A Road to NOMADS - A Perspective on Distributed Computing

Miroslaw Malek
www.informatik.hu-berlin.de/rok
May 12, 2005

Humble Beginnings

From 4000 BC to 1200 BC :
Sumerians used clay tables to register major business transactions (trade records)

The Analogy

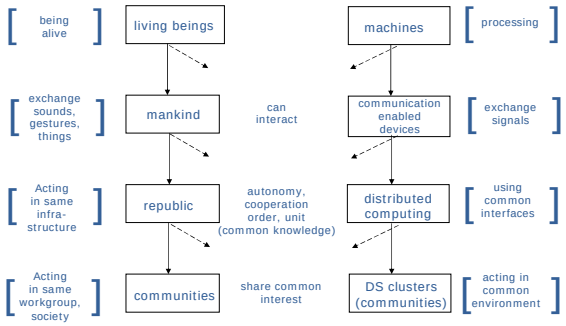

Evolution

- Uniprocessor
- Multiprocessors (parallel systems)
- Multicomputers (distributed systems)

Uniprocessor - Differential Engine

Charles Babbage: The engine could solve sixth degree polynomials with 20 digit precision

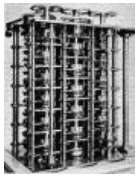

Babbage received £ 17.000 to build the machine but the development was stopped in 1842 and never finished.

Analytic Engine

Babbage contemporary, L. G. Menebrea, developed a hypothetical program to solve a set of linear equations and Ada Augusta Byron, later Lady of Lovelace programmed this application.

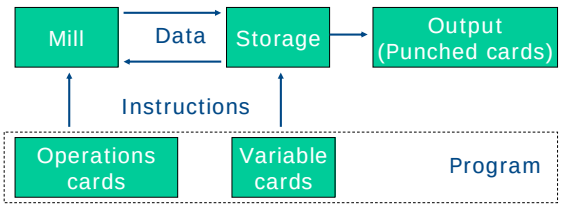

Telephone

1854 Charles Bourseilles (B-F) described how to transfer voice over wires

1857 Antonio Meucci constructs “Teletrofono“ and three years later demonstrates it to the public

1961 Philip Reis builds a device for transferring sounds

1871-74 Meucci reserves a right to file a patent

1876 Alexander Graham Bell files a patent and…

Bell and Western Union

tries, as Meucci did, to sell it to Western Union. The Commission of Experts turns his proposal down with the following explanation:

Technically, we do not see that this device will be ever capable of sending recognizable speech over a distance of several miles. Bell wants to install one of their telephone devices in every city. The idea is idiotic on the face of it. Furthermore, why would any person want to use this ungainly and impractical device when he can send a messenger to the telegraph office and have a clear written message sent to any large city in the United States?

MEMEX

- 1930–45 Vannevar Bush, MIT Professor
- MEMEX: theoretical computer described in „As we may think“
- MEMEX = memory extender electronically linked to a library and able to display books and films
- First concept of hypertext/ -media

- 50‘s IBM‘s concept of distribution and ease of use:
  computing = electricity

First Remote-Access to a Computer

- 1939: Bell Telephone Labs: full-scale electromagnetic relay calculator for solving equations with complex numbers: „Complex Number Calculator“ (later: „Bell Labs Model 1“)
- 1940: the next version was used remotely over telephone lines, and so has been the first „Server“
- The essence of DS:

computers and communication

ILLIAC Computer - 1952

- 5 K main memory and 64 K drum storage
- Weight 5 tons, 2800 vacuum tubes

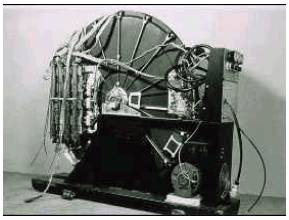

Technology Push

PROCESSOR	PEAK PERFORMANCE	YEAR
N10 (i860)	40 MFLOPS	1990
N11 (i870)	100 MFLOPS	1991
ALPHA	200 MFLOPS	1993
Pentium, Power PC	300 MFLOPS	1997
Pentium Pro	500 MFLOPS	1999
Itanium	10 GFLOPS	2006

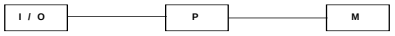

Parallel Computing

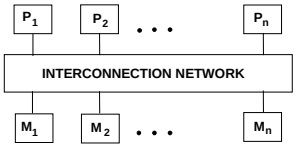

50's Burroughs 4x4 Parallel Computer

60's ILLIAC IV

70's Texas Reconfigurable Array Computer

80's Research Parallel Processor Project (RP3)

Connection Machine, Butterfly Monarch, GF-11, SP-2

Genesis,Touchstone, n-cube, GRID, Sequent

Burroughs Parallel Computer

4x4 crossbar-switched
shared memory system

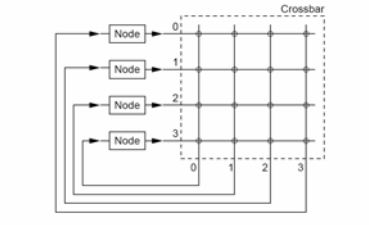

ILLIAC IV, 1965-1972

- 64 processor system on grid with neighbor shared memory,
- Performance: 200 MIPS, 15 MFLOPS (expected 1000MFLOPS)
- Discontinued in 1972 (about a quarter was actually built) as the cost reached $31M

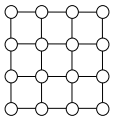

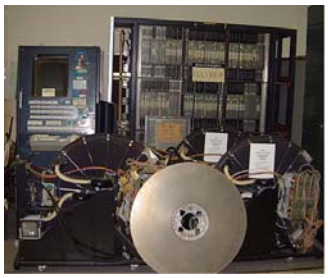

TRAC - Main Issues

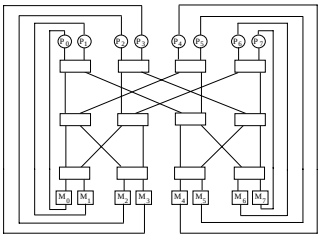

Texas Reconfigurable Array Computer

Multistage Interconnection Networks

Circuit-switching vs packet switching

Shared memory vs private memory

Operational mode (Flynn's Taxonomy):
SISD
SIMD
MISD
MIMD

Testing

Finite State Machine

Finite State Machine of the ALB of a 4 x 4 switch using round-robin priority

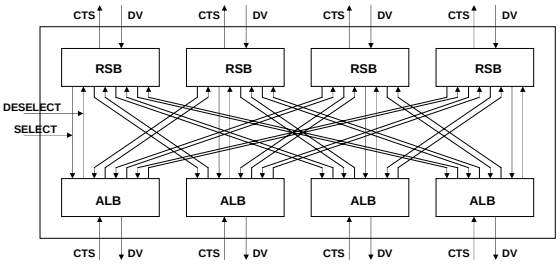

ESTIMATION OF TESTING TIME FOR THE SIGMA-1 COMPUTER

- SIGMA-1 interconnection network:
- L = 2 two levels
- 10 x 10 switches configured as 8 x 8's
- Round-robin priority
- Time assumed for traversal of the network and memory access: t = 120 ns
- Estimated testing time using a pseudoexhaustive method: ~20 hours
- Actual testing time: ~22.5 hours
- Estimated testing time using our method: 26 seconds

OVER 3000 TIMES BETTER!

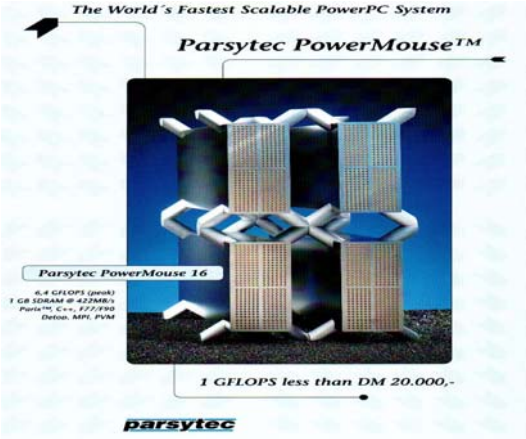

Heterogeneous Multicomputing

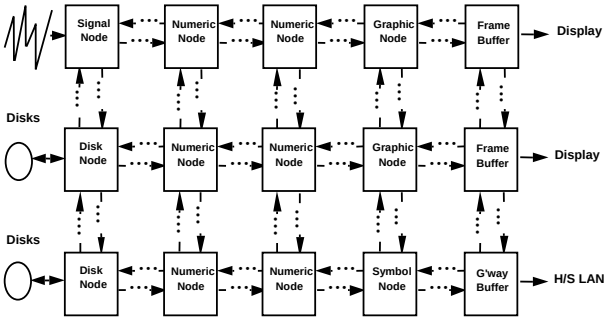

Multicomputers – distributed systems

A **distributed system** is a set of connected independent entities with a processing/storage/communication (computers) capability that appears to the users as a single system.

Synchronization

- The essence of cooperation
- Synchronization based on the actual time (clock sychronization)
- Synchronization based on relative ordering (logical synchronization)
- Snapshot – recording a state

Naming, Addressing and Routing

- A key to performance and scalability
- Name identifies an entity (be it computer, file)
- The name of an access point to an entity is called an address
- A plethora of protocols

Network

- Paul Baran of RAND develops the idea of distributed, packet-switching networks
- decentralized network linking computers which communicate using small packages of data
- searching the best possible route in the network

Ethernet

- 1973: Local Area Networks
  - Ethernet, Xerox PARC by Robert Metcalfe
  - Token-passing ring, IBM and Sweden
- 1973: Robert Kahn and Vinton Cerf develop the basic ideas of the Internet (*connectivity, distribution, black box design, error recovery*)

ARPANET, Internet - Timeline

- 1969: ARPANET connecting University of Utah, Stanford Research Institute, UCLA and UCSB
- 1982: TCP/IP (Transmission Control Protocol and Internet Protocol) is established as the standard for ARPANET
- 1984: ARPANET was divided into two networks: ARPANET and MILNET :
  - ARPANET to support the advanced research component
  - MILNET was to serve the needs of the military
- 1992: ARPANET becomes de facto Internet
- 1994: The main U.S. Internet backbone traffic begins routing through commercial providers

WWW

- 1987: The number of network hosts breaks 10,000, then 100,000 in 1989. It reaches 1,000,000 in 1992, 100 M in 2000 and 350 M in 2005
- 1991: Tim Berners-Lee develops the World Wide Web (URL, HTML, HTTP)
- CERN releases the first Web server

Evolution of Computing

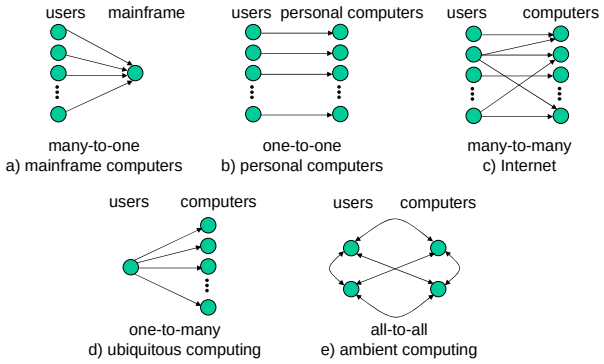

The NOMADS Republic

- The Status:The largest nation on Earth
- Population: 20 - 100 B citizens, maybe 1T
- Key qualities: Mobility, Adaptivity & Dependability

The Need

Need for a unifying paradigm that covers:

- embedded systems
- sensor networks
- Personal computing
- server farms
- GRID computing

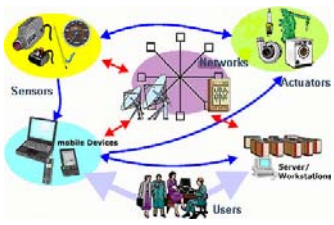

The Goal: Taming the Chaos

- Interacting entities
- Examples:
  - *Physics*: Molecules, atoms, planets, galaxies
  - *Biology*: cells, organisms
  - *Society*: dictatorship, republic, anarchy
- Challenge: „orchestration of the communities“

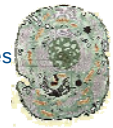

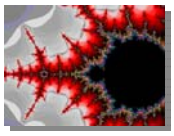

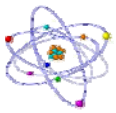

Ambient Computing

- Humans in charge
- Societal model – computer networks grow along humans
- Comfort level not disturbed (e-mail today, services overrun tomorrow)
- Service and user-friendliness-oriented
- Ubiquitous but not overbearing
- Non-invasive
- Anticipating user functions, tolerating mistakes
- Supporting MAD properties

Organization of the Republic

Organization of the NOMADS Republic:

- citizens
- laws
- economy
- social structure
- services

Through the Eyes of Our Experience

- Consensus, leader election, consistency:
  - Unstoppable orchestra, robots
  - Fault prediction leading to reconfiguration
- Communication
  - Ad hoc routing
  - Remote experiment
- Resource allocation
  - Dynamic scheduling
- Composability
  - Real-time
  - Service

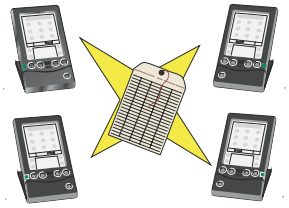

Magic Map – Location-Based Services

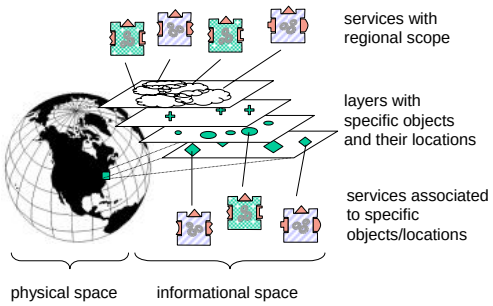

Campus Berlin-Adlershof – A Case Study

Wireless networks:

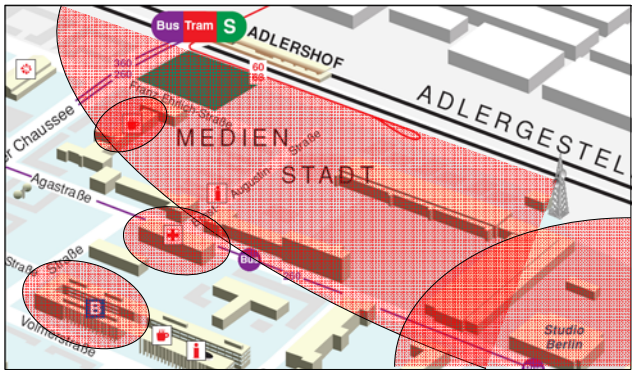

RFID chips:

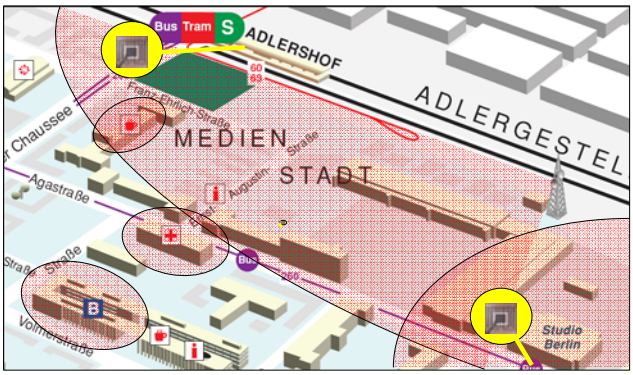

Content providers:

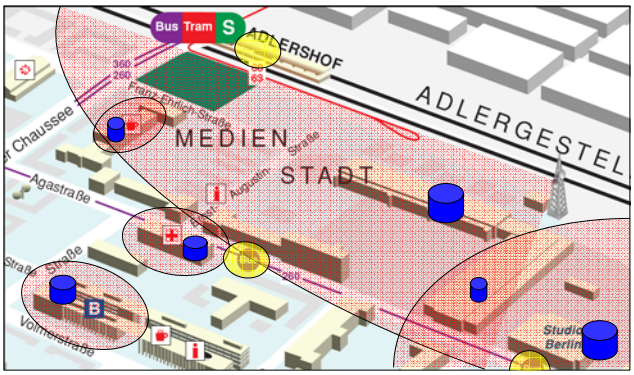

GPS satellites:

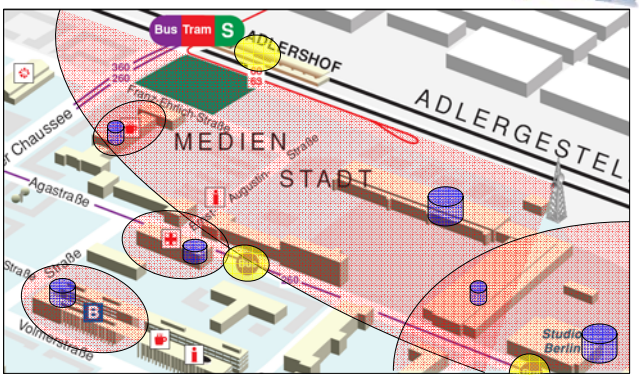

Mapping services to real world objects:

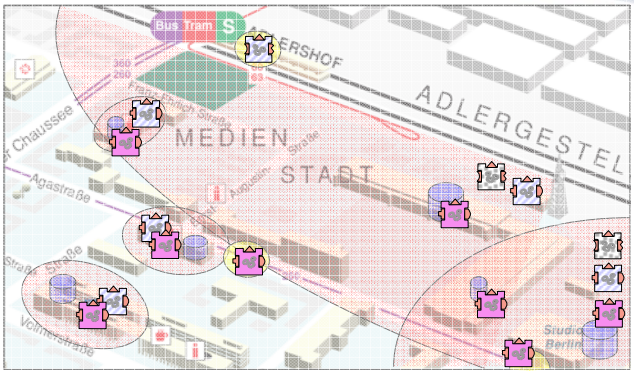

Adaptive Service Composition:

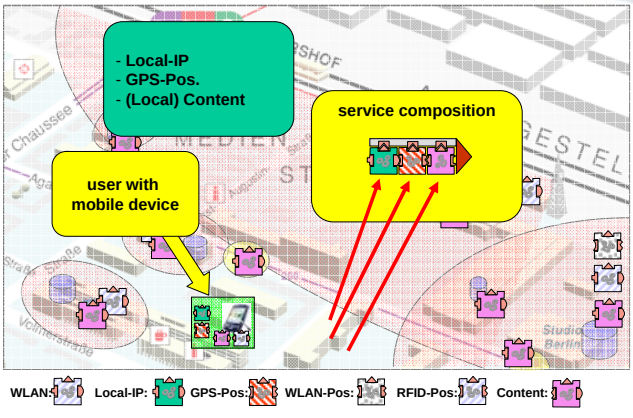

Adaptive Service Composition:

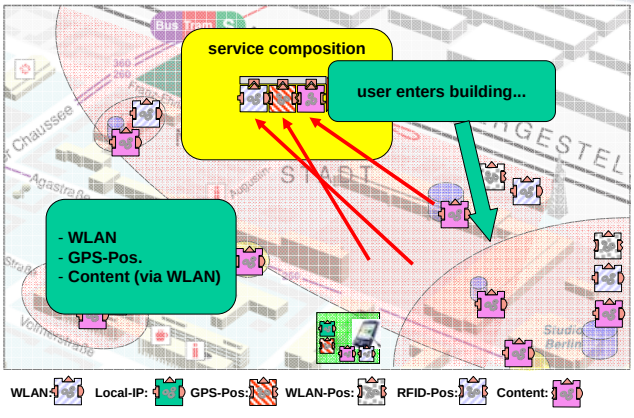

Inside the building:

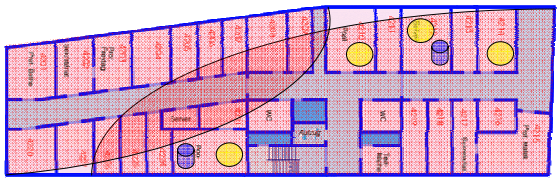

WLAN: RFID: Content: GPS:

Services inside the building:

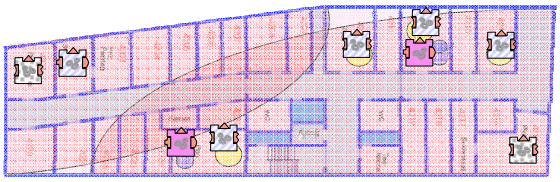

WLAN: Local-IP: WLAN-Pos: RFID-Pos: Content:

Adaptive service composition:

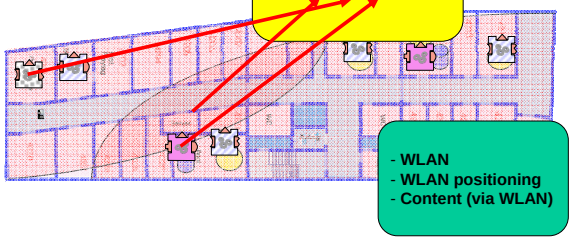

WLAN: Local-IP: WLAN-Pos: RFID-Pos: Content:

Adaptive service composition:

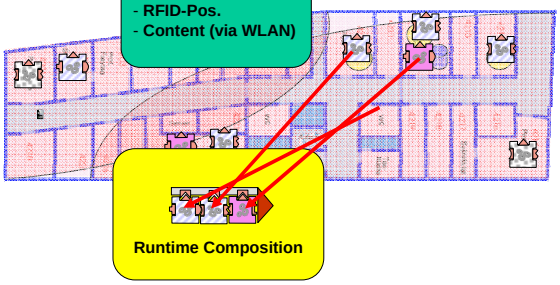

WLAN: Local-IP: WLAN-Pos: RFID-Pos: Content:

Services

- Hiding complexity behind services: “everything” is a service
- W-questions:
  - Who are you?
  - Where are you?
  - What do you offer?
- Services expose extended interfaces:
  - functional properties
  - non-functional properties
  - semantics
- interoperation between systems
- composable architectures
- composition, decomposition, adaptation
- Existing Approaches:
  - WSCI (Web Service Choreography Interface)
  - WSFL (Web Service Flow Language)
  - BPEL (Business Process Execution Language)

Trends

Zero-Mania: low power, time-to-market, downtime, etc.

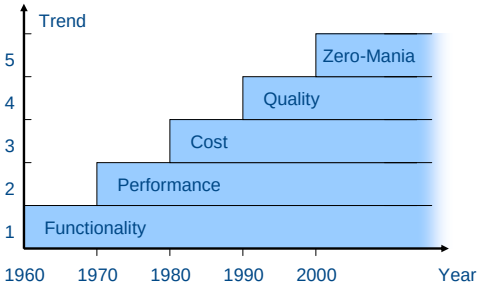

Then, computers became a rarity…

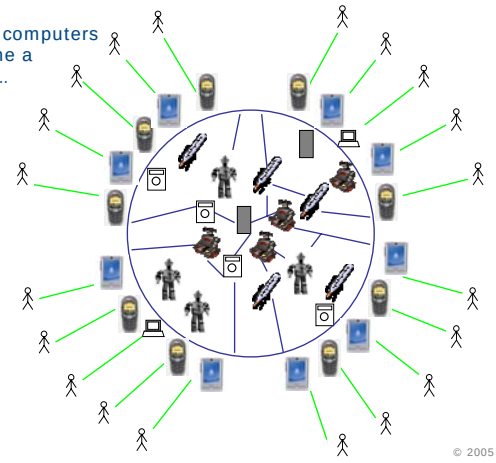

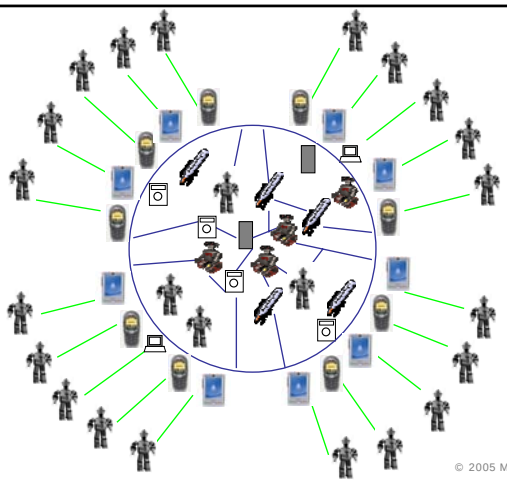

Conclusion

- A viable organization concept for the emerging largest nation on earth is an imperative for this decade
- Convergence of computing and communication will continue
- An infrastructure for service architectures is evolving
- Humans will regain the power of using computers and ambient computing might become a reality

We‘ve only just begun…

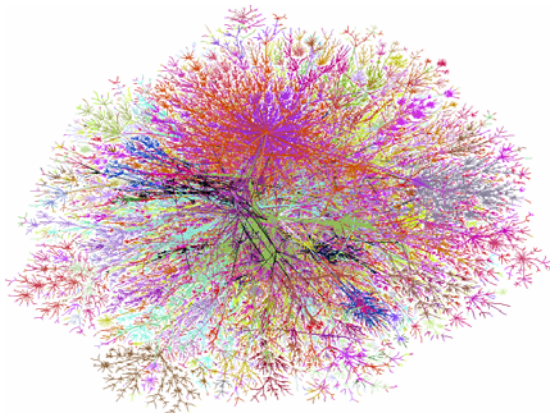

Acknowledgments:

- IEEE – Photos on pages 5, 9, 10, 11, 15, 24, 25 and 27
- Parsytec GmbH – p. 19
- Unknown Source – p. 53