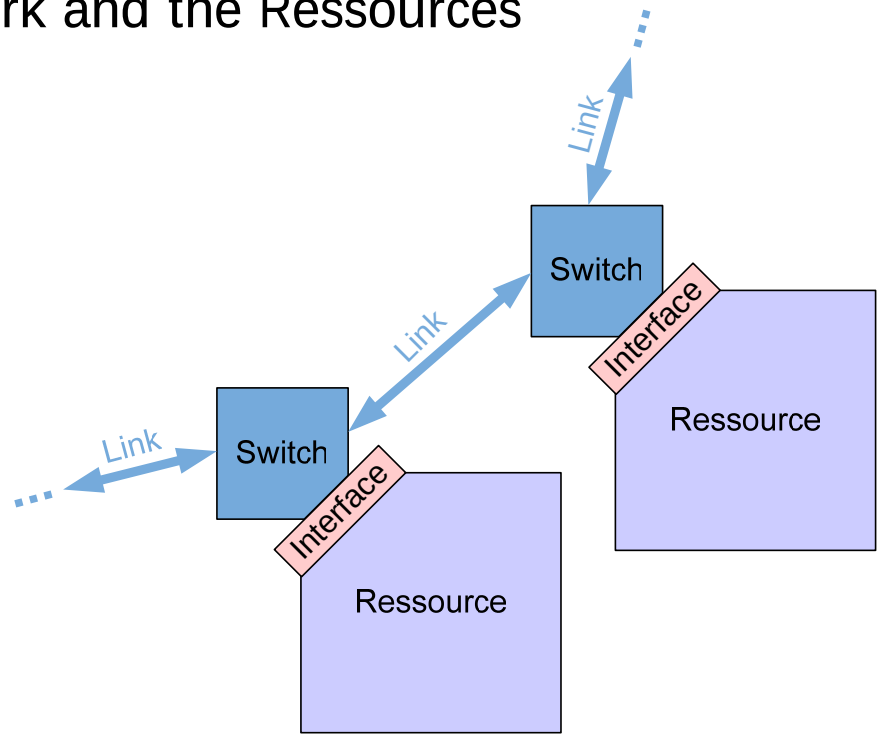

Network-on-Chip

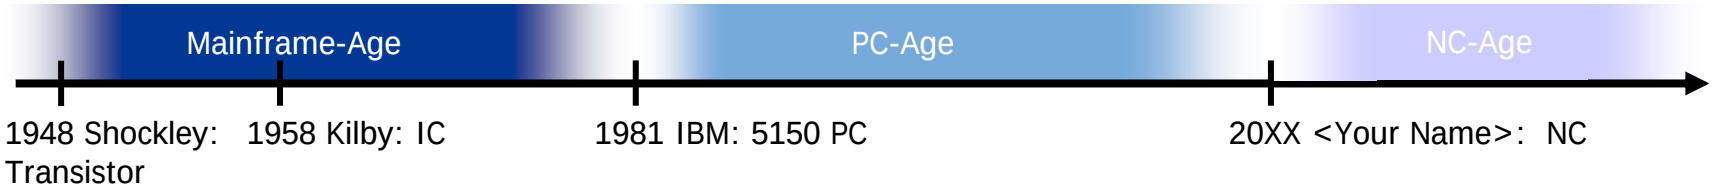

From PC to NoC-Computer (NC)

NoC-Examples:

- Nostrum at KTH, Stockholm
- Æthereal at Philips Research Laboratories, Eindhoven

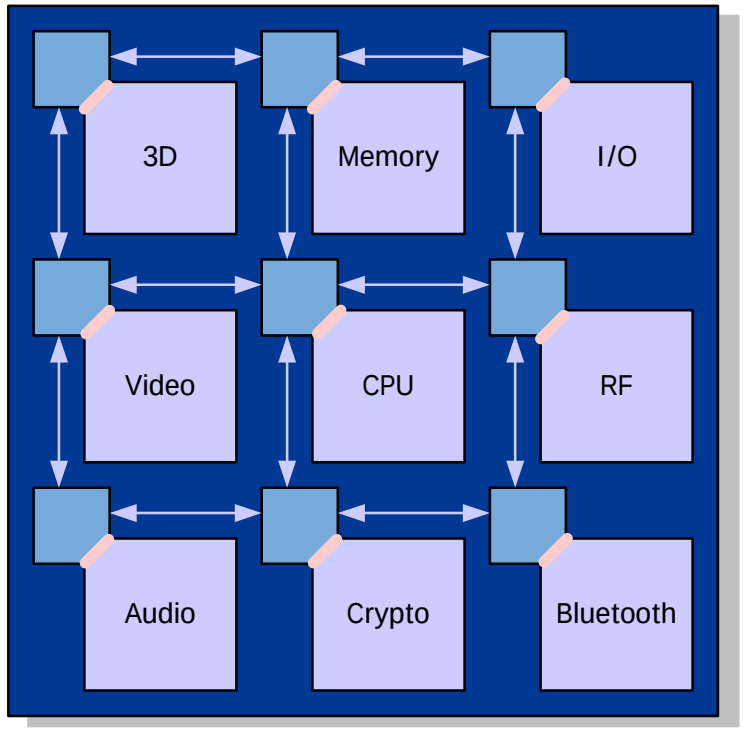

Network-on-Chip

➔ Programmable Communication

➔ Programmable Computation

NoC - Summary

- Hard- and Software-Components
- Communication-oriented
- Modularity of Communication and Computation
- Scalable
- Extensive traffic between the components
- Parallelization/Decentralization
- Large Chip-Area
- Von Neumann-Architecture
  - Memory-Bottleneck remains, but is eased by Parallelization

- Overhead due to Standardization

Outstanding questions for NoCs

- Monitoring and Administering
- Dynamic power management
- Dynamic information-flow management
- Methods for Test and Verification
- Parameter islands, Power-Down, Clock-Gating
- Types of topology
- Routing-Strategies
- Switch-Architectures
- Data transmission/-coding
- Error detection/-correction
- Clock- and Power-Distribution
- Physical Bus-Layout
- Soft-errors (Radiation, EMI)

Summary

Network-on-Chip: Hype or new System-Paradigm?

Change of Paradigms

Computation → Communication

Computing-in-time → Computing-in-time-and-space

- Outline
  - Origin of current System development
  - Problems of Productivity and Architecture
  - Approach for a Network-on-Chip
- Bottom line:
  - NoC offers promising opportunities
  - Many problems are still to solve