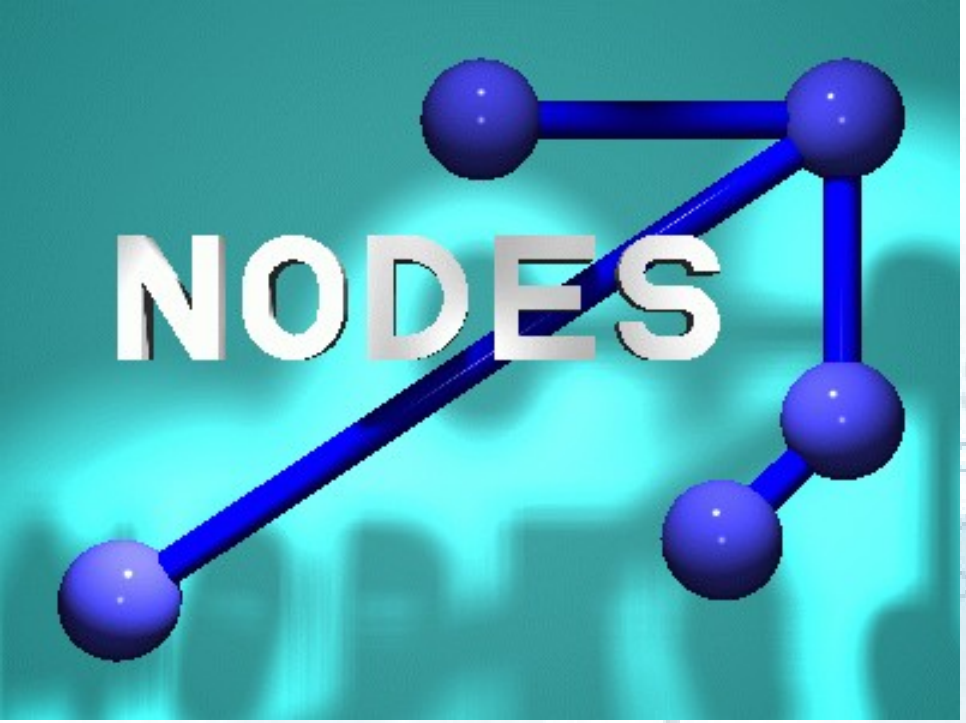

Benchmarking Telecommunication Systems

Kimmo Raatikainen

Helsinki University Computer Science Department

kimmo.raatikainen@cs.helsinki.fi

http://www.cs.helsinki.fi/Kimmo.Raatikainen/

Presentation Outline

- Fundamentals of Benchmarking
- Networking Benchmarks
- Other Benchmarks
- Summary

Benchmarking Challenge

- Benchmarking is about measuring performance under a test suite
- In order to get adequate information, benchmark must be:
  - ***Relevant***: the measurements must reflect tylical operations in the selected problem domain
  - ***Portable***: the benchmark should be easy to implement in different developments environments, system architecturers, etc.
  - ***Scalable***: the same benchmark should apply to small and large systems alike, including different flavors of parallel and distributed architectures
  - ***Simple***: the benchmark and its results should be easy to understand

Rat Holes in Performance Analysis

1. No Goals
2. Biased Goals
3. Unsystematic Approach
4. Analysis without Understanding the Problem
5. Incorrect Performance Metrics
6. Unrepresentative Workload
7. Wrong Evaluation Technique
8. Overlooking Important Parameters
9. Ignoring Significant Factors
10. Inappropriate Experimental Design
11. Inappropriate Level of Detail
12. No Analysis
13. Erroneous AnalysisNo Sensitivity Analysis
14. Ignoring Errors in Input
15. Improper Treatment of Outliers
16. Assuming No Change in the Future
17. Ignoring Variability
18. Too Complex Analysis
19. Improper Presentation of Results
20. Ignoring Social Aspects
21. Omitting Assumptions and Limitations

Summarizing Benchmark Results

- All text books give a warning about using single number to characterizing a computing system performance
- However, single number is what most people want
- You have two options:
  - Measure a single index
  - Calculate a summary based on multiple measured indices
- The first option is naïve
- The second option gives you all the freedom to obtain the ordering you want
  - However, you can be honest and not to use the possibility.
  - You should, before any measurement, decided how you calculate the summary (and explain why you selected the formula)
- There are interesting articles that discuss using weighted arithmetic mean, geometric mean, harmonic mean, ratios

World is full of Benchmark Suites

- Google got in January 2006 2.9 million hits on “computer performance benchmark”
- IEEE Xplore finds 426 articles published since 1.1.2000 in that category
- ACM Digital Library has 11,519 entries in that category
- I have found about 120 different benchmarks that has a name
  - Some on them are good, some of them are terrible
-
- ***The problem is not to find benchmarks but to select suitable ones!***

UMTS Reference Architecture

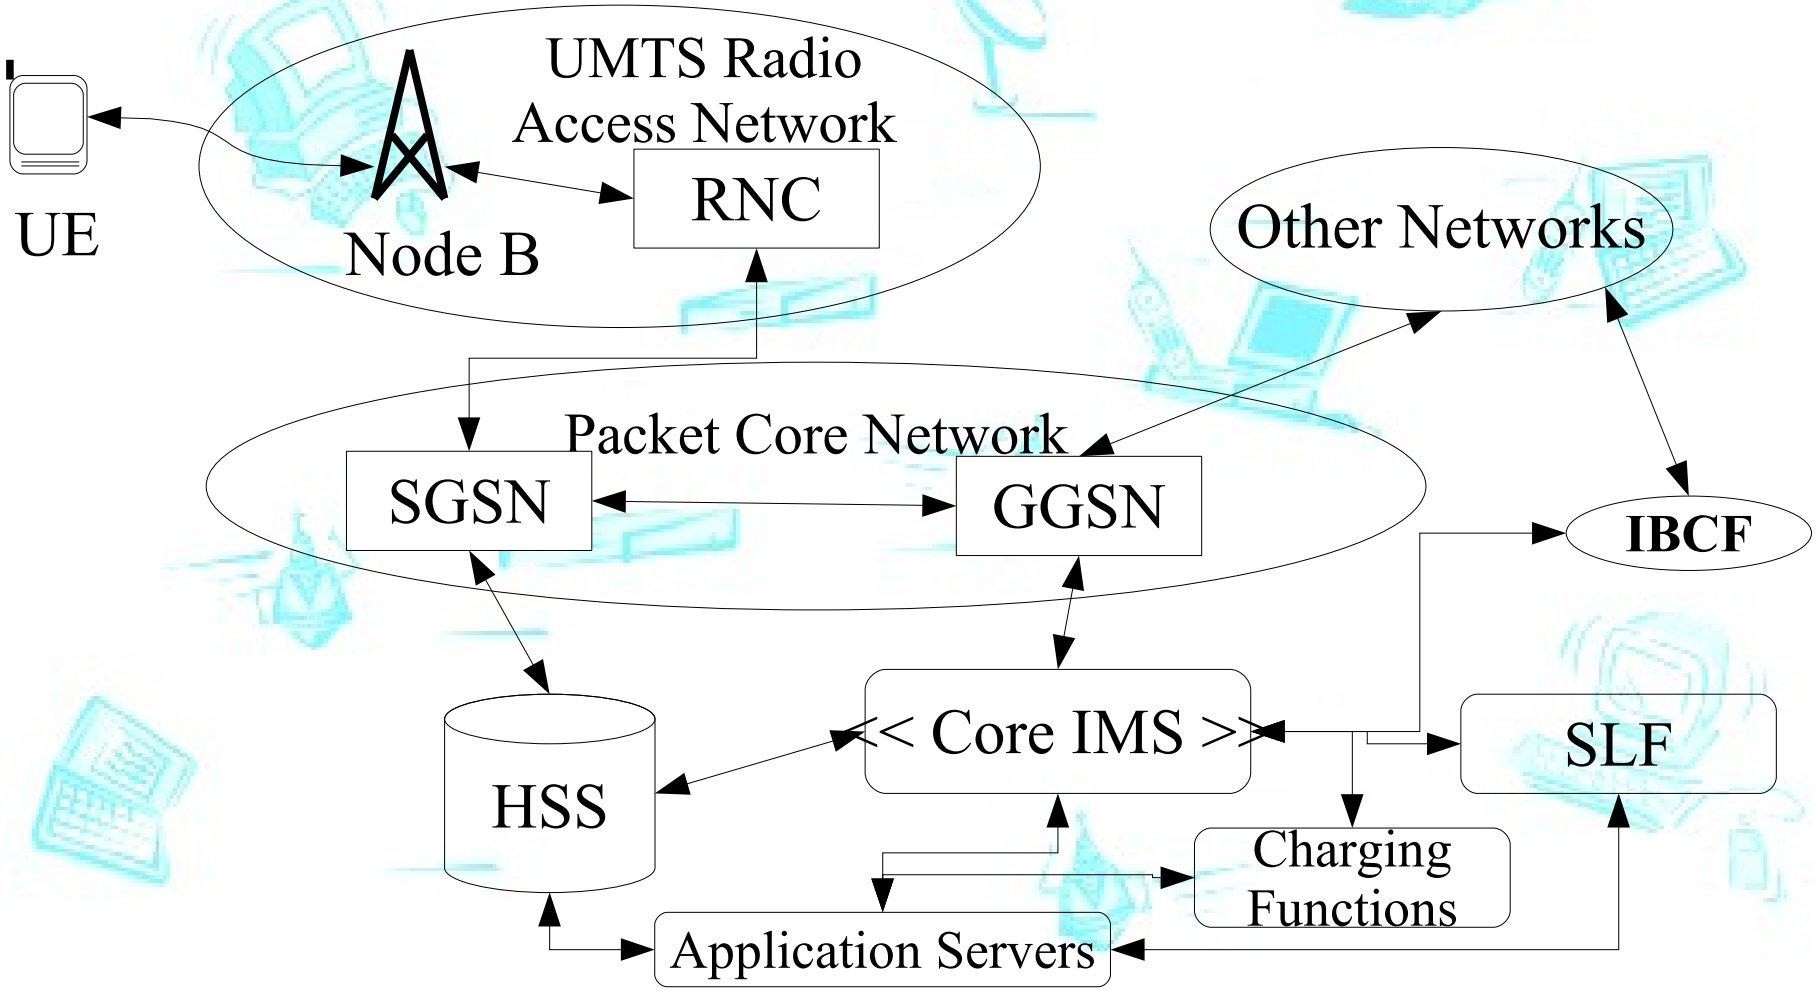

Selection Framework

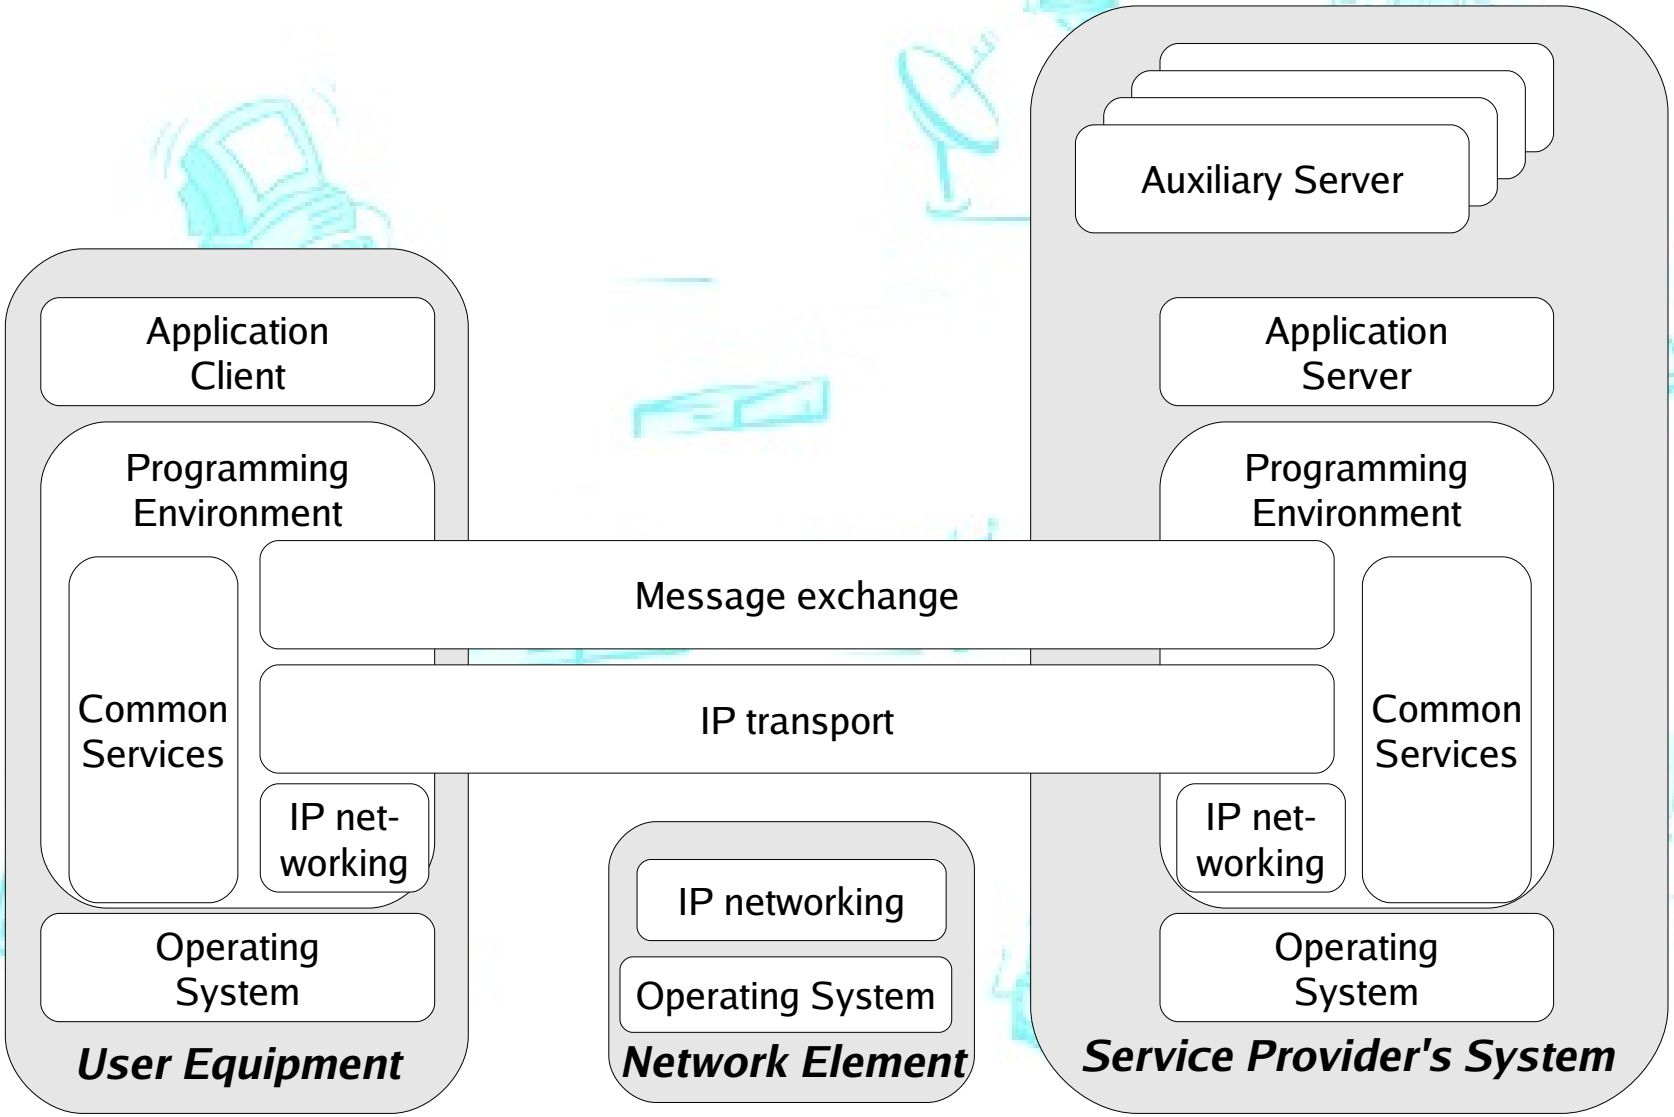

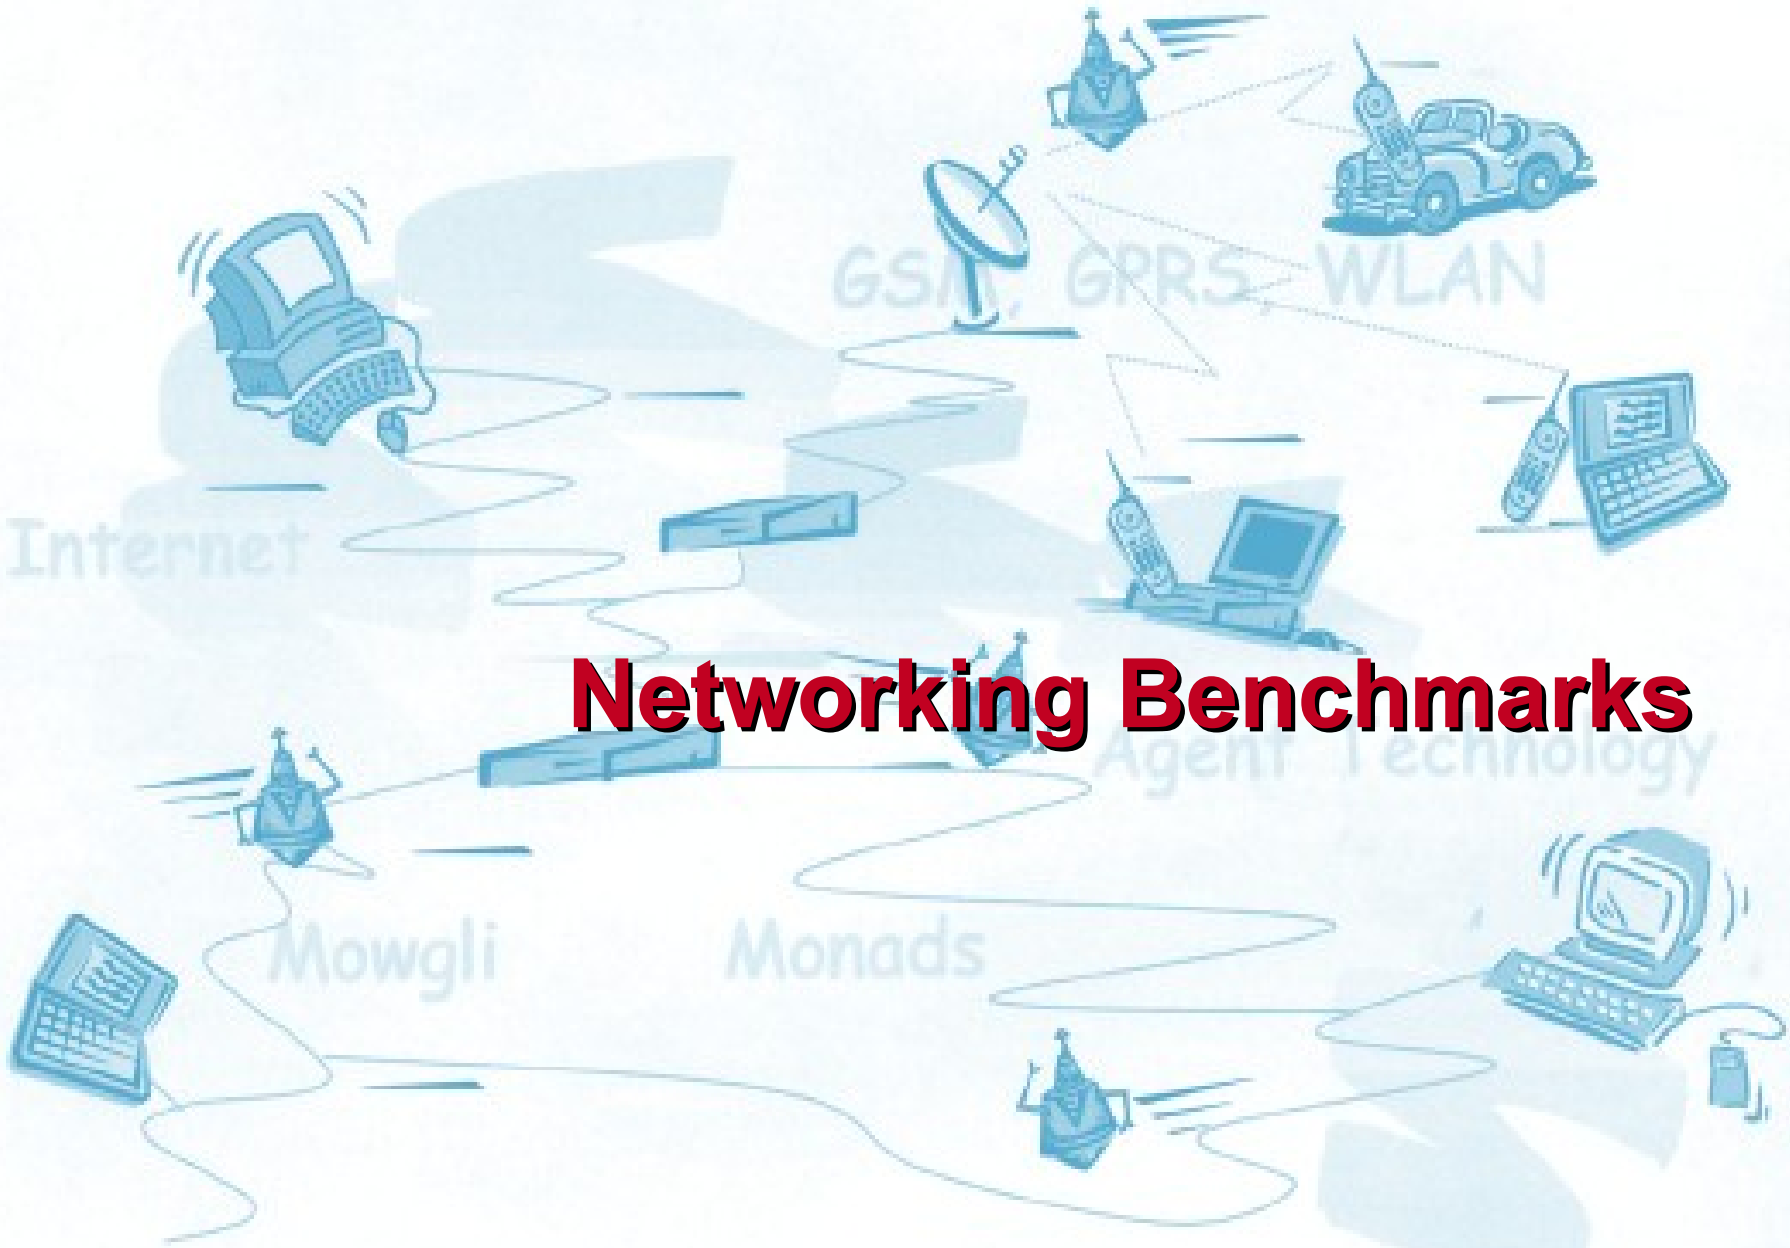

Networking Benchmarks

Networking Benchmarks

- CommBench
- NpBench
- BTU or Bits To User
- NetBench
- SIPstone
- Benchmarking Implementation Agreements by Network Processing Forum
- IMS/NGN Performance Benchmark by ETSI TISPAN
- CplaneBench
-

CommBench - 1/2

- Header processing kernels:
  - lookup operations on tree data structure based on a radix-tree;
  - packet header modification and checksum computation based on an application called FRAG;
  - queue maintenance and packet scheduling for fair resource utilization based on deficit round-robin fair scheduling algorithm; and
  - pattern matching on header data fields based on the tcpdump application.

CommBench - 2/2

- Payload processing kernels:
  - encryption arithmetic based on the CAST-128 block cipher algorithm;
  - data compression based on the Lempel-Ziv algorithm as implemented in ZIP;
  - redundancy coding using the Reed-Solomon FEC; and
  - DCT and Huffmann coding, based on JPEG code.

NpBench

- Traffic-management and QoS group
  - WFQ algorithm,
  - RED algorithm,
  - SSL dispatcher, and
  - multi-protocol label switching (MPLS)
- Security and media group
  - media transcoding,
  - AES,
  - MD5, and
  - Diffie-Hellman key exchange
- Packet processing group
  - packet fragmenting (FRAG) and
  - CRC calculation.

BTU

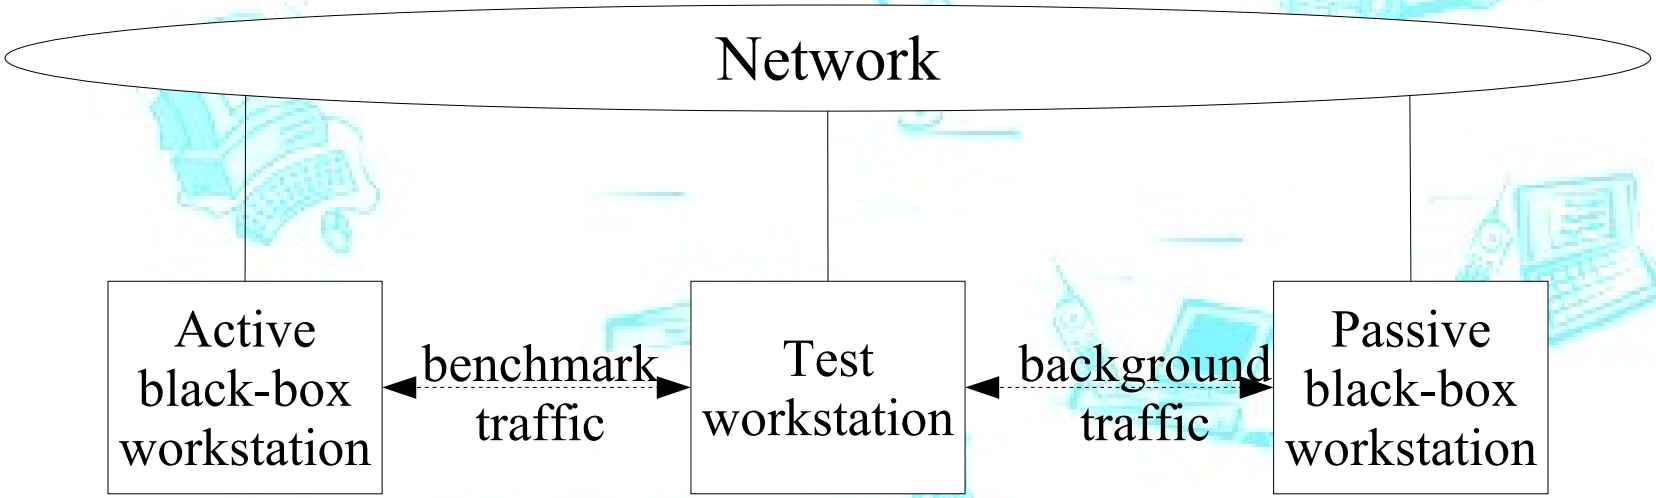

- host activity is CPU, network, or I/O intensive
- benchmarked communication represents different application classes:
  - Telnet, FTP, Audio, Netscape, and Video

NetBench (ACM TOECS) - 1/3

- three-level categorization:
- low- or ***micro-leve***l routines containing operations nearest to the link or operations that are part of more complex tasks;
- ***routing-level*** applications, which are similar to traditional IP level routing and related tasks; and
- ***application-level*** programs, which have to parse the packet header and sometimes a portion of the payload and make intelligent decisions about the destination of the packet.

NetBench - 2/3

- Micro-Level Programs
  - CRC: The CRC-32 checksum calculation (ISO 3309)
  - TL: TL is the table lookup routine (radix-tree routing table)

- Routing-Level Programs
  - make a decision depending on the source or destination IP address of the packet.
  - ROUTE: IPv4 routing (RFC 2644)
  - DRR: Deficit-round robin scheduling
  - IPCHAINS: a firewall application
  - NAT: Network Address Translation

NetBench - 3/3

- Application-Level Programs
  - DH: Diffie-Hellman public key encryption/decryption
  - MD5: Message Digest algorithm (MD5) creates a signature for each outgoing packet
  - SNORT: an open-source network intrusion detection system
  - SSL or Secure Sockets Layer is the secure transmission package
  - URL: implements URL-based destination switching

SIPstone - 1/3

- benchmarks for SIP proxy, redirect, and registrar servers
- to measure the request handling capacity of SIP servers

- The benchmark environment:
  - SUT: SIP proxy, redirect or registrar server,
  - load generators (in essence, SIP user-agent clients) that generate the requests
  - call handlers, which simulate user-agent servers, and
  - a central benchmark manager, which coordinates the execution of the benchmark

SIPstone - 2/3

- Registration: the load generator sends REGISTER messages using digest authentication to the SUT.
- Outbound proxy: The load generator sends INVITE requests to the SUT acting as an outbound proxy.
- Redirect: The load generator sends INVITE requests to the SUT acting as a redirect server.
  - The delay from sending an INVITE request to receiving the 3xx response is measured.

SIPstone - 3/3

- Proxy 480: The load generator sends INVITE requests to the SUT acting as a redirect server.
  - the destinations have not registered and the server returns a 480 (temporarily unavailable) response
- Proxy 200: The load generator simulates a call setup and teardown by sending an INVITE and a BYE immediately after the INVITE is completed.

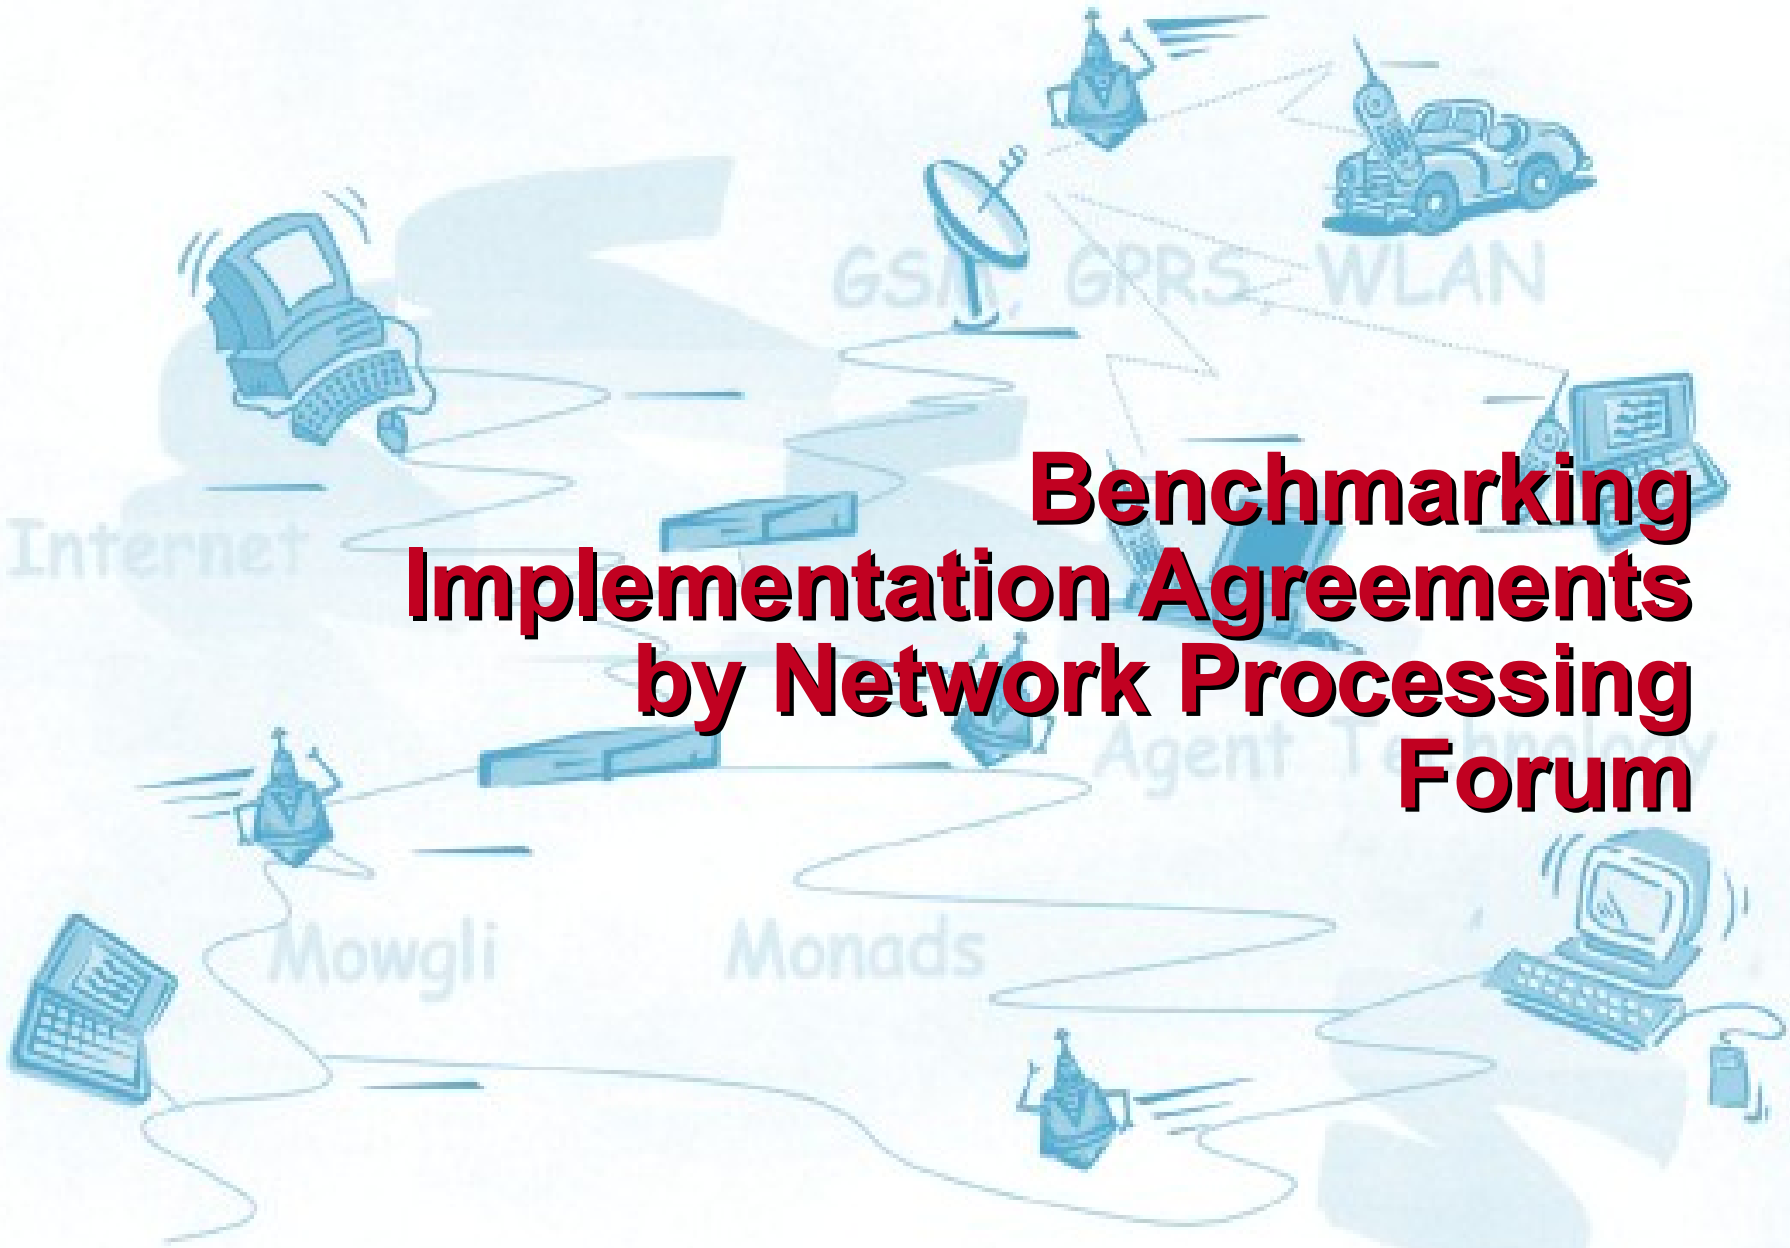

Benchmarking Implementation Agreements by Network Processing Forum

NPF Benchmark Suites

- TCP Proxy Application Level Benchmark
- IPSec Forwarding Application Level Benchmark
- Switch Fabric Benchmark
- IP Forwarding Benchmark
- MPLS Application Level Benchmark
- IPv4 Forwarding Benchmark

TCP Proxy Application Level Benchmark

- TCP Proxy Goodput on existing connections,
- Goodput with connection setup/teardown,
- connection setup rate,
- connectionsetup and teardown rate,
- SYN/ACK latency, and
- connectionsetup latency.

Note: TCP Proxy terminates TCP connections

IPSec Forwarding Application Level Benchmark

- IPSec forwarding rate,
- IPSec throughput, and
- IPSec latency.

IP Forwarding Benchmark - 1/2

- native IPv4, native IPv6 and mixed IPv4/IPv6 traffic
- IPv4 routing tables: 10k, 120k and 1M routes
- IPv6 routing tables: 400 and 1200 routes

- The tests are grouped into three categories:
  - user (data) plane tests,
  - control plane tests, and
  - concurrent user (data) plane and control plane tests

IP Forwarding Benchmark - 2/2

- User plane tests:
  - the aggregate forwarding rate,
  - Throughput,
  - Latency,
  - loss ratio,
  - overload forwarding rate, and
  - system power consumption.
- Control plane tests:
  - forwarding table update rates
- User plane and control plane tests:
  - concurrent forwarding table updates and the forwarding rate

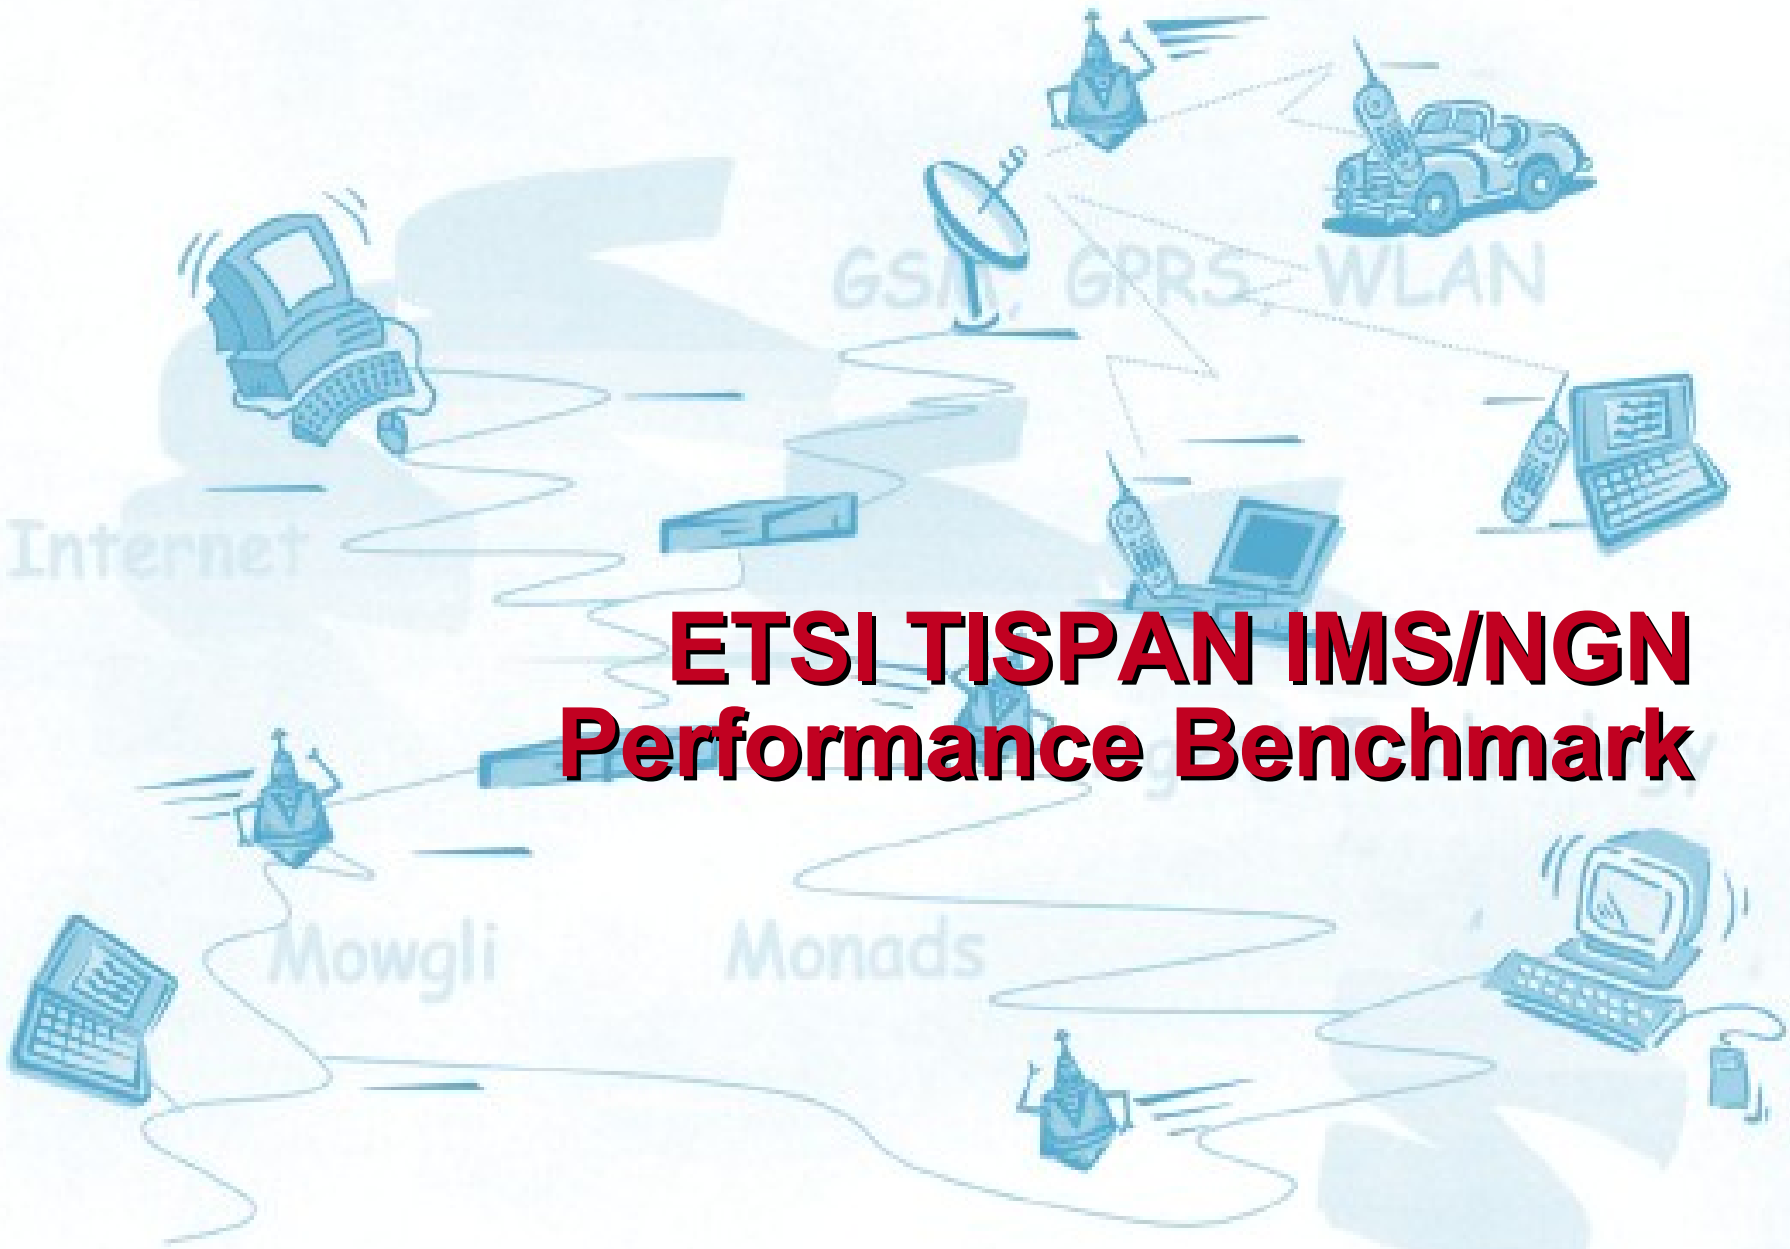

ETSI TISPAN IMS/NGN Performance Benchmark

IMS Benchmark Information Model

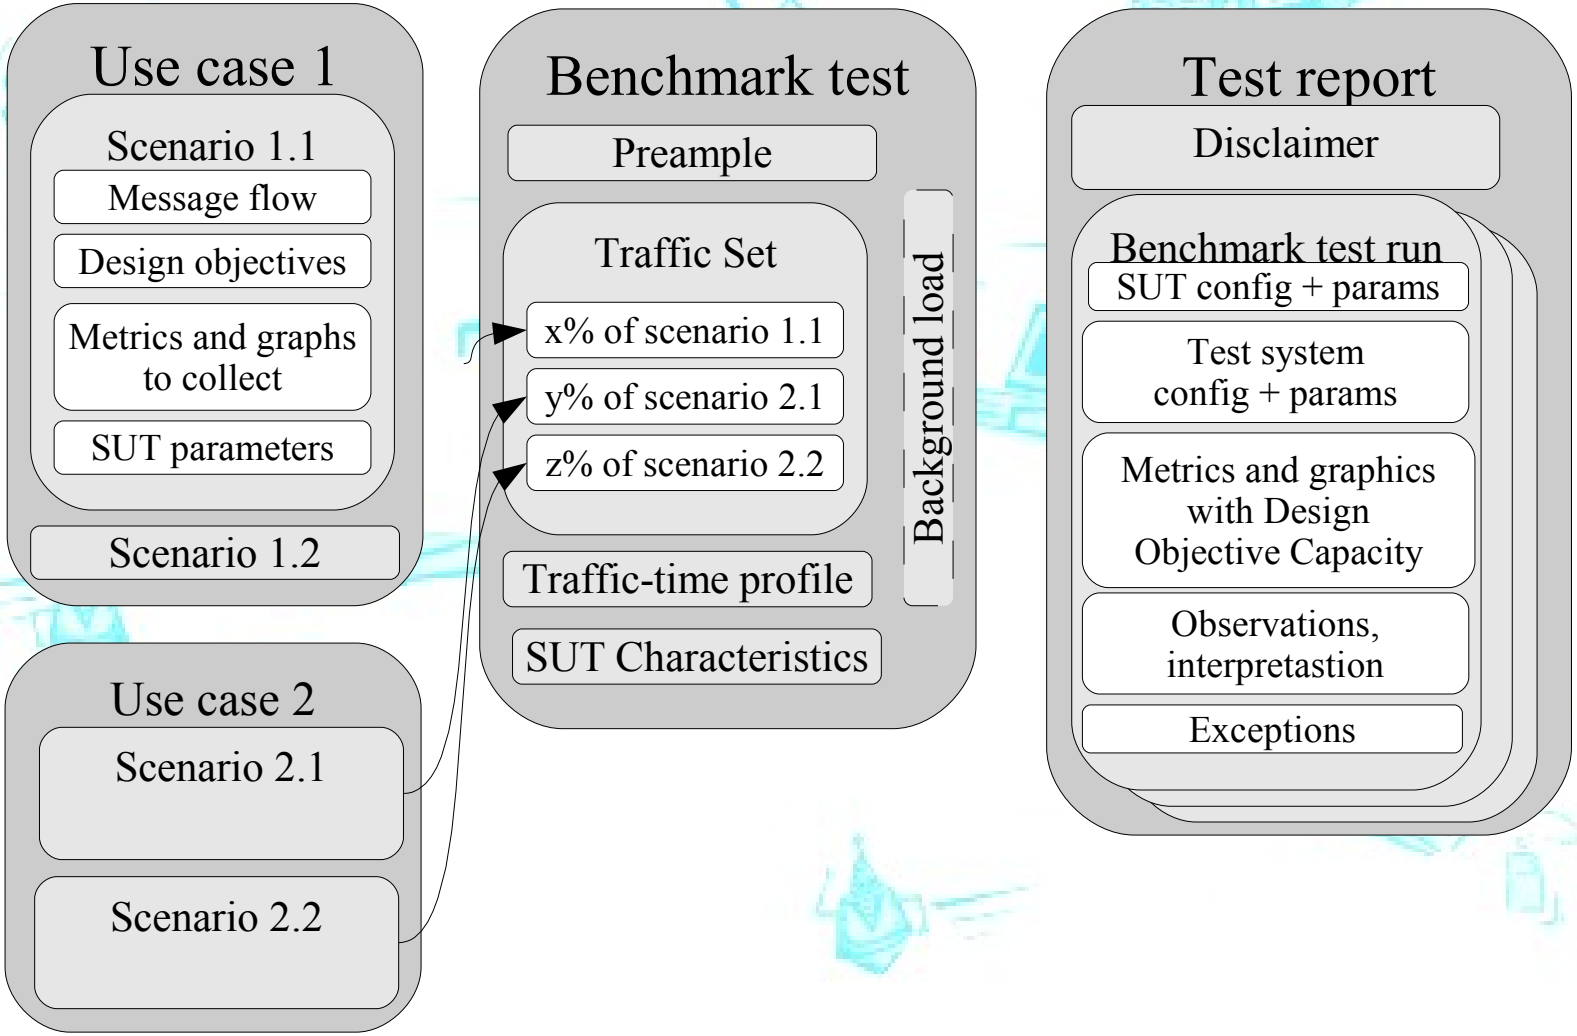

IMS Reference Architecture

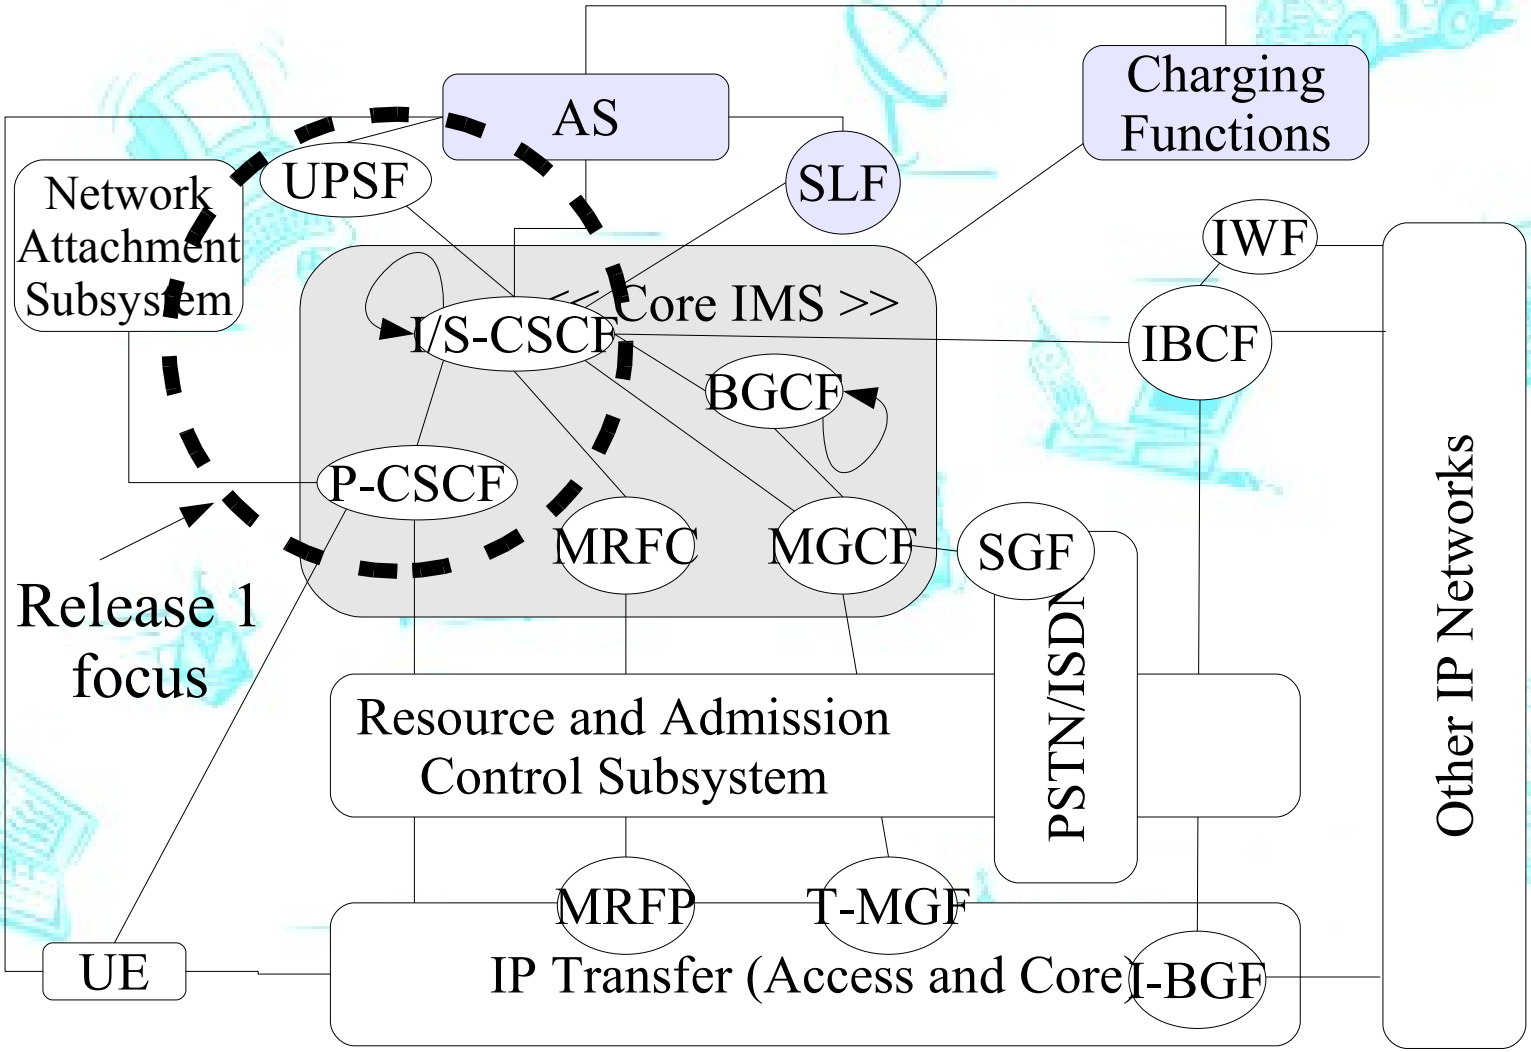

SUT Topologies

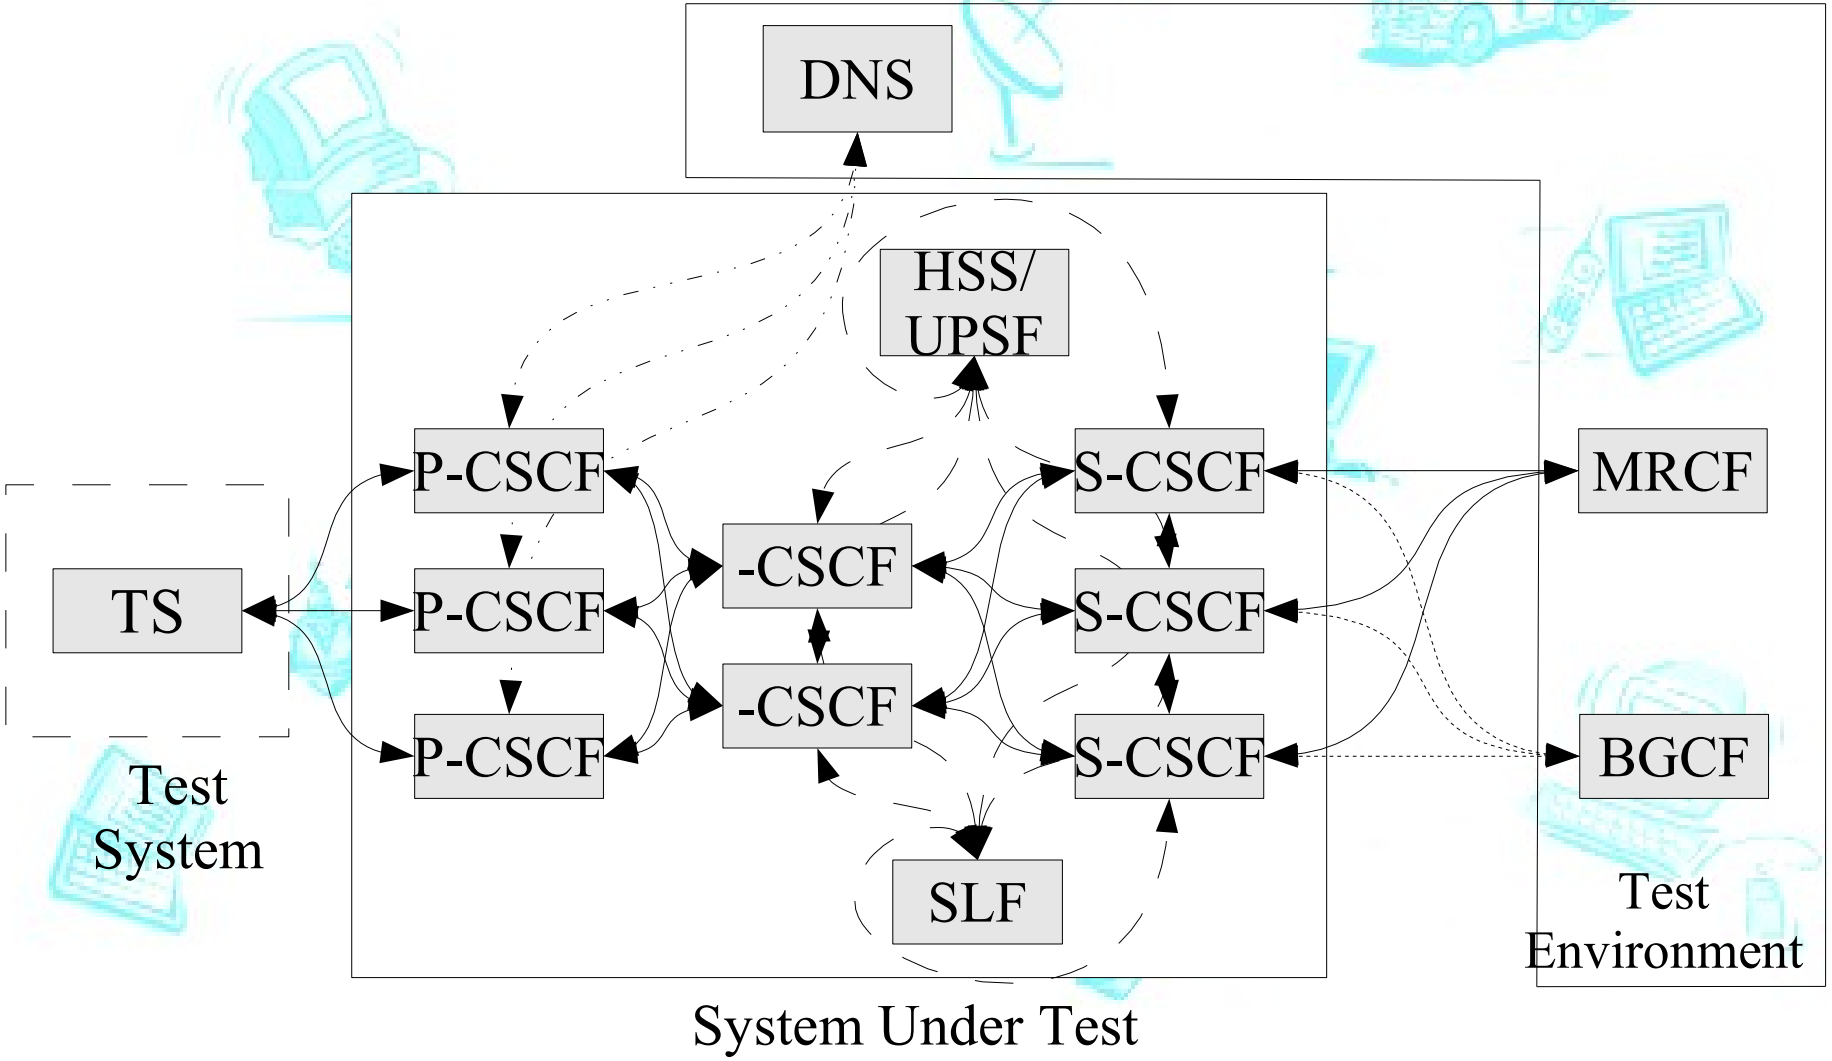

Test System and SUT Interactions

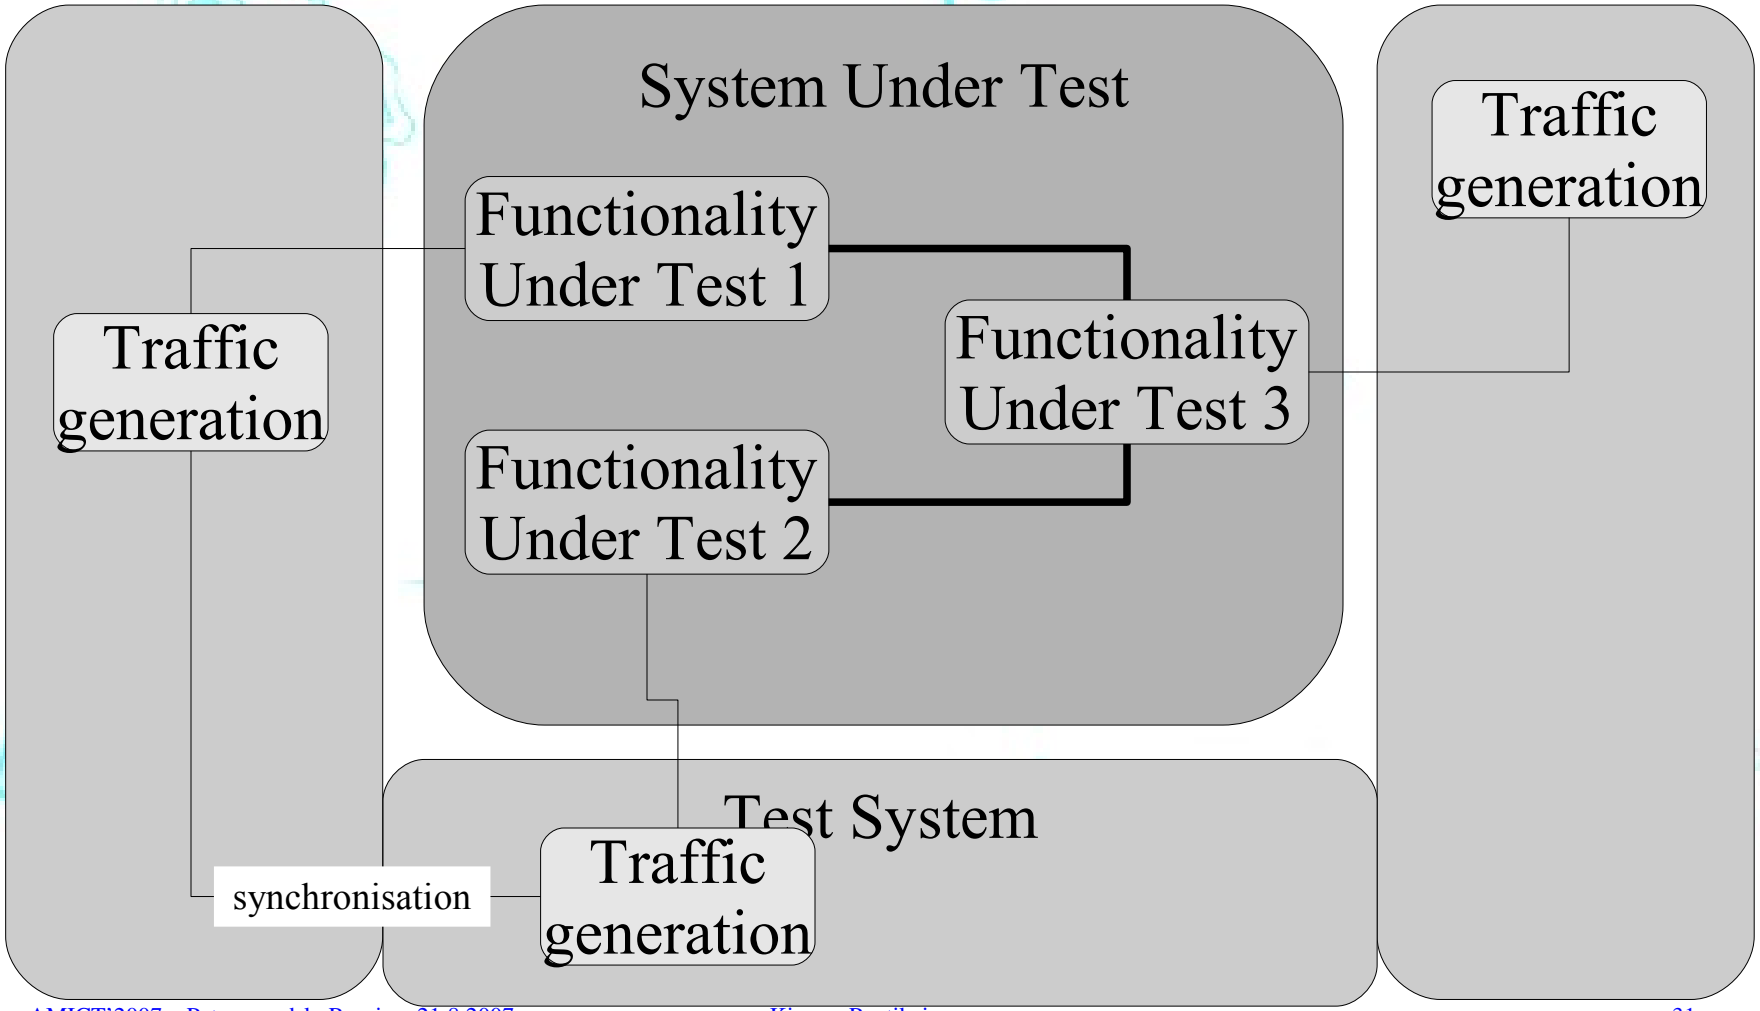

Benchmark Metrics - 1/2

- SAPS: (Scenario Attempts Per Second)
  - The average rate in one second period at which scenarios are attempted (not necessarily successful).
- TRT: (Transaction Response Time)
  - Defined as the time elapsed from the first message sent to initiate the transaction until the message ending the transaction is received.
  - Part 2 defines exact measurement points in the message sequence chart for each scenario.
  - The maximum response times of adequately handled scenarios are also specified.
  - In some scenarios there are separate thresholds for different segments in the scenari
- CPU: (CPU usage ratio)
  - The ratio of used CPU time to the total CPU time available.

Benchmark Metrics - 2/2

- MEM: (Memory usage ratio)
  - The ratio of used memory to the total memory available.
- RETR: (Retransmission Rate)
  - Applies to UDP transport.
- SIMS: (Simultaneous Scenarios)
  - Number of scenarios that are active at the same time.
- %IHS: (Percent Inadequately Handled Scenarios)
  - The ratio of inadequately handled scenarios to the total number of attempted scenarios.
  - Design Objective Capacity (DOC) defines maximum total round-trip time (TRT) of the scenario.
  - Under nominal load %IHS must be less than 1% and under stress condition (overload) less than 10%.

Use-Cases - 1/3

- Registration/De-registration
  - The assigned S-CSCF challenges the UE using authentication vectors obtained from the HSS or UPSF
  - 9 scenarios:
    - Successful Initial Registration without Synchronization;
    - Successful Initial Registration with Synchronization;
    - Re-Registration – User Currently Registered;
    - Re-Subscription – User Currently Registered;
    - Re-Registration – User Roaming;
    - UE Initiated De-Registration;
    - Network Initiated De-Registration;
    - Network Initiated De-Registration upon Roaming or Expiration;
    - Network Initiated Re-Authentication

Use-Cases - 2/3

- Session Set-Up/Tear-Down
  - 25 scenarios
  - four types of calls:
    - Successful,
    - abandoned (B party does not answer in time),
    - Rejected (B party immediately rejects the invitation),
    - failed
  - Eight scenarios for each of the first three call types:
    - Four combinations of resource reservation (yes or no) when both sides are IMS users.
    - Two combinations of resource reservation (yes or no) on originating side when the terminating side is non-IMS.
    - Two combinations of resource reservation (yes or no) on terminating side when the originating side is non-IMS.

Use-Cases - 3/3

- Page-mode Messaging
  - simple message exchange between two peers
  - normal call set-up and tear-down takes a minimum of five SIP messages and typically seven
  - simple messaging employs just two SIP messages.

  - Scenarios:
    - Successful Message Exchange
    - Unsuccessful Message Exchange – Called User Not Found

ETSI Specifications

- Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN);
- IMS/NGN Performance Benchmark.
  - Part 1: Core Concepts.
  - Part 2: Subsystem Configurations and Benchmarks.
  - Part 3: Traffic Sets and Traffic Profiles.
- Draft Technical Specification ETSI TS 186 008-1/2/3
- V 0.0.98
- European Telecommunications Standardisation Institute
- January 2007

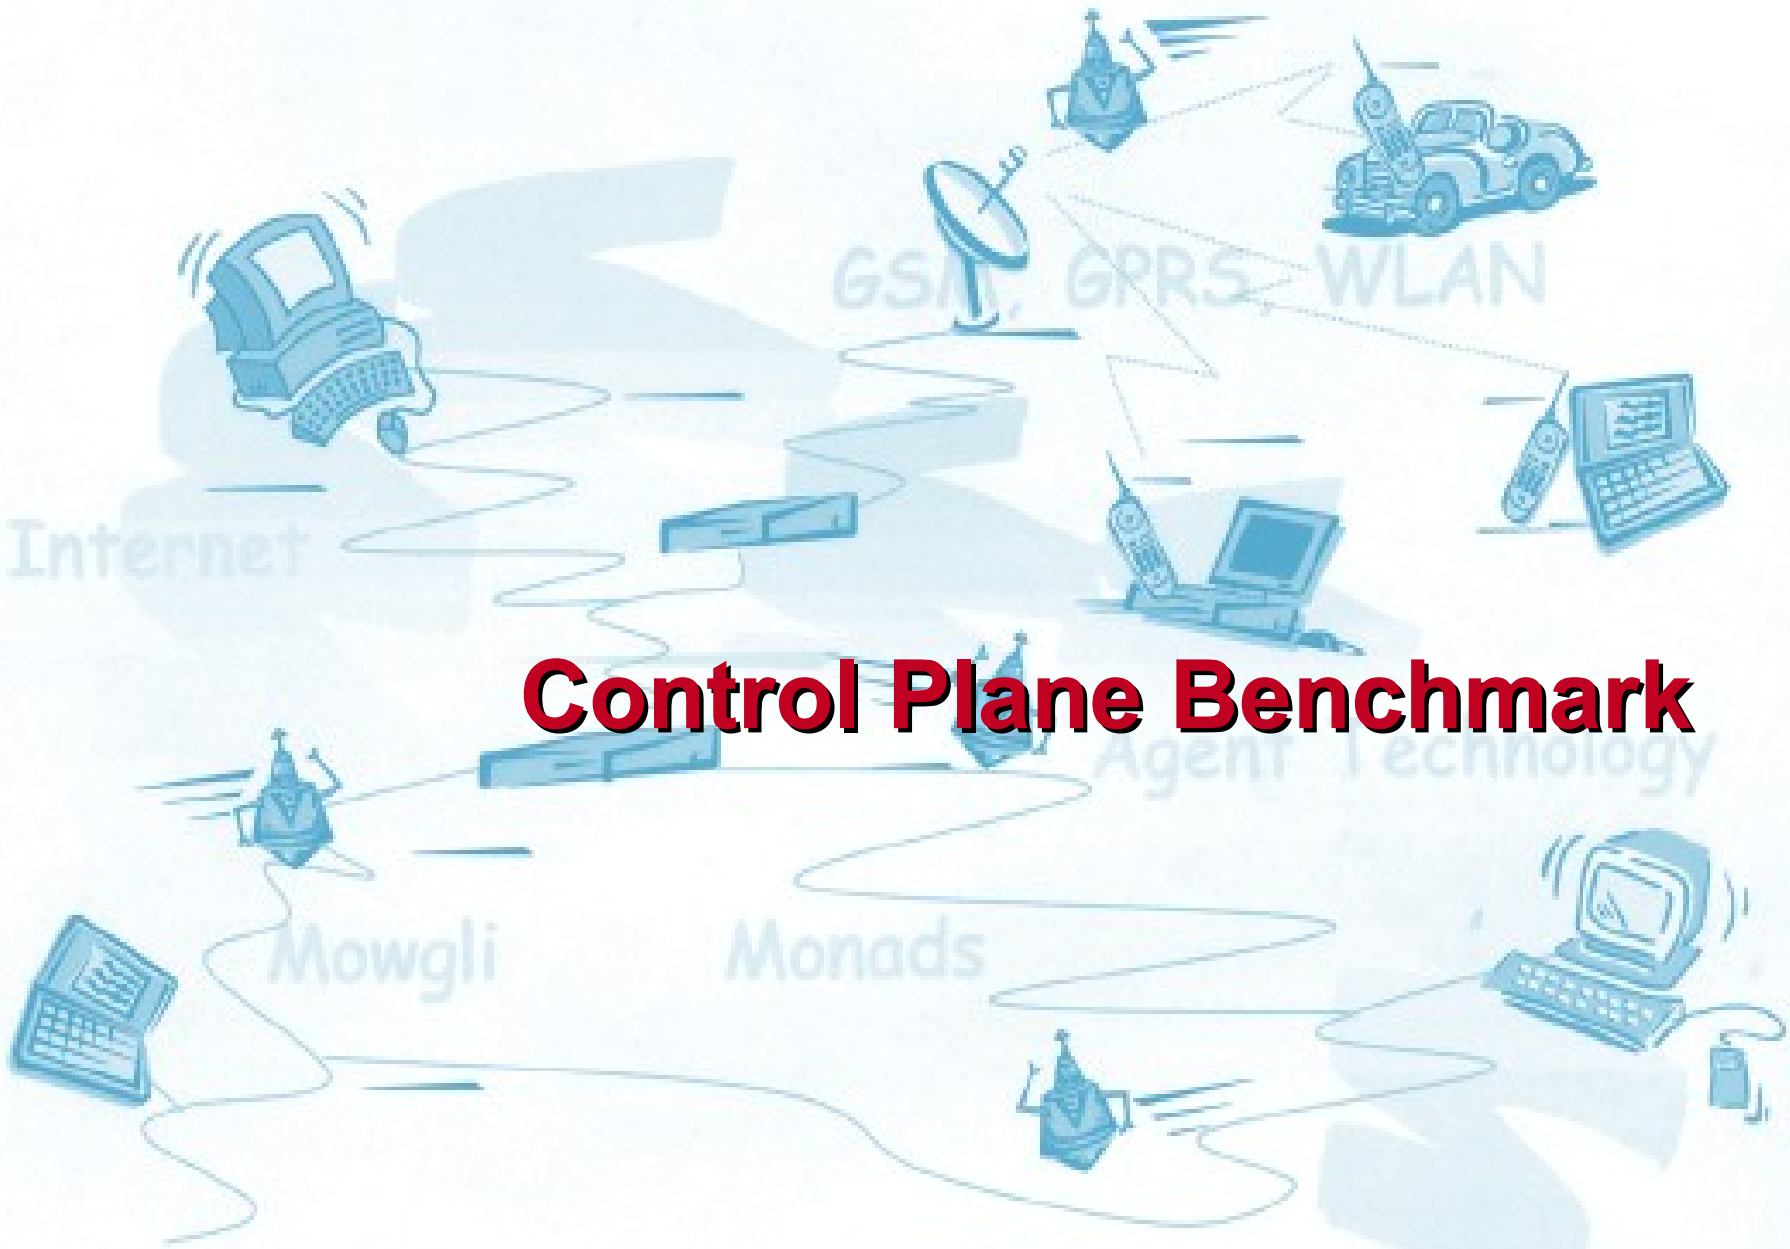

Control Plane Benchmark

Contol Plane Benchmark Framework

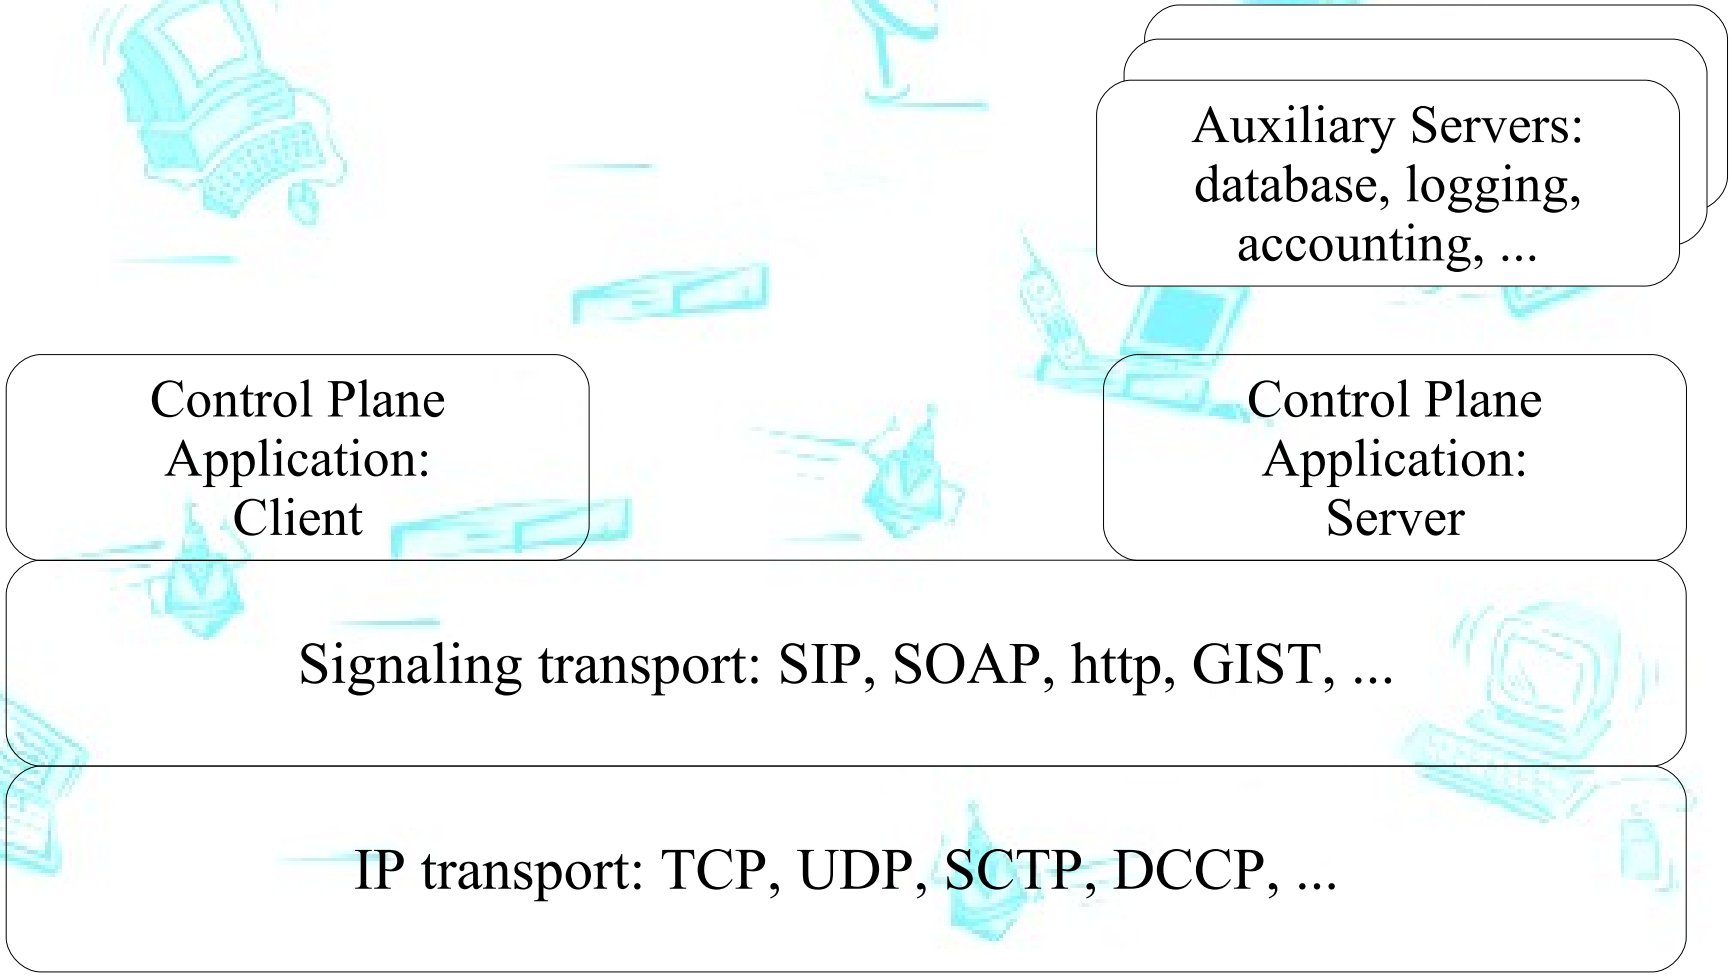

Control Plane Benchmark Architecture

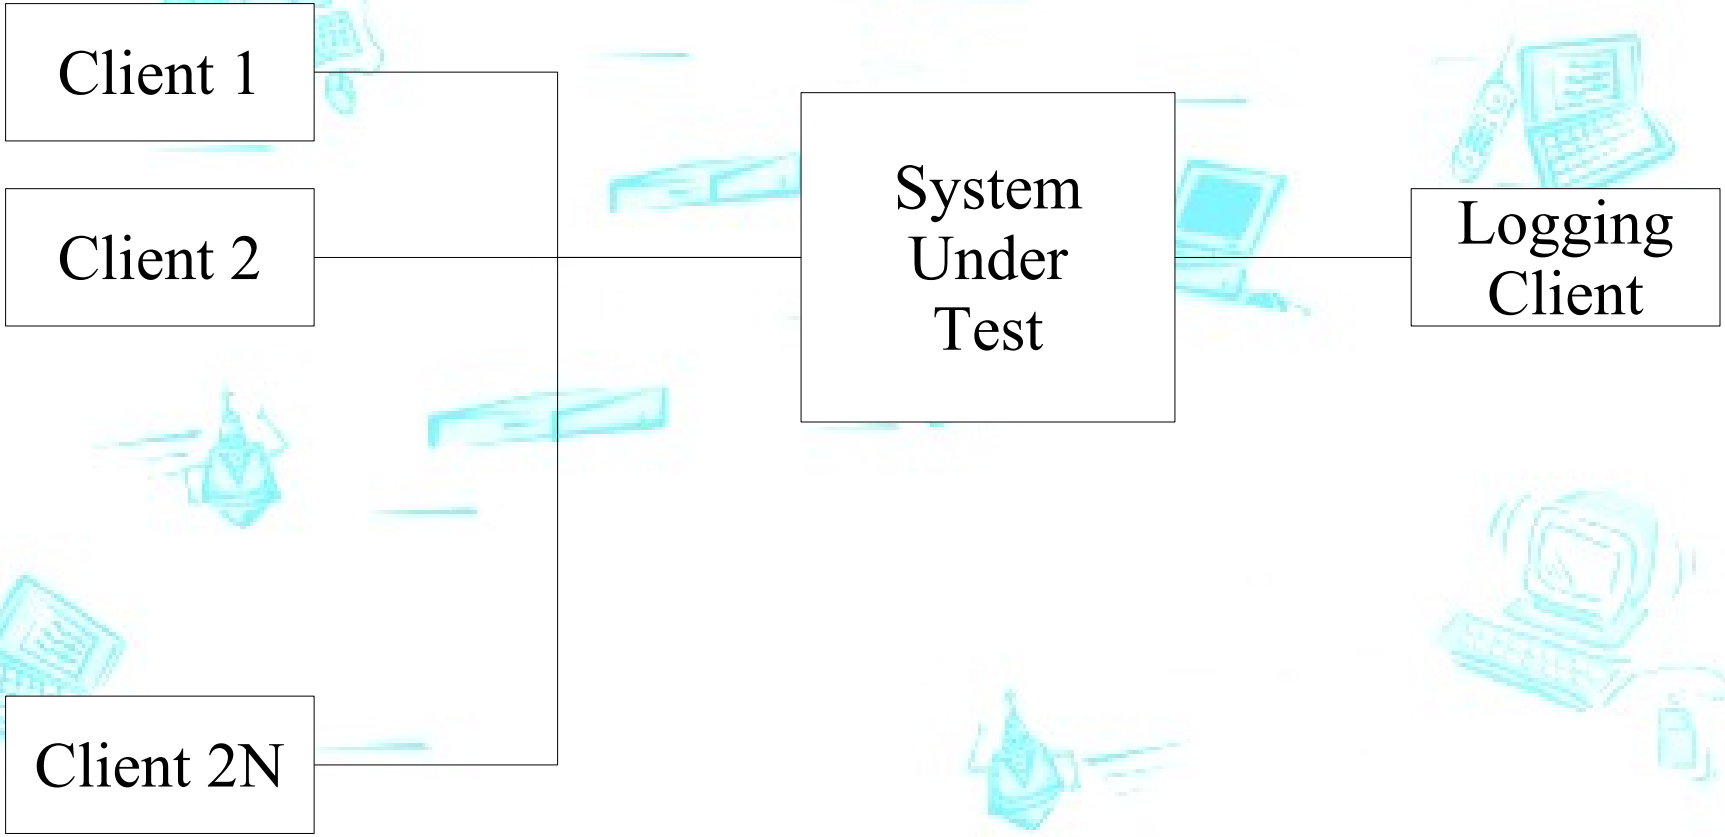

Basic Call Sequence Diagram

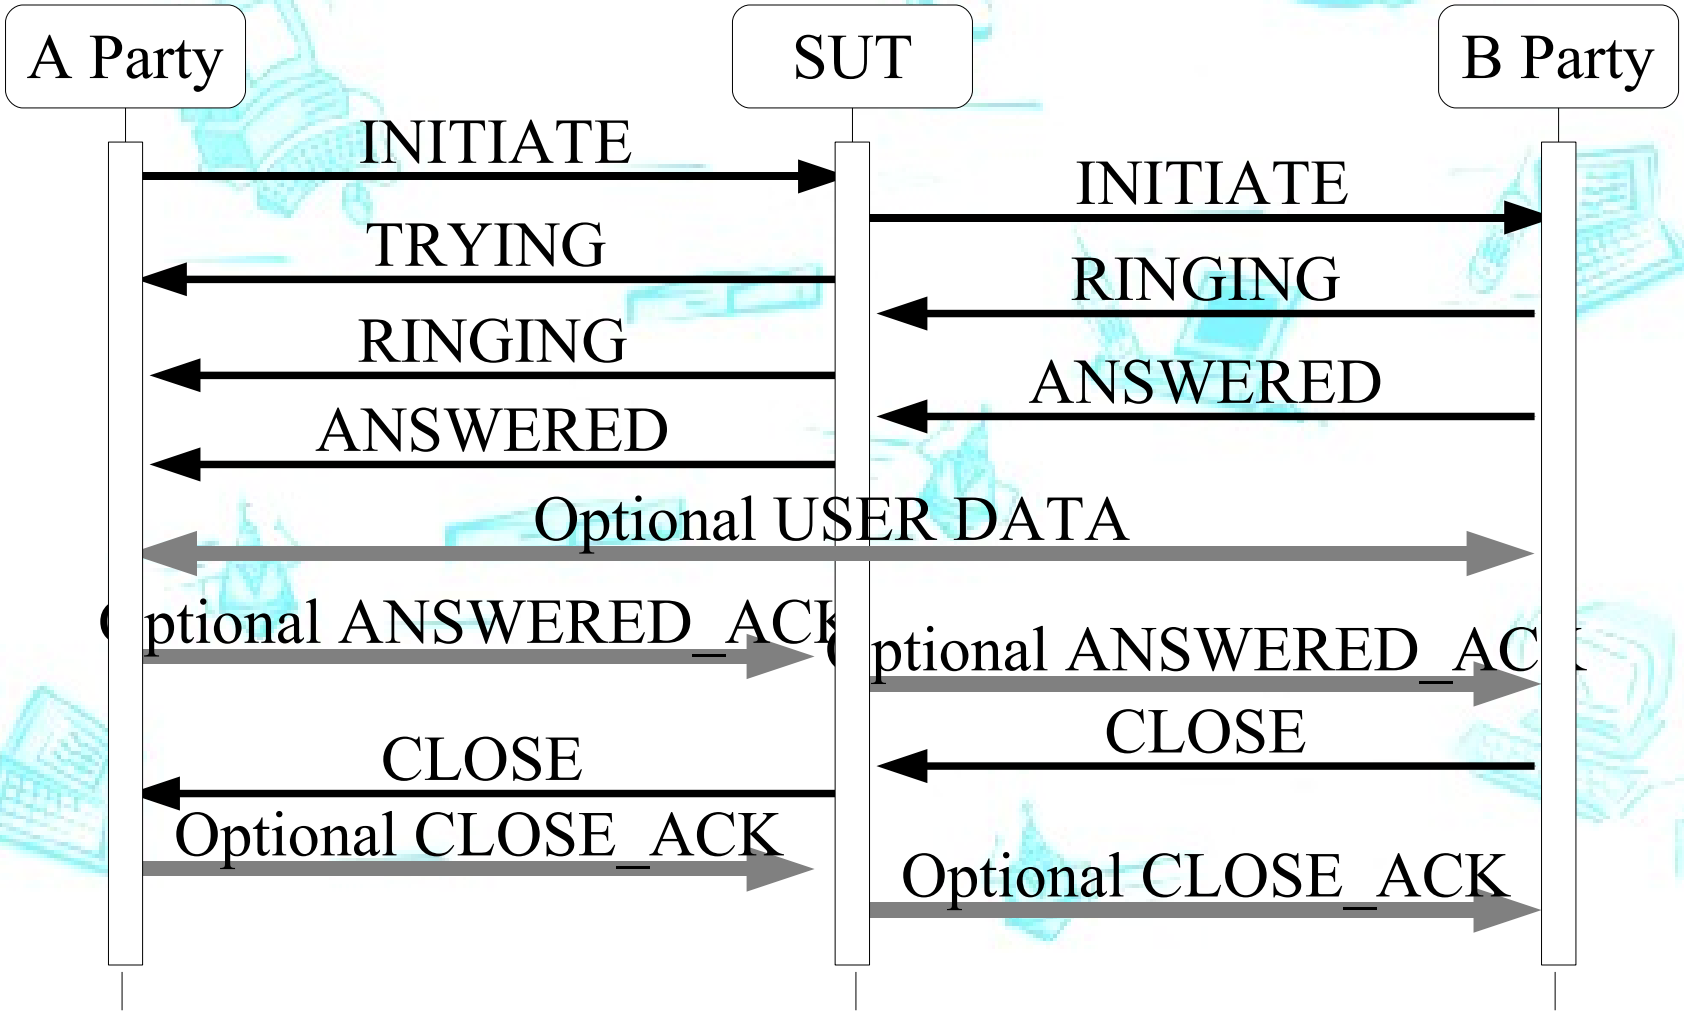

Echo Test Sequence Diagram

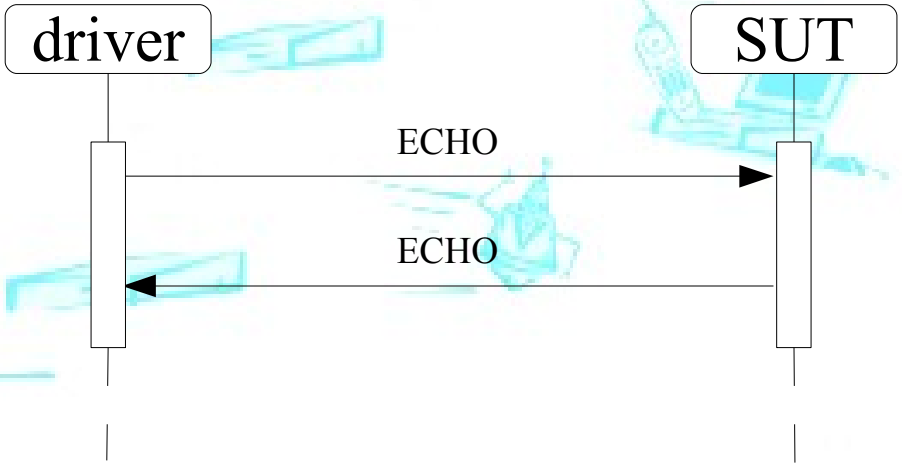

Telephony Server State Machine

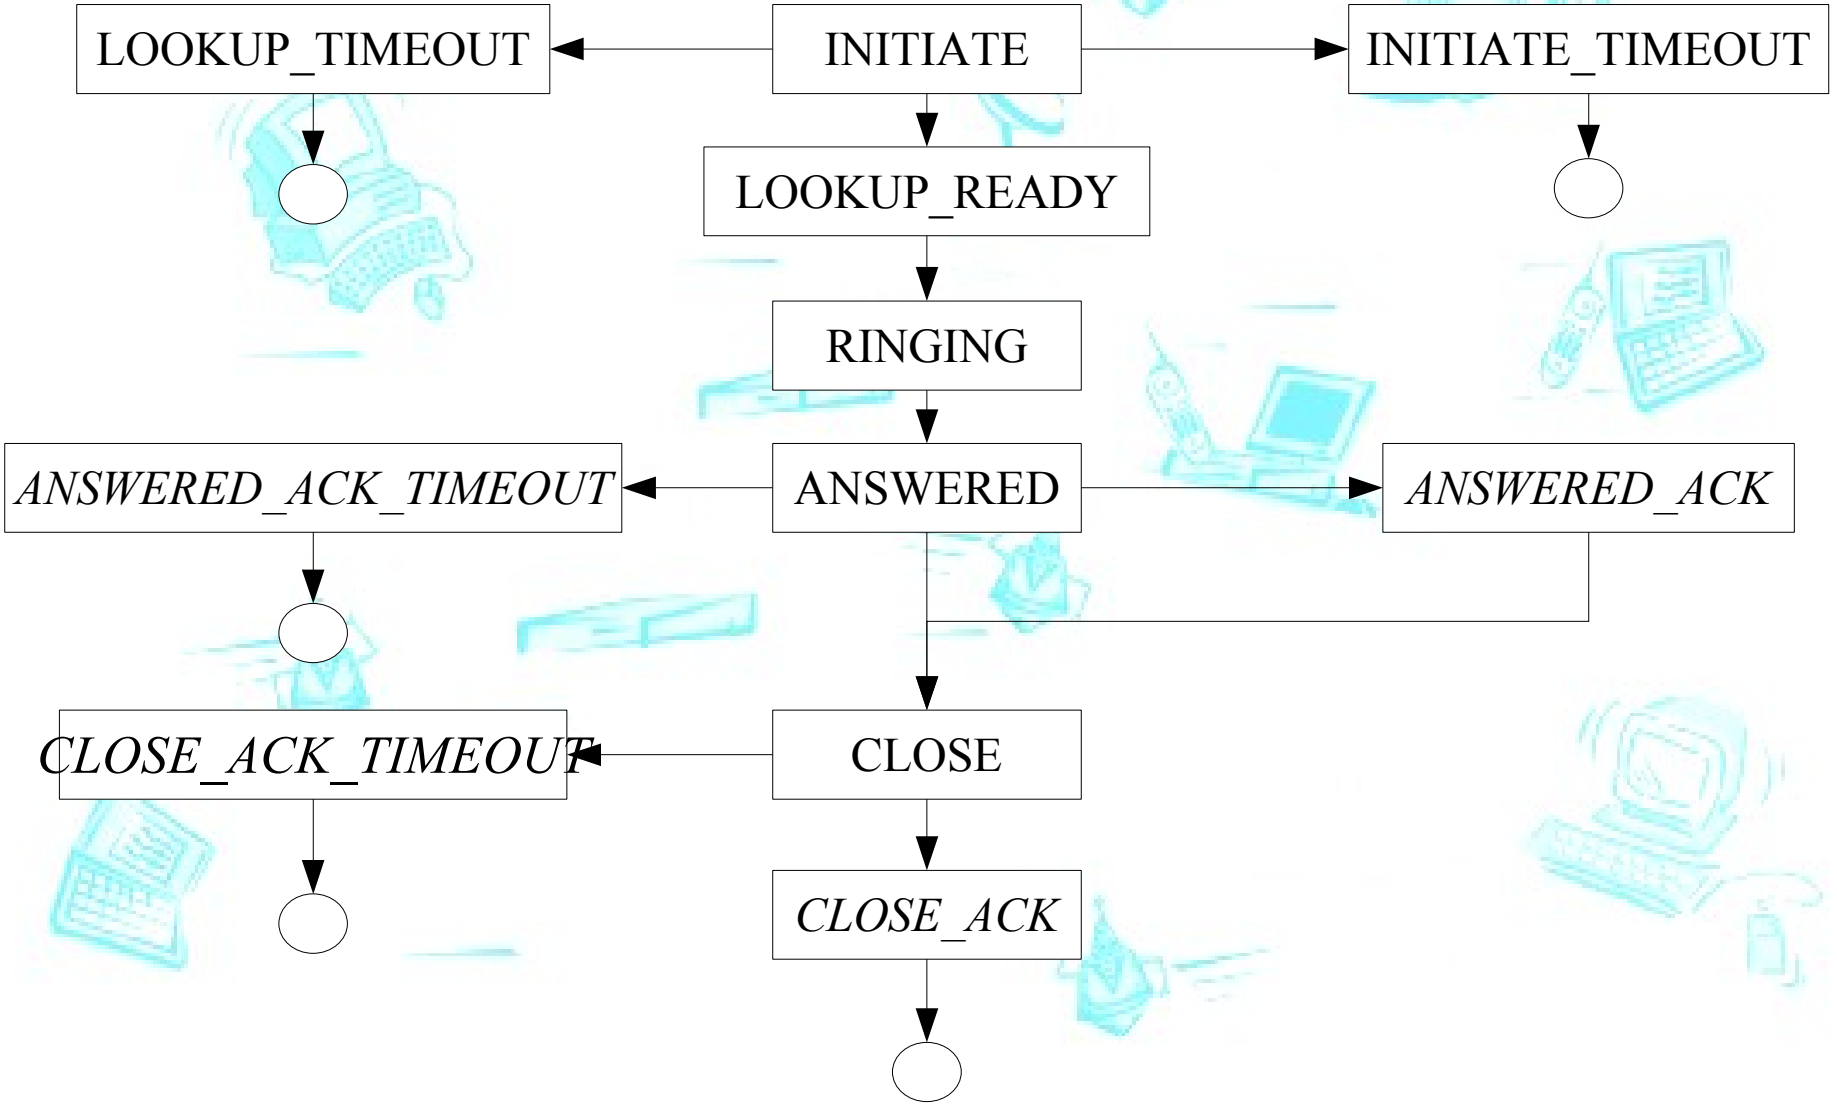

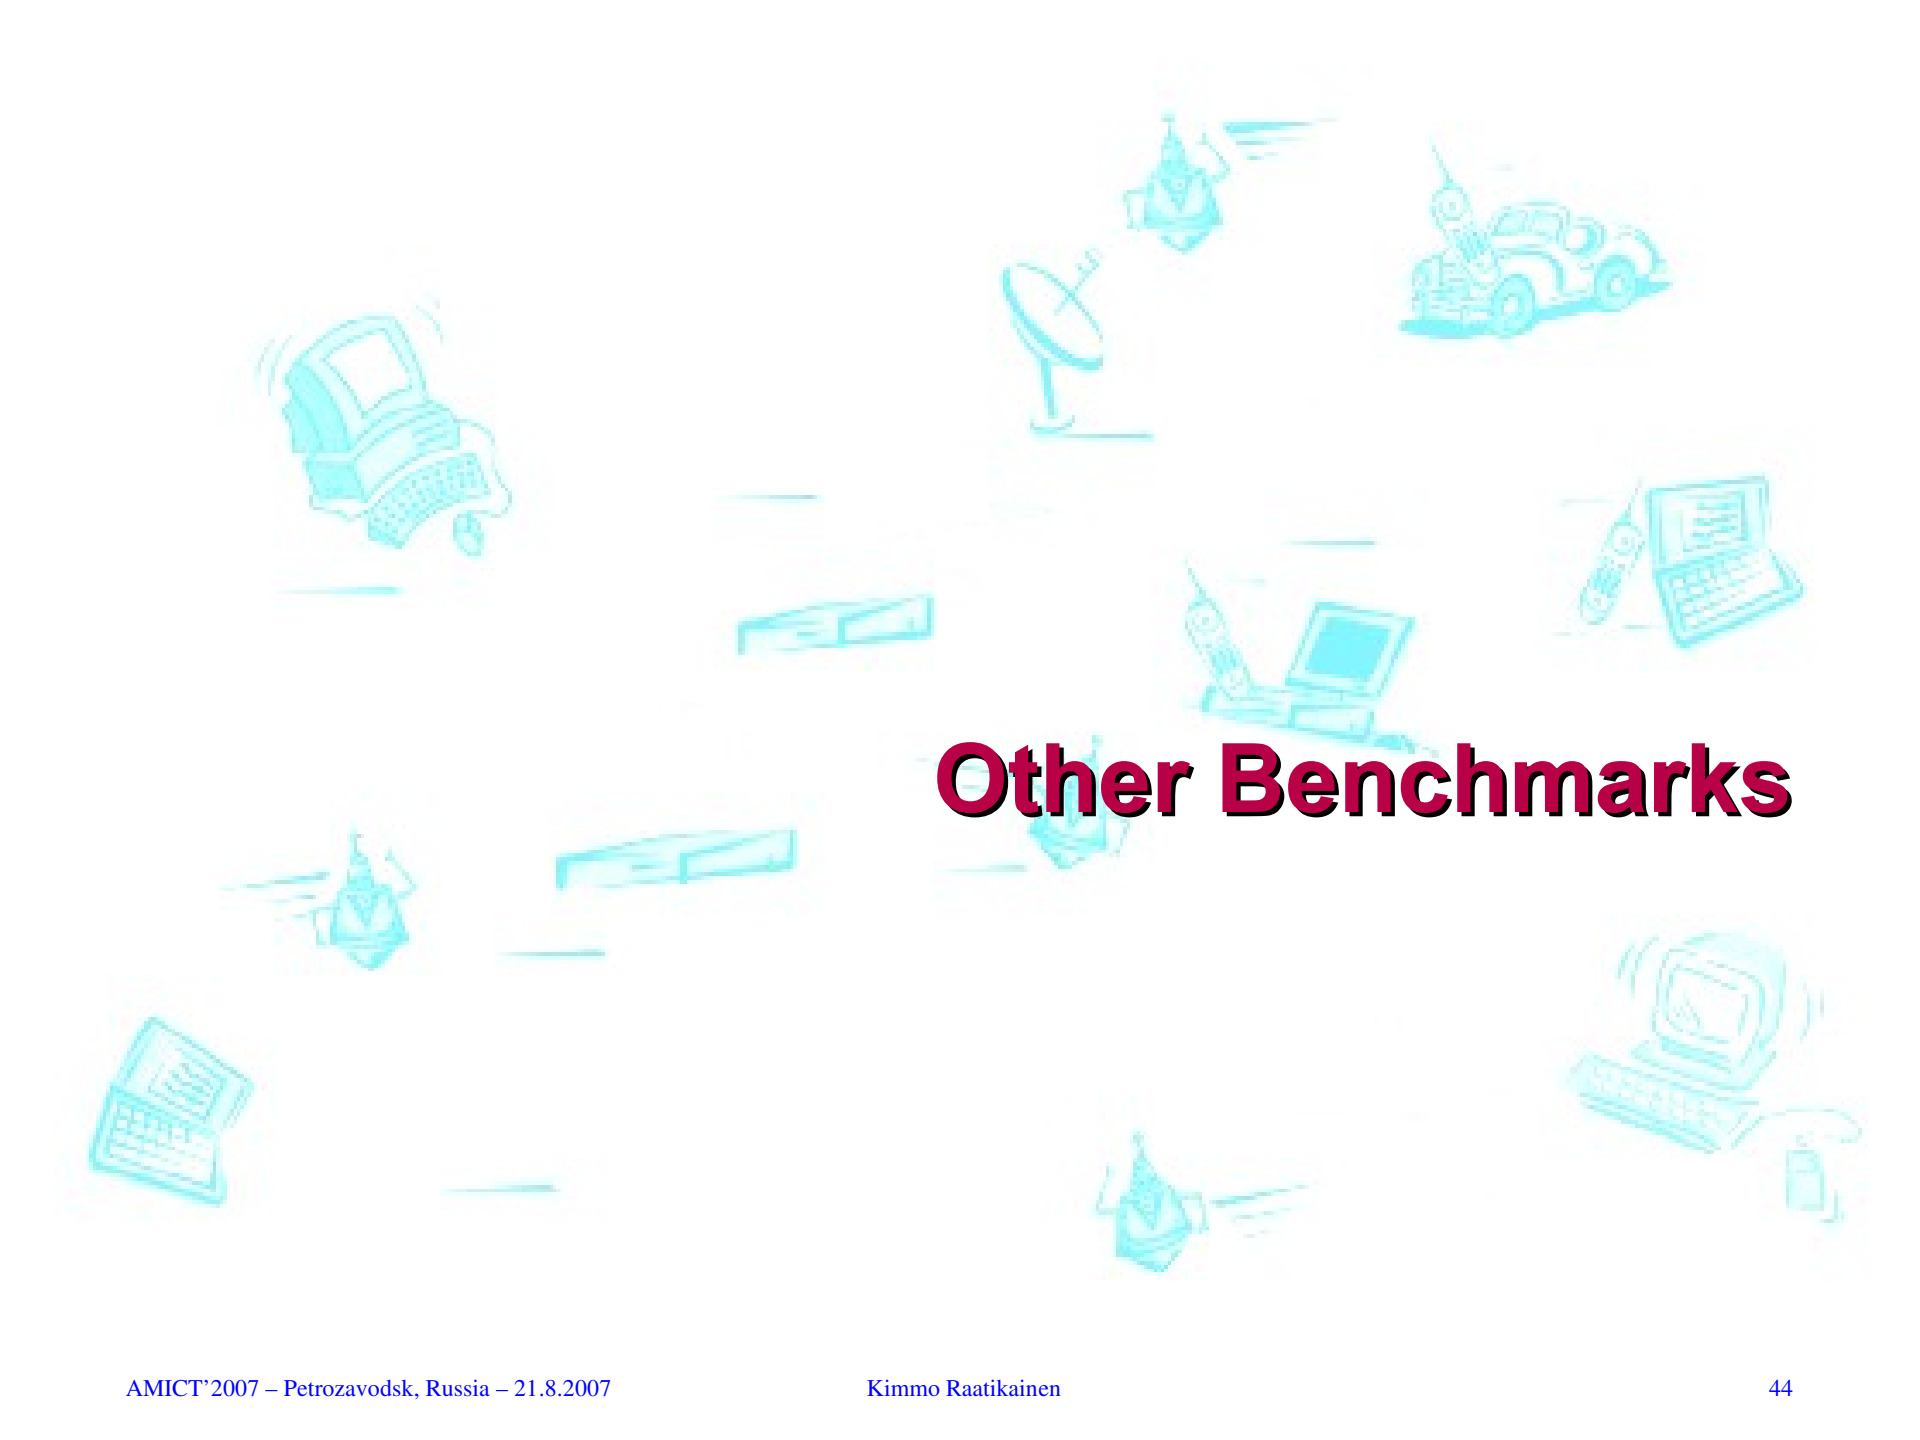

Other Benchmarks

XML Processing Benchmarks

- XMLTest
- XML Benchmark
- W3C EXI Measurement Framework

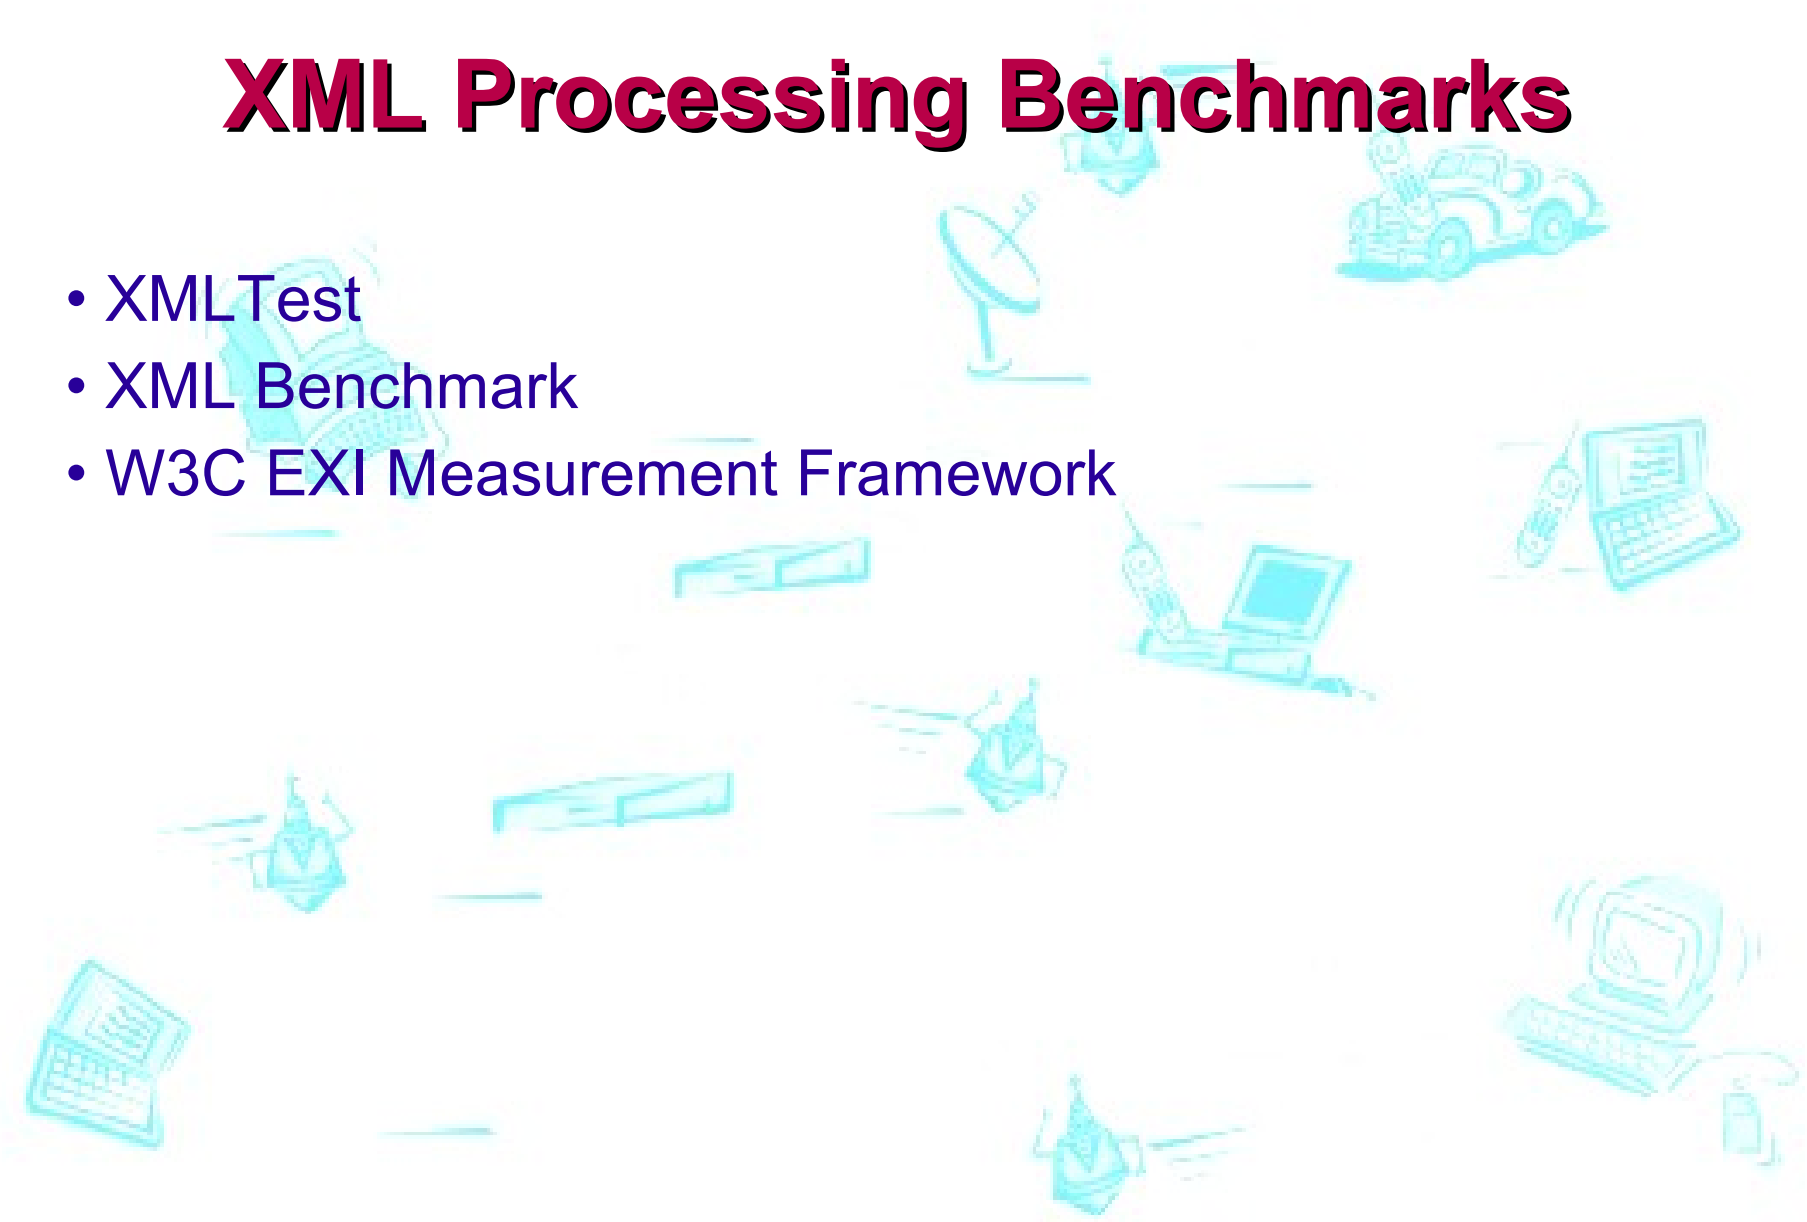

Database Benchmarks

- Transaction Processing Performance Council (TPC)
  - Classical OLTP (known as TPC-A) is outdated – current version TPC-C
  - TPC-E – OLTP Brokering
  - TPC-H — Decision Support for Ad Hoc Queries
  - TPC-App — Application Server
- OSDL (now Linux Foundation) have open source implementations (but different query rates)
- Telecom specific
  - Based on M.Sc. Thesis by Toni Strandell (Univ. Helsinki)
- XML Database benchmarks:
  - Xmach, Xmark, Xbench, X007, mbench

TPC-App Set-Up

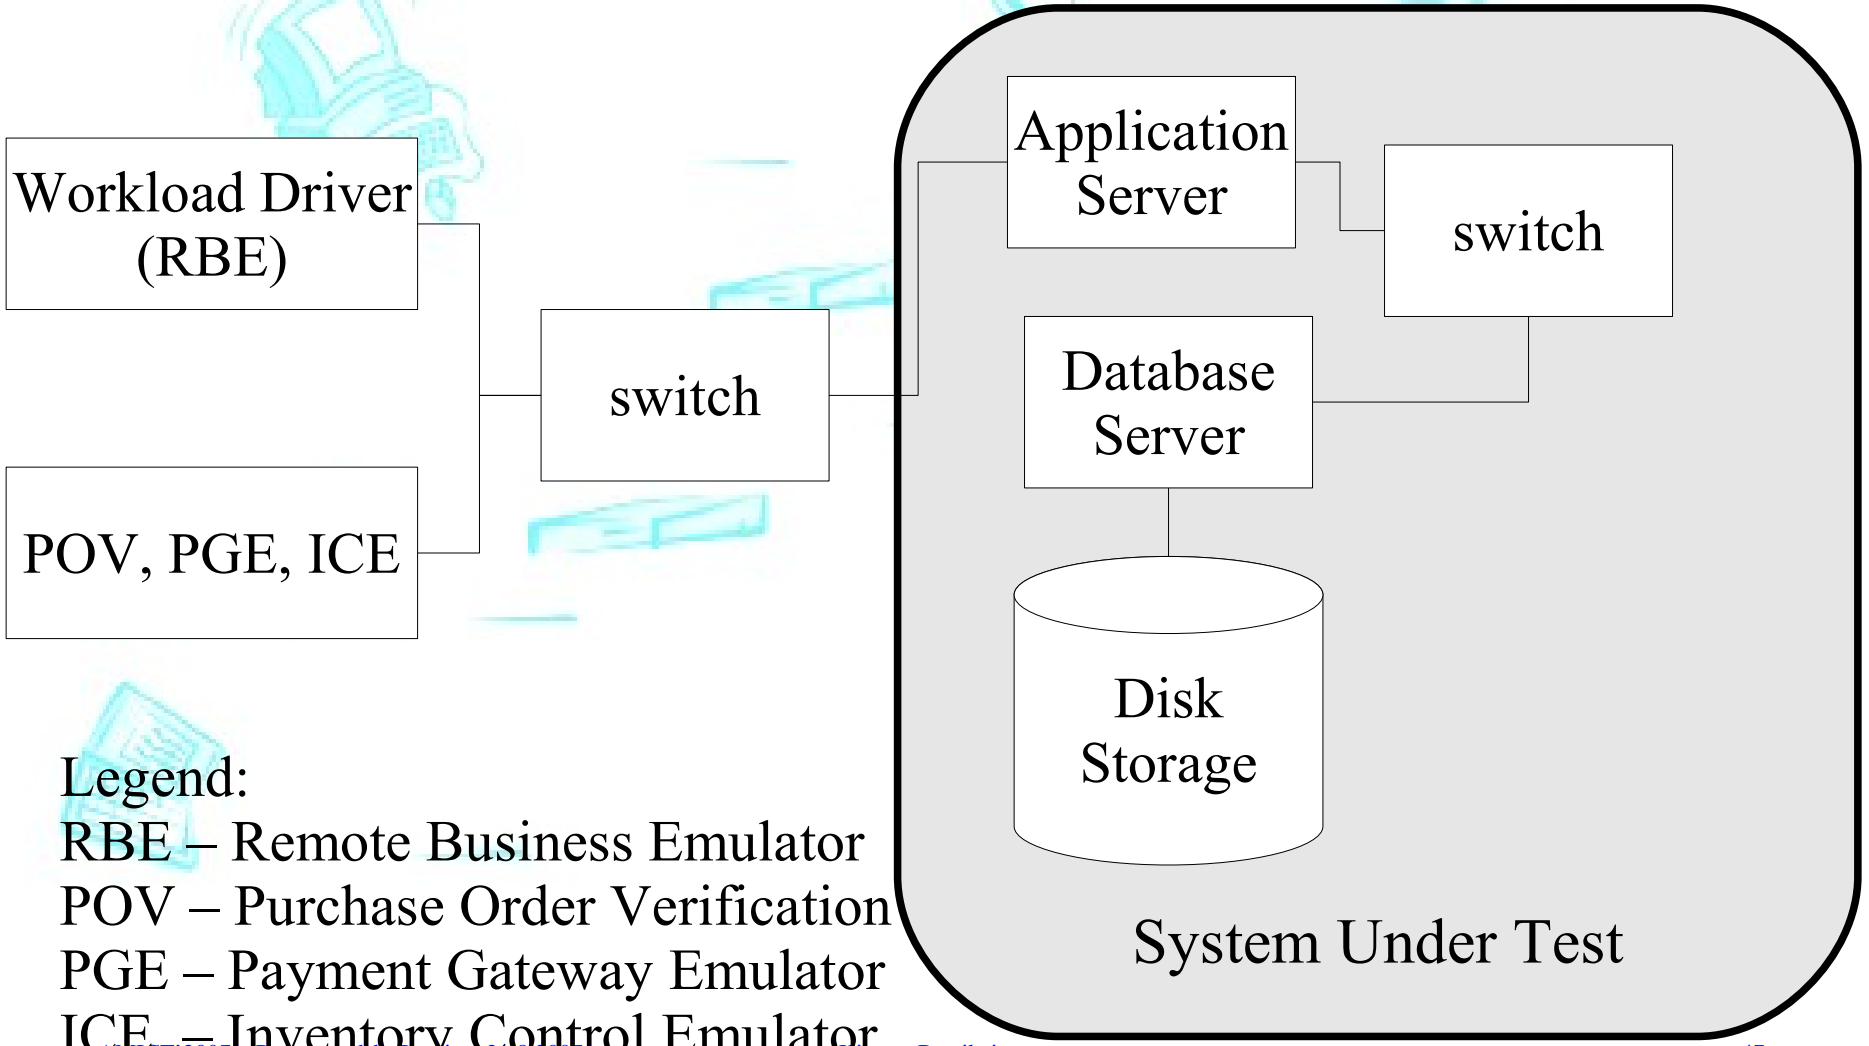

Legend:
RBE – Remote Business Emulator
POV – Purchase Order Verification
PGE – Payment Gateway Emulator
ICE – Inventory Control Emulator

Operating System Benchmark: lmbench

- Low level: integer operations , float operations, double operations, memory latencies
- OS-1: processor and processes: null call, null I/O, stat, open/close, slct TCP, sig inst, sign hndl, fork proc, exec proc, sh proc
- OS-2: Context Switch: 2p/0K, 2p/16K, 2p/64K, 8p/16K, 8p/64K, 16p/16K, 16p/64K
- OS-3: Local communication latencies: 2p/0K ctxsw, pipe, AF Unix, UDP, RPC/UDP, TCP, RPC/TCP, TCP conn
- OS-4: File and VM latencies: 0K create, 0K delete, 10K create, 10K delete, mmap latency, prot fault, page fault, 100fd select
- OS-5: Local communication bandwidths: pipe, AF Unix, TCP, File reread, Mmap reread, Bcopy (libc), Bcopy (hand), Mem read, Mem write

Comparing lmbench results

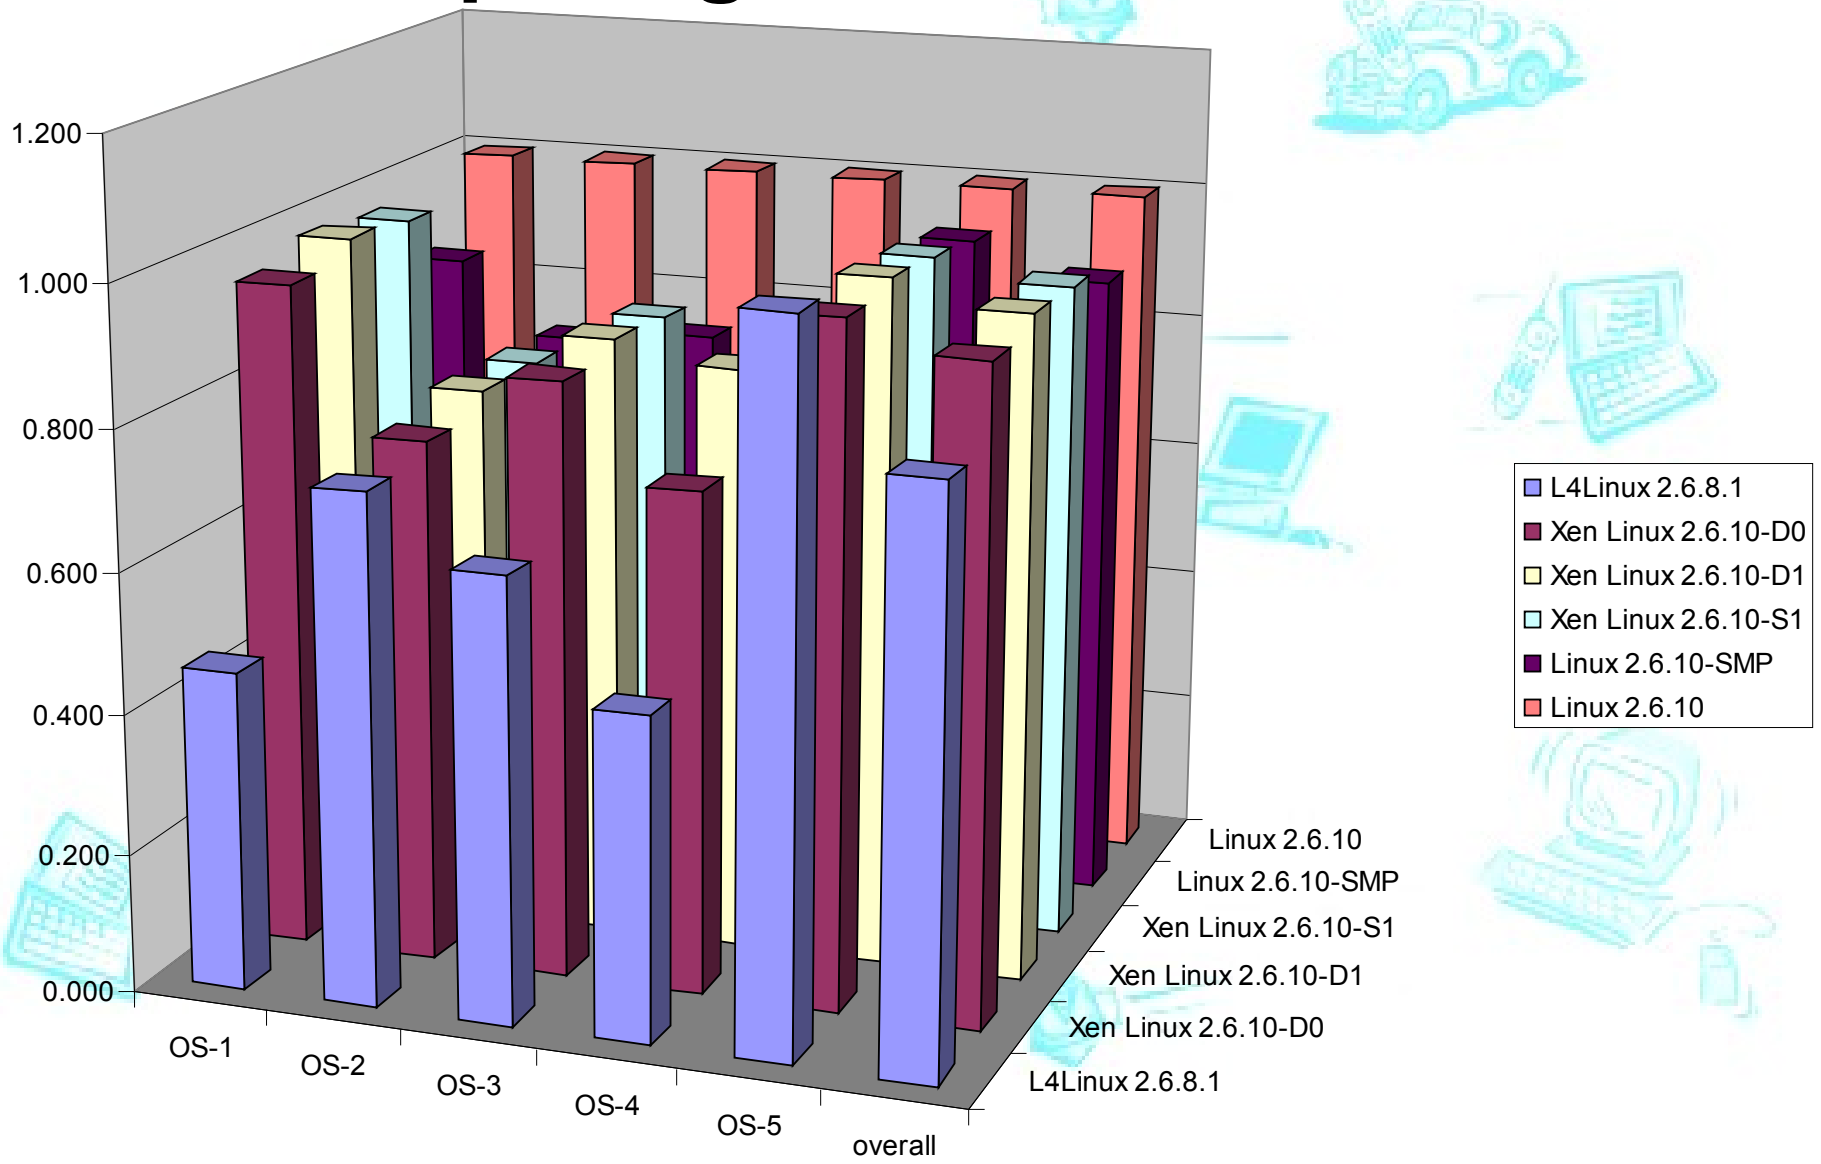

EEMBC: Embedded Microprocessor Benchmark Consortium

- Current Benchmark suites:
  - Automotive/Industrial Version 1.1
  - Consumer Version 1.1
  - Digital Entertainment
  - GrinderBench for the Java 2 Micro Edition (J2ME) Platform
  - Networking Version 1.1 and Version 2.0
  - Storage
  - Office Automation Version 1.1
  - Telecom Version 1.1
  - Energy Consumption

Some EEMBC Benchmark Suites

- J2ME (GRINDERmark)
  - Chess
  - Cryptography
  - kXML
  - ParallelBench
  - PNG Decoding
  - Regular Expression
- Consumer Version 1.1 (CONSUMERmark)
  - High Pass Grey-Scale Filter
  - JPEG
  - RGB to CMYK Conversion
  - RGB to YIQ Conversion
  -
- Networking 2.0 (NETmark)
  - IP Packet Check
  - IP Network Address Translator (NAT)
  - Open Shortest Path First (OSPF)
  - Quality of Service (QOS)
  - Route Lookup
- Telecom Version 1.1 (TELEmark)
  - Autocorrelation
  - Bit Allocation
  - Convolutional Encoder
  - Fast Fourier Transform (FFT)
  - Viterbi Decoder

Standard Performance Evaluation Corporation (SPEC)

SPEC

- CPU2006 (see ACM SIGARCH CAN March 2007)
  - 12 integer (CINT2006) and 17 float-point (CFP2007) benchmark programs
  - CPU2006: 3334 kLOC and 5261 modules
  - CPU2000: 811 kLOC and 1432 mpdules
- Graphics
- High Speed Computing
  - HPC2002, OMP2001, MPI2006
- SPECjAppServer 2004
- SPECjbb 2005
- SPECweb2++5
- SFS97_R1 (Network File System)
- MAIL2001

SPECweb2005 Set-up

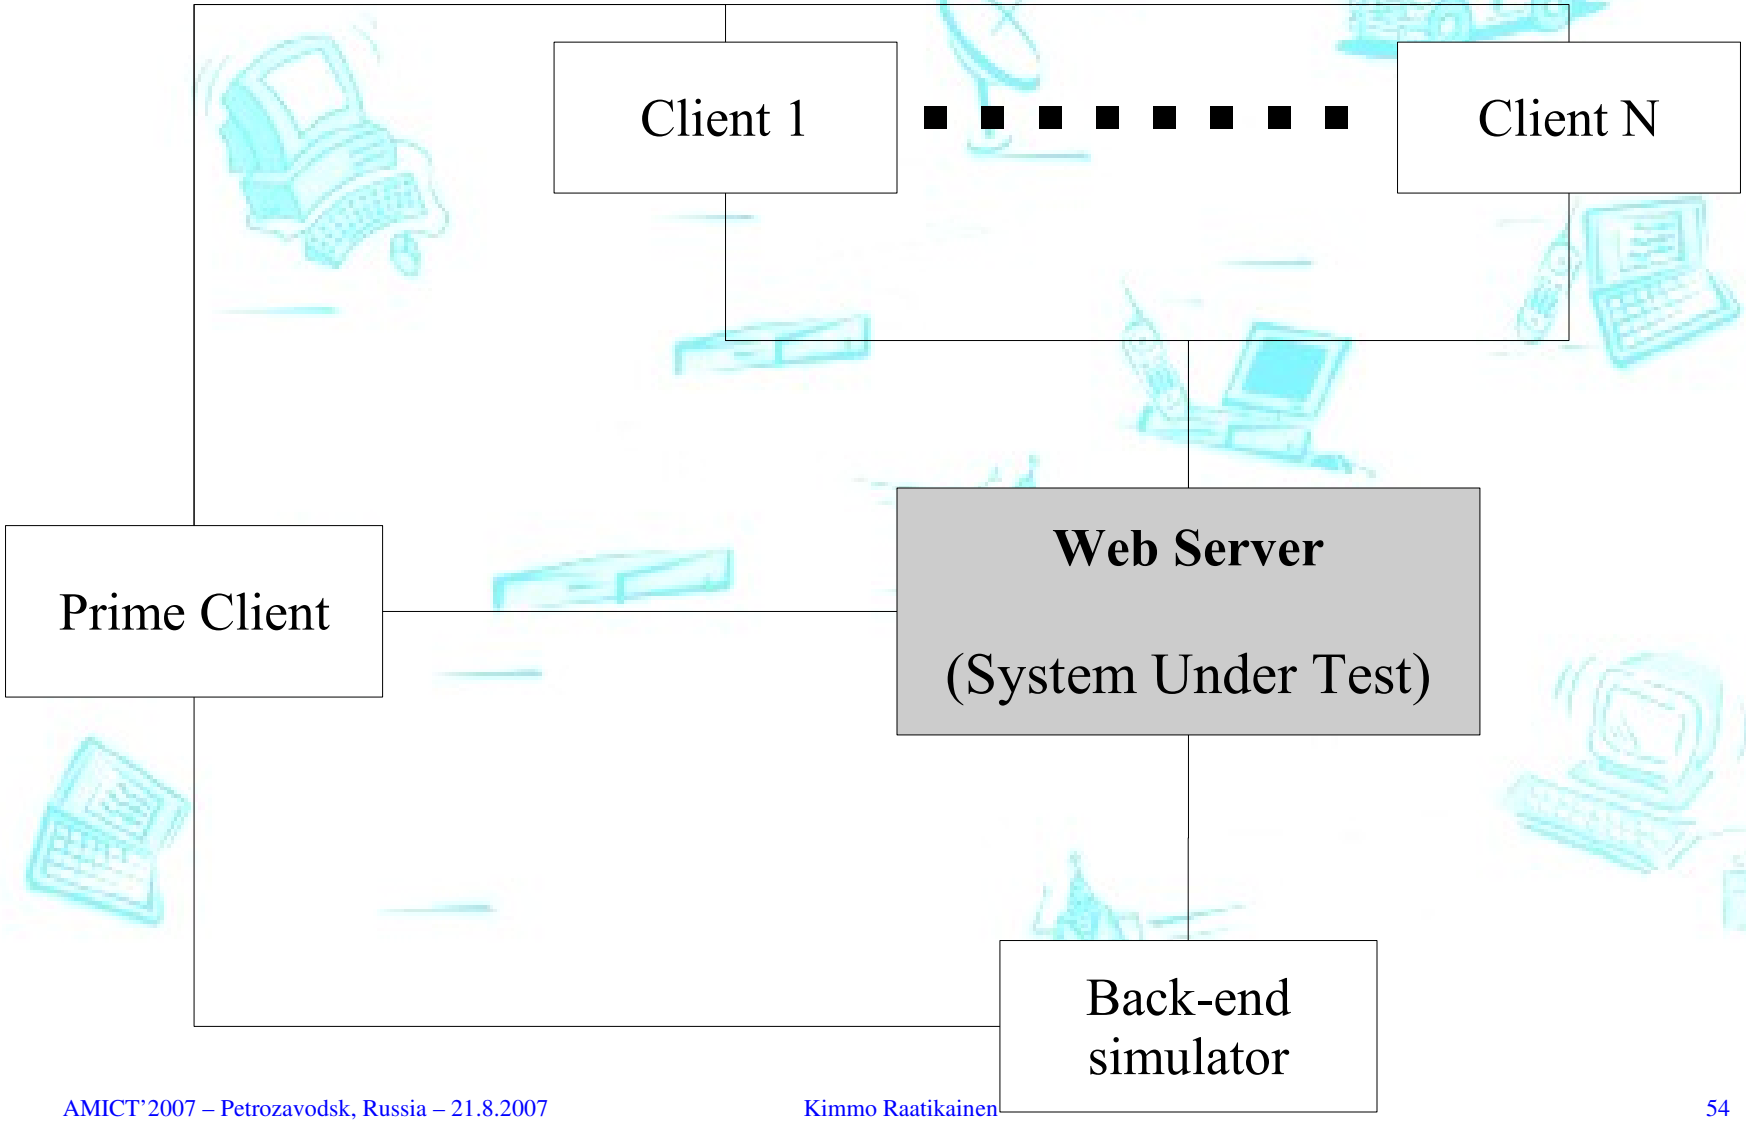

MIPS

- MIPS, that is, millions of instructions per second
  - are only meaningful in the context of a single processor family

- Other explanations
  - meaningless indicators of performance for systems,
  - meaningless information of performance for salesmen, or
  - meaningless information from pushy salesmen.

Out-dated Classics

- Dhrystone
  - Too small, fits most L1-caches
  - Good analysis in EEMBC White paper
- Whetstone
- Perfect (Digital Review)

Summary

Summary

- Our strong recommendation is to use standardized and well-established benchmarks
  - there is usually a large number of certified results available.
- Benchmark execution and reporting requirements by EEMBC, NPForum, SPEC and TPC are detailed enough to trust and exploit available benchmarking results.
- quite often an implementation is also available
  - benchmark programming is challenging
  - If you are not careful, you end-up measuring efficiency of your measurement system, not the phenomenon
  - You must try to avoid changing the program behavior by your instrumentation code.
  - The most important issues are related to the I/O-behavior and caches.

Selection Framework

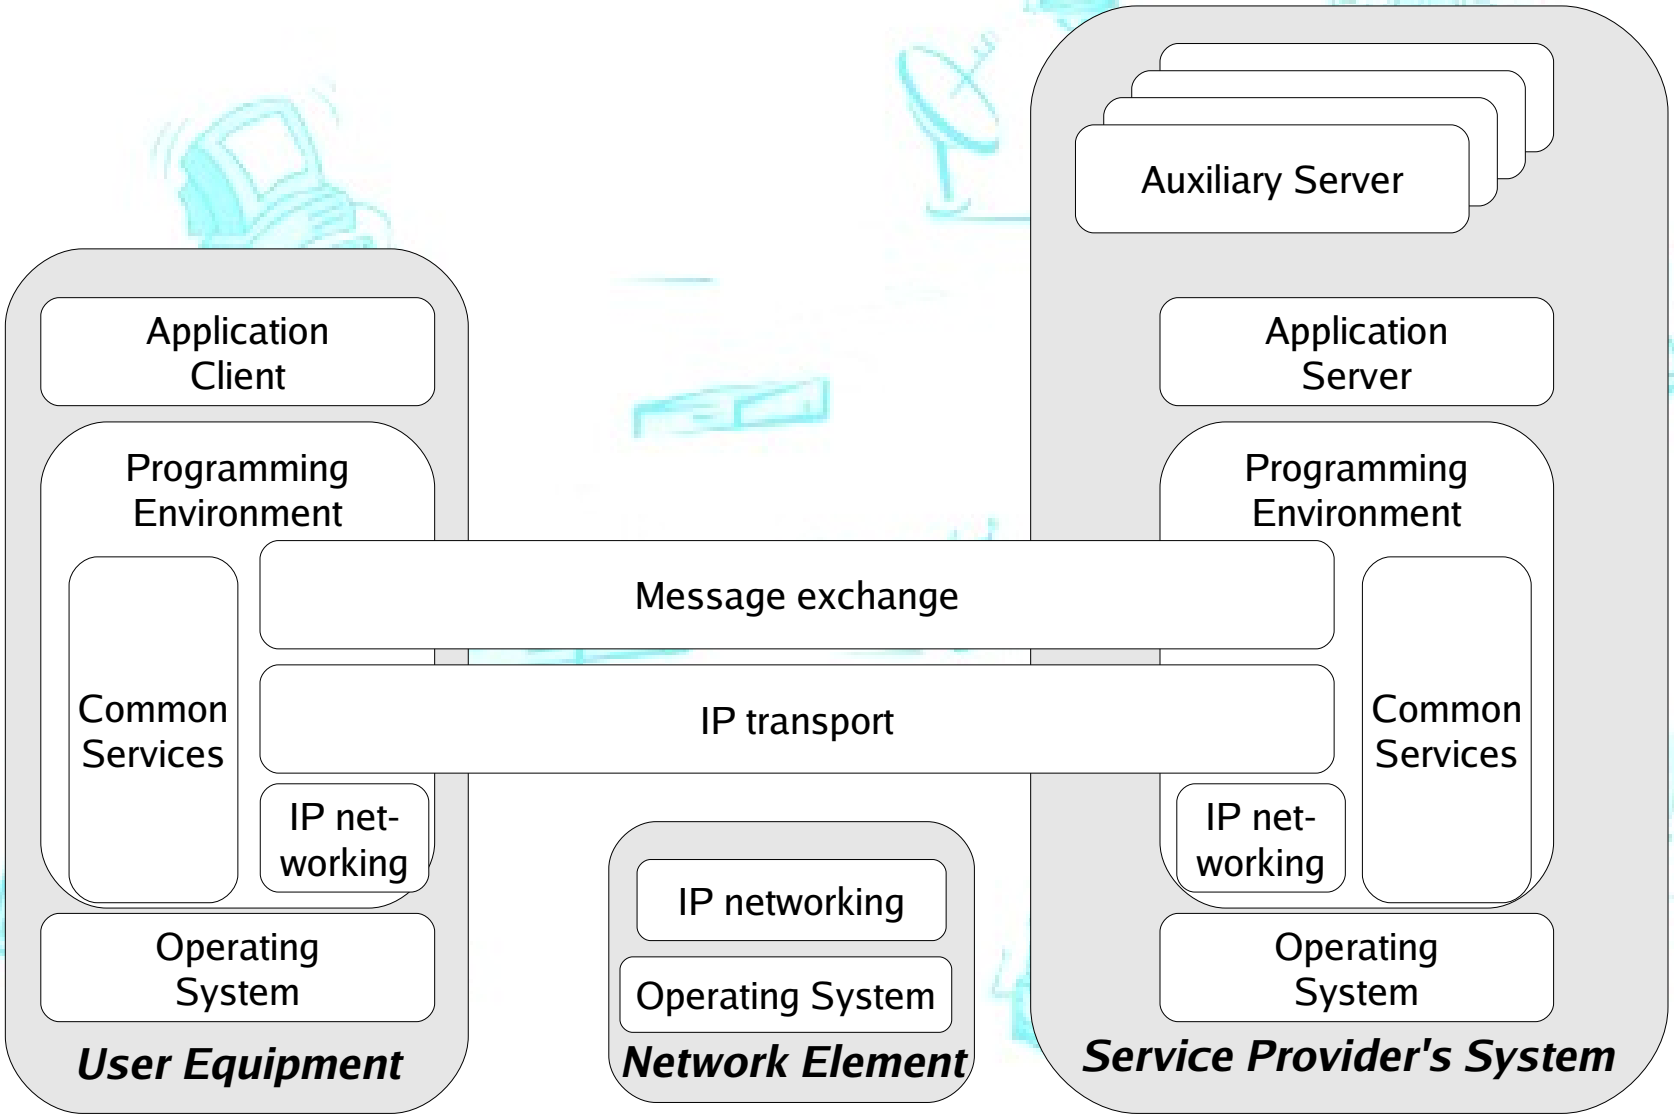

Recommended selections - 1/3

- The Imbench is a de facto standard operating system benchmark.
- IP networking bechmarks include IPSec Forwarding Application Level Benchmark, IP Forwarding Benchmark, MPLS Application Level Benchmark, and Ipv4 Forwarding Benchmark by Network Processing Forum
- TCP Proxy Application Level Benchmark belongs to IP Transport benchmarks.
- NetBench by EEMBC is a good suite of microbenchmarks for IP Networking and IP Transport.

Recommended selections - 2/3

- Common Services covers
  - XML processing,
  - format encoding and decoding,
  - Compression, transforms,
  - filtering and cryptography.
- Most of EEMBC benchmark suites (AutoBench, ConsumerBench, DENBench, GrinderBench and TeleBench) belong to Common Services
- In Programming Environments we have Web Services and Java (particularly JVM and J2xE) benchmarks
  - SPECjAppServer2004, SPECjbb2005, SPECweb2005
  - TPC-App

Recommended selections - 3/3

- The Open Source Telecommunications Database Benchmark is a database benchmark tailored to a telecommunication use-case
- TPC-C is a general purpose database (on-line transaction processing) benchmark.
-
- Control Plane applications benchmarks for telecoms:
  - ETSI TISPAN IMS/NGN Performance Benchmark
  - CplaneBench

Advertisement

- M.Sc. level course on Performance Issues in Mobile Computing and Communications
  - Part I: Current trends and recent developments in hardware and software for MCC
  - Part II : Performance evaluation
- Lectures on Mondays from 5pm to 7pm (Finnish time)
  - Sep 10 - Oct 8 and Oct 29 – Nov 26
  - two examinations
  - recorded lectures will be available in Web
- For details, consult my teaching Web-page
  - http://www.cs.helsinki.fi/u/kraatika/Courses/