

Evaluation and Validation

Peter Marwedel
TU Dortmund, Informatik 12
Germany

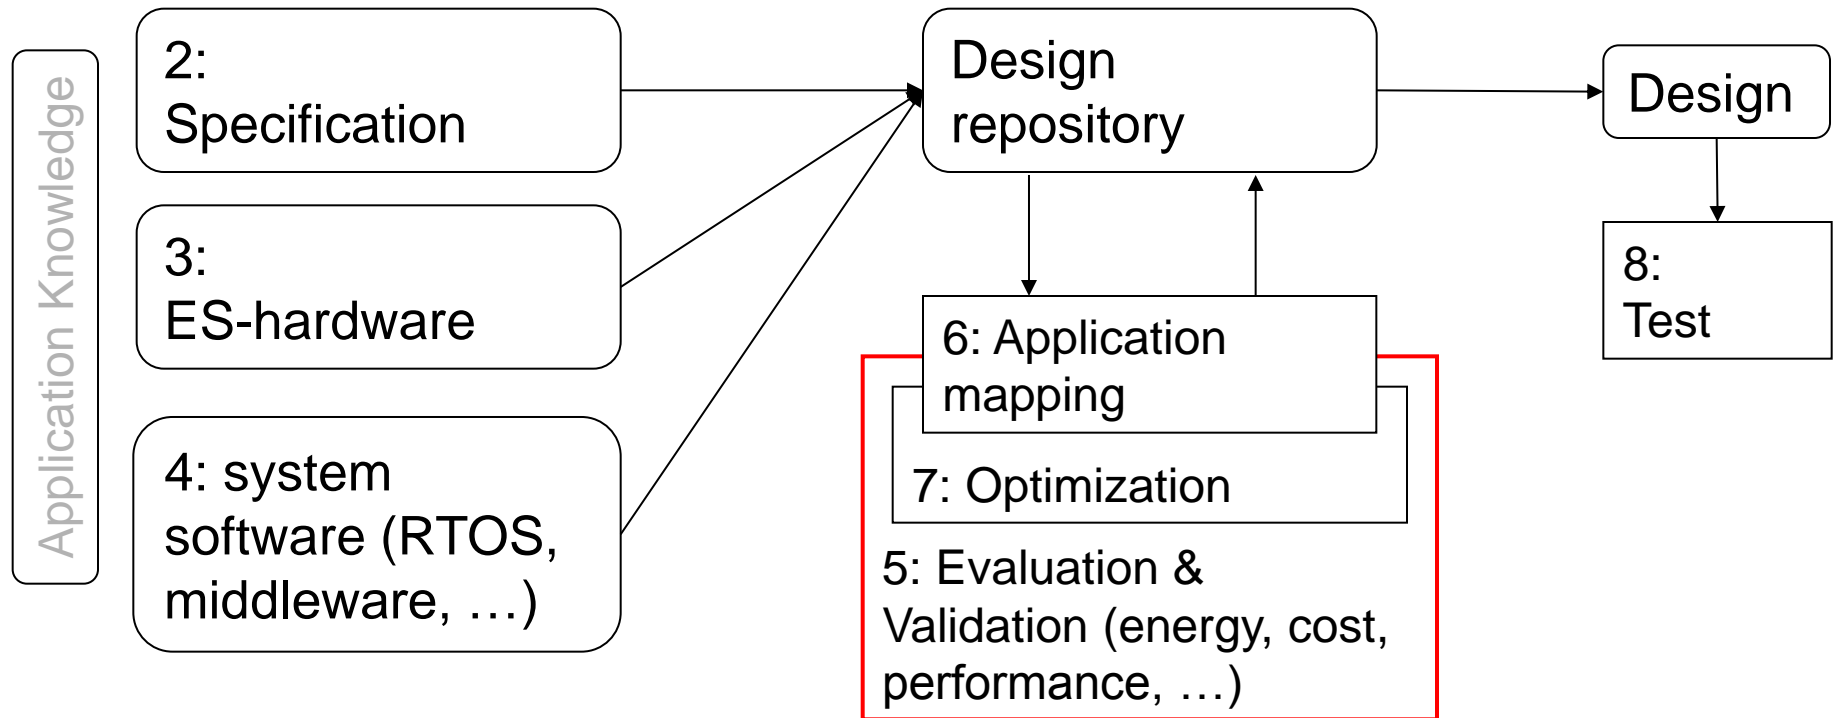


© Springer, 2010

2013 年 12 月 03 日

These slides use Microsoft clip arts. Microsoft copyright restrictions apply.

Structure of this course




Numbers denote sequence of chapters

How to evaluate designs according to multiple criteria?

Many different criteria are relevant for evaluating designs:

- ➔ Average & worst case delay
- power/energy consumption
- thermal behavior
- reliability, safety, security
- cost, size
- weight
- EMC characteristics
- radiation hardness, environmental friendliness, ..




How to compare different designs?
(Some designs are “better” than others)


Real-time calculus (RTC)/ Modular performance analysis (MPA)

Streams of events important: Examples

periodic event stream





periodic event stream with jitter



Thiele et al. (ETHZ): Extended **network calculus**:


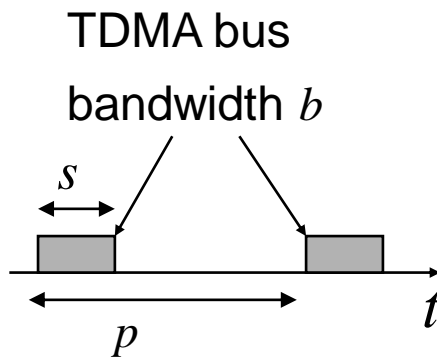
Arrival curves describe the maximum and minimum number of events arriving in some time interval Δ .



RTC/MPA: Service curves

Service curves β^u resp. β^l describe the maximum and minimum service capacity available in some time interval Δ


Example:



RTC/MPA: Workload characterization

γ^u resp. γ^l describe the maximum and minimum service capacity required as a function of the number e of events.

Example:



RTC/MPA:

Workload required for incoming stream

Incoming workload

$$\alpha^u(\Delta) = \gamma^u(\bar{\alpha}^u(\Delta)) \qquad \alpha^l(\Delta) = \gamma^l(\bar{\alpha}^l(\Delta))$$

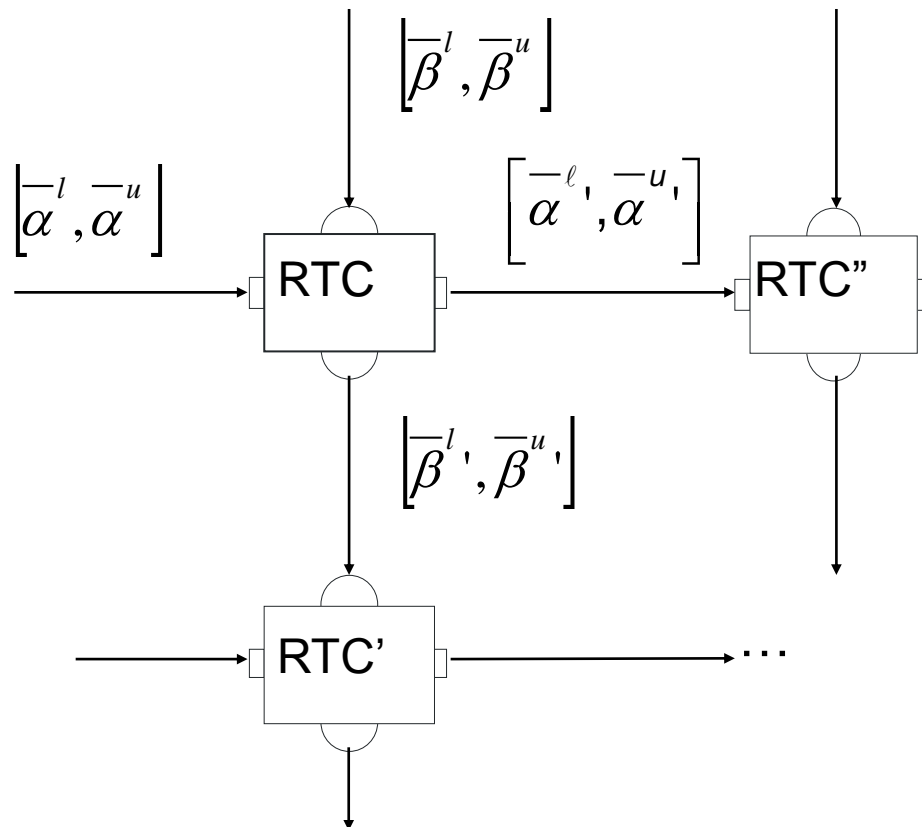
Upper and lower bounds on the number of events


$$\bar{\beta}^u(\Delta) = (\gamma^l)^{-1}(\beta^u(\Delta)) \qquad \bar{\beta}^l(\Delta) = (\gamma^u)^{-1}(\beta^l(\Delta))$$

RTC/MPA:

System of real time components

Incoming event streams and available capacity are transformed by real-time components:



Theoretical results allow the computation of properties of outgoing streams 

Resulting arrival curves:

$$\bar{\alpha}'^u = \min \left(\left[\left(\bar{\alpha}^u \underline{\otimes} \bar{\beta}^u \right) \underline{\oplus} \bar{\beta}^l \right], \bar{\beta}^u \right)$$
$$\bar{\alpha}'^l = \min \left(\left[\left(\bar{\alpha}^l \underline{\oplus} \bar{\beta}^u \right) \underline{\otimes} \bar{\beta}^l \right], \bar{\beta}^l \right)$$

Remaining service curves:

$$\bar{\beta}'^u = \left(\bar{\beta}^u - \bar{\alpha}^l \right) \underline{\oplus} 0$$
$$\bar{\beta}'^l = \left(\bar{\beta}^l - \bar{\alpha}^u \right) \underline{\otimes} 0$$

Where:

$$(f \underline{\otimes} g)(t) = \inf_{0 \leq u \leq t} \{f(t-u) + g(u)\} \quad (f \bar{\otimes} g)(t) = \sup_{0 \leq u \leq t} \{f(t-u) + g(u)\}$$

$$(f \underline{\oplus} g)(t) = \inf_{u \geq 0} \{f(t+u) - g(u)\} \quad (f \bar{\oplus} g)(t) = \sup_{u \geq 0} \{f(t+u) - g(u)\}$$

RTC/MPA: Remarks

- Details of the proofs can be found in relevant references
- Results also include bounds on buffer sizes and on maximum latency.
- Theory has been extended into various directions, e.g. for computing remaining battery capacities

Application: In-Car Navigation System

Car radio with navigation system

User interface needs to be responsive

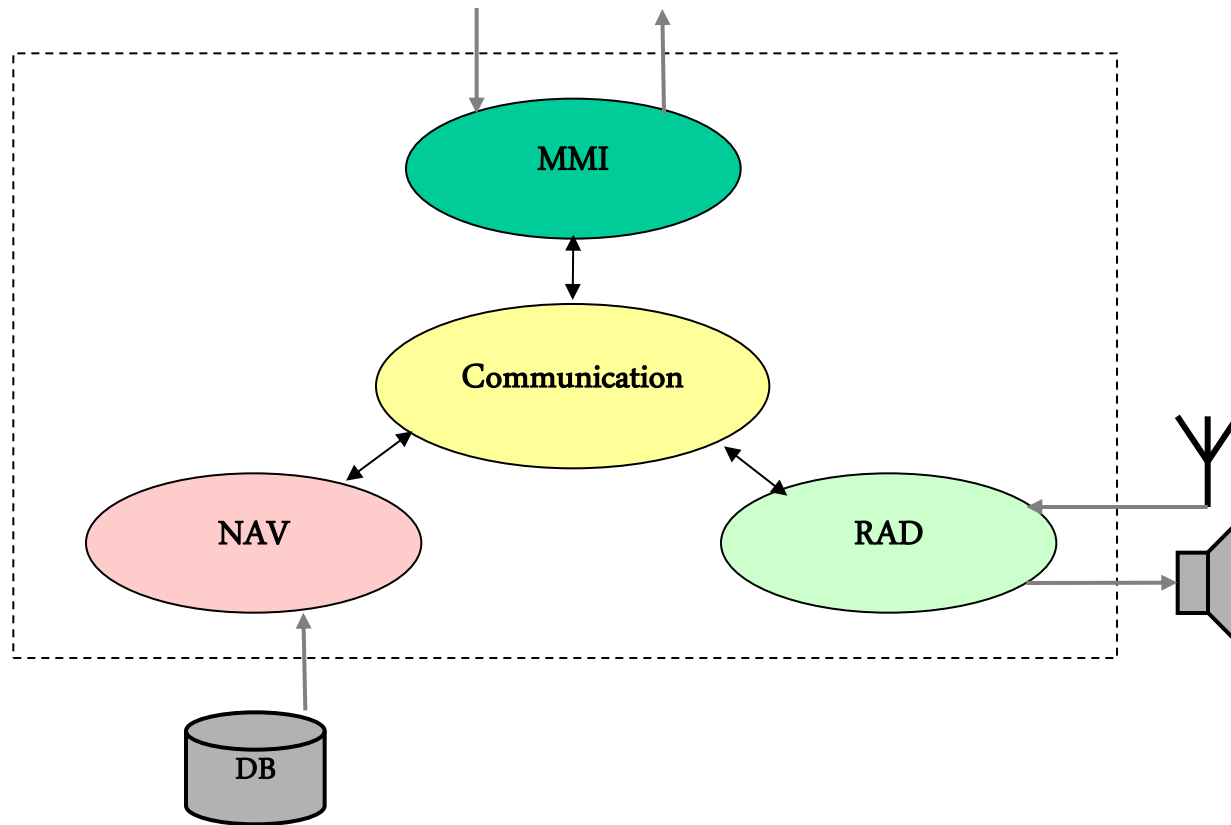
Traffic messages (TMC) must be processed in a timely way

Several applications may execute concurrently

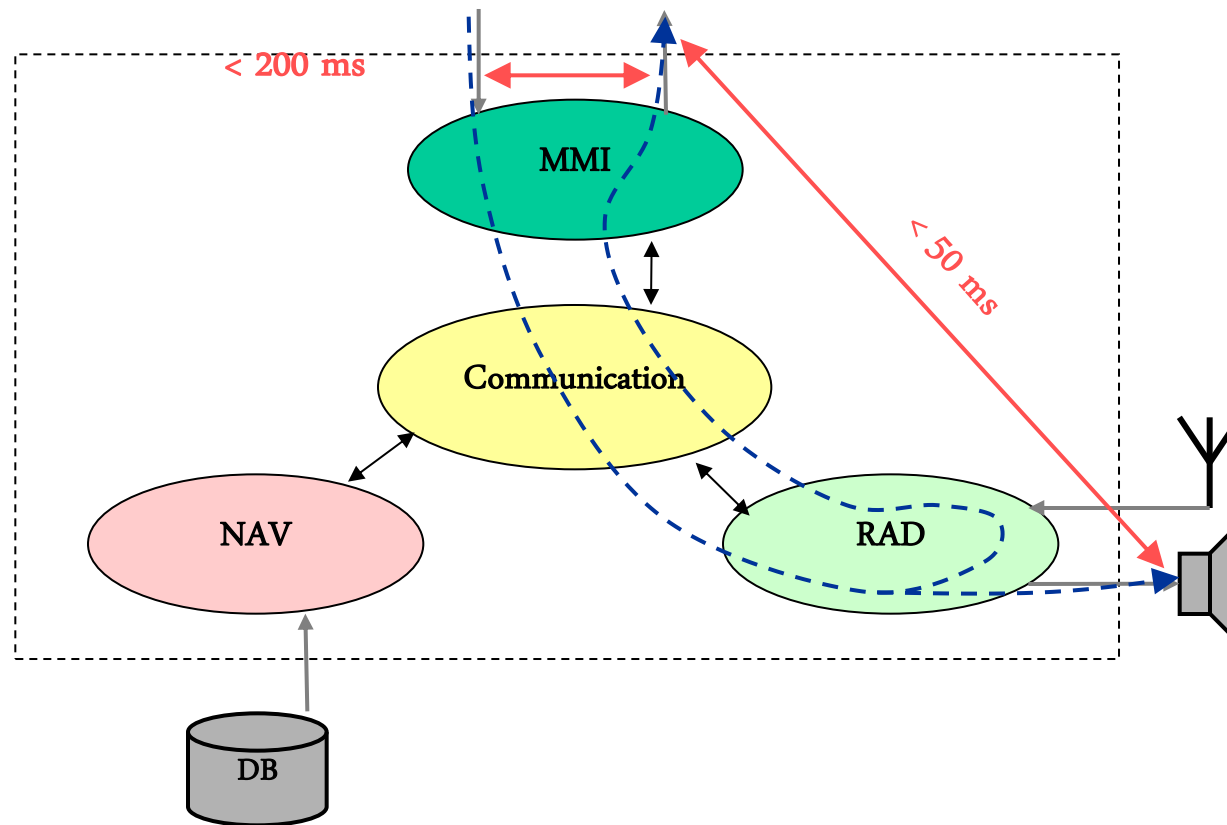


© Thiele, ETHZ

System Overview

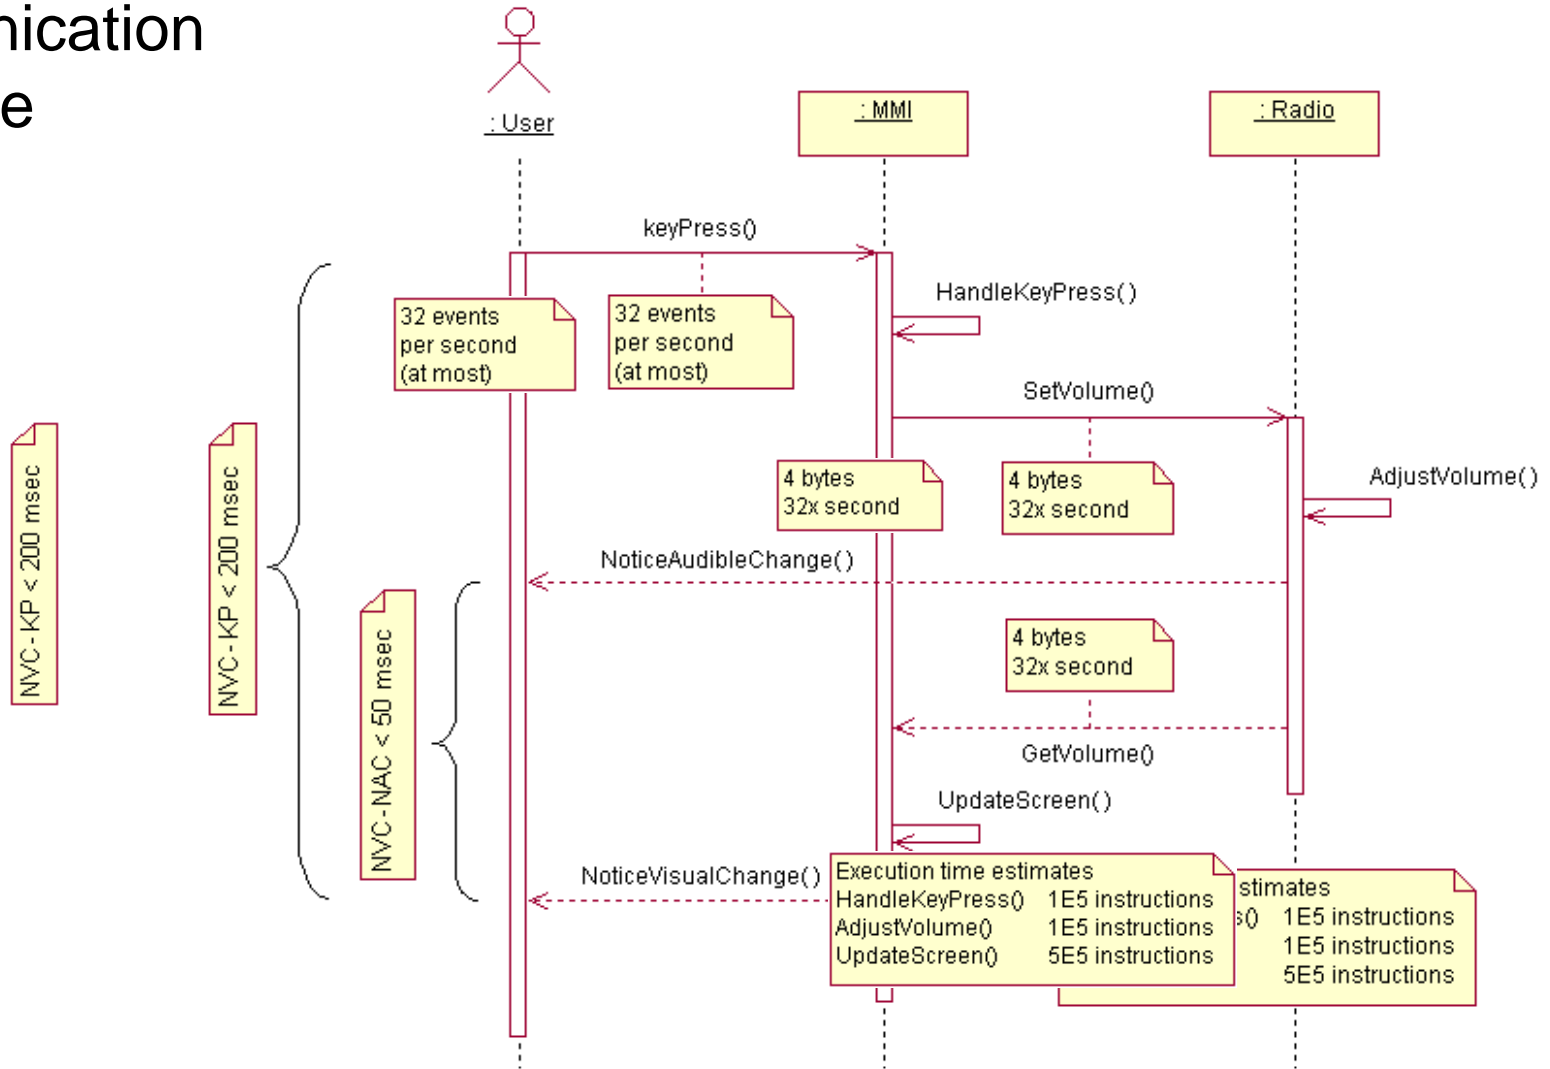


Use case 1: Change Audio Volume

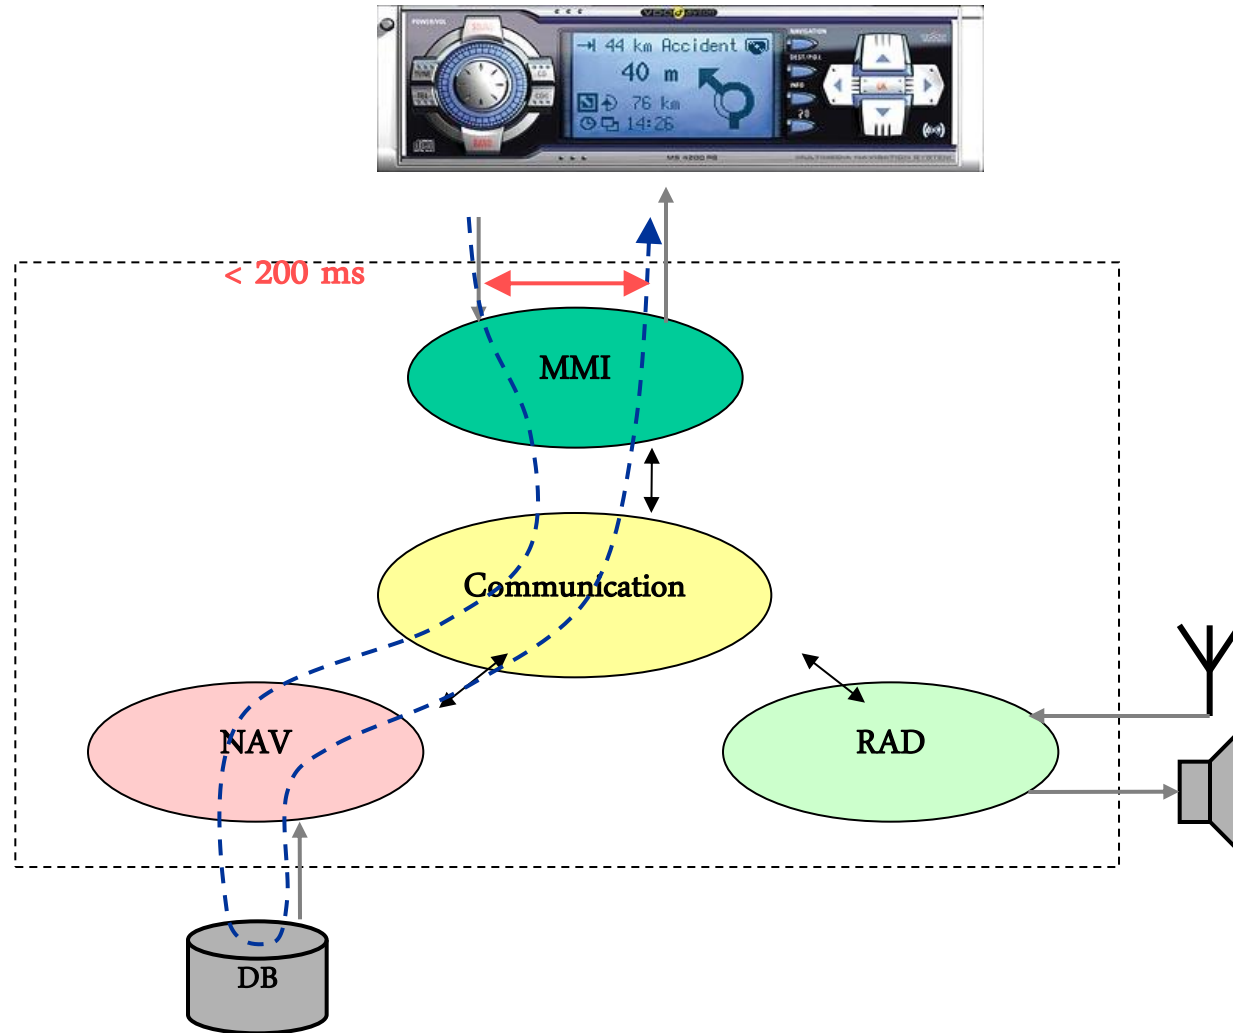


Use case 1: Change Audio Volume

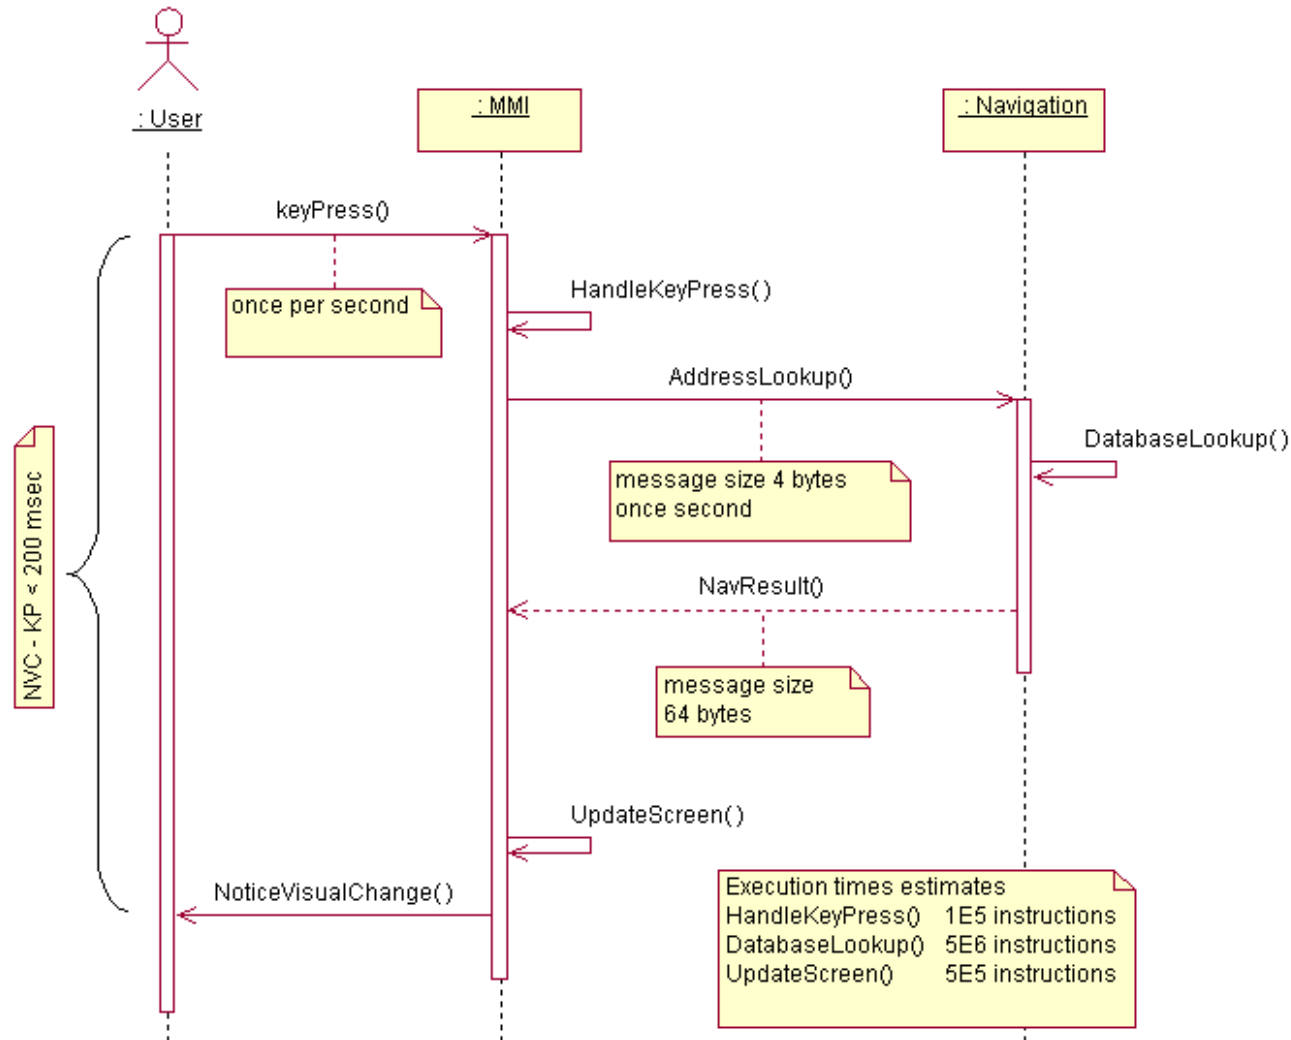
Communication
Resource
Demand




Use case 2: Lookup Destination Address



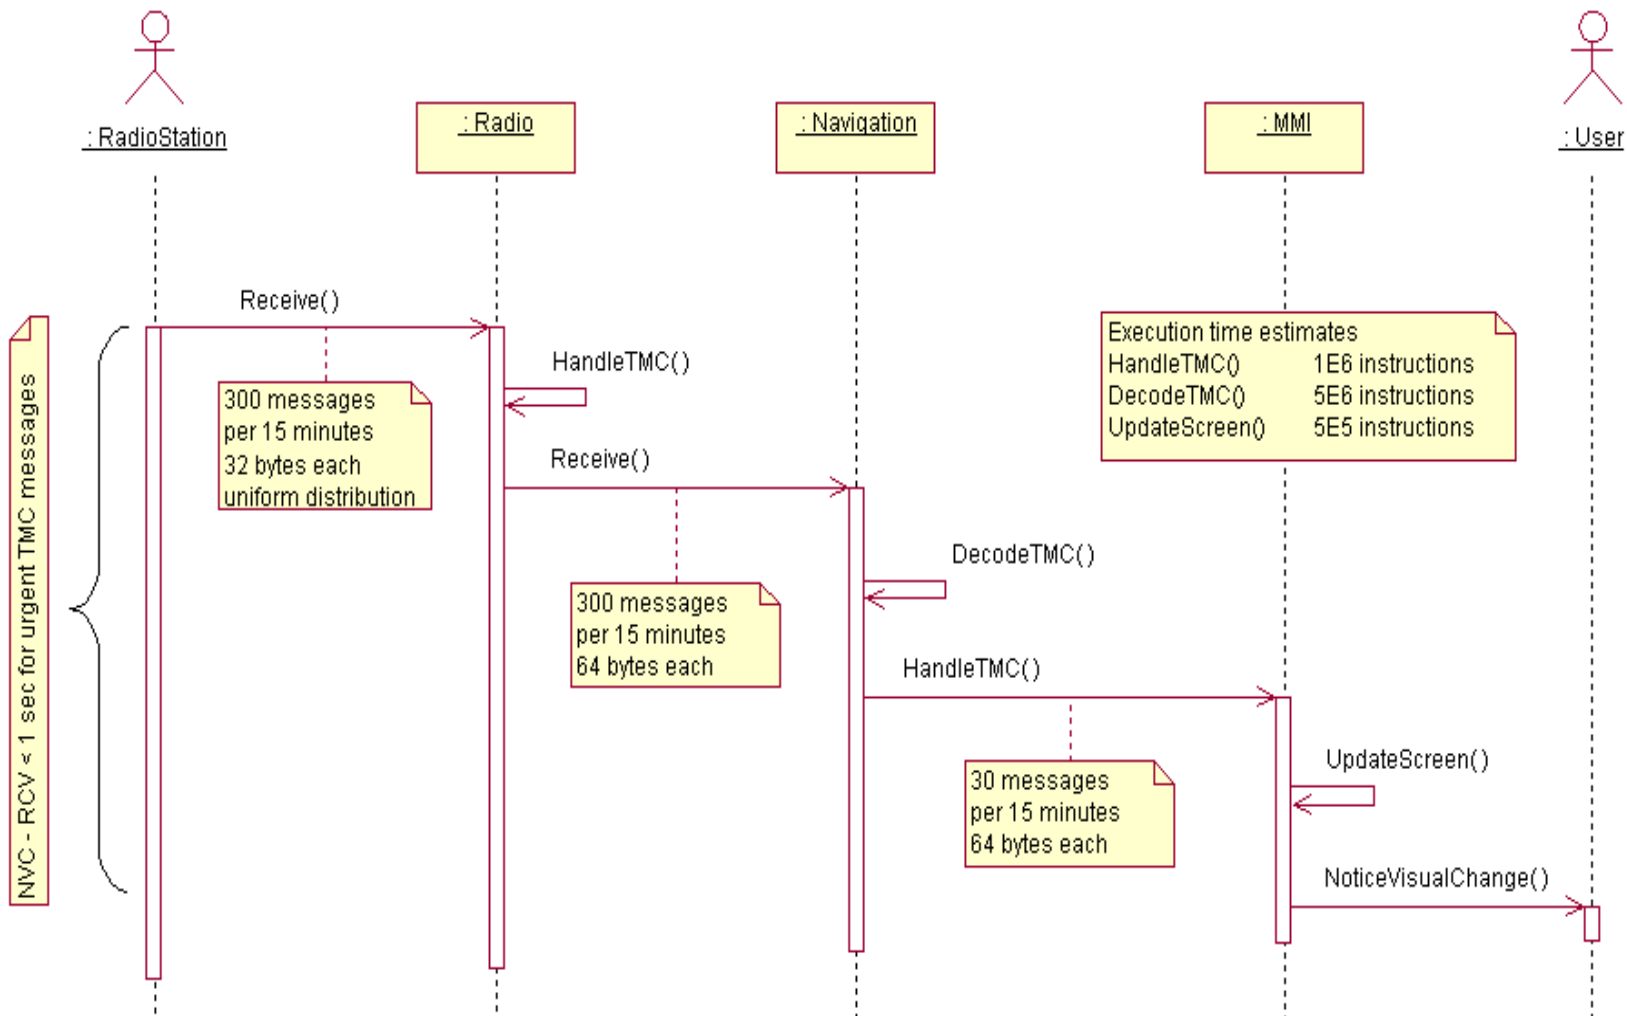
Use case 2: Lookup Destination Address



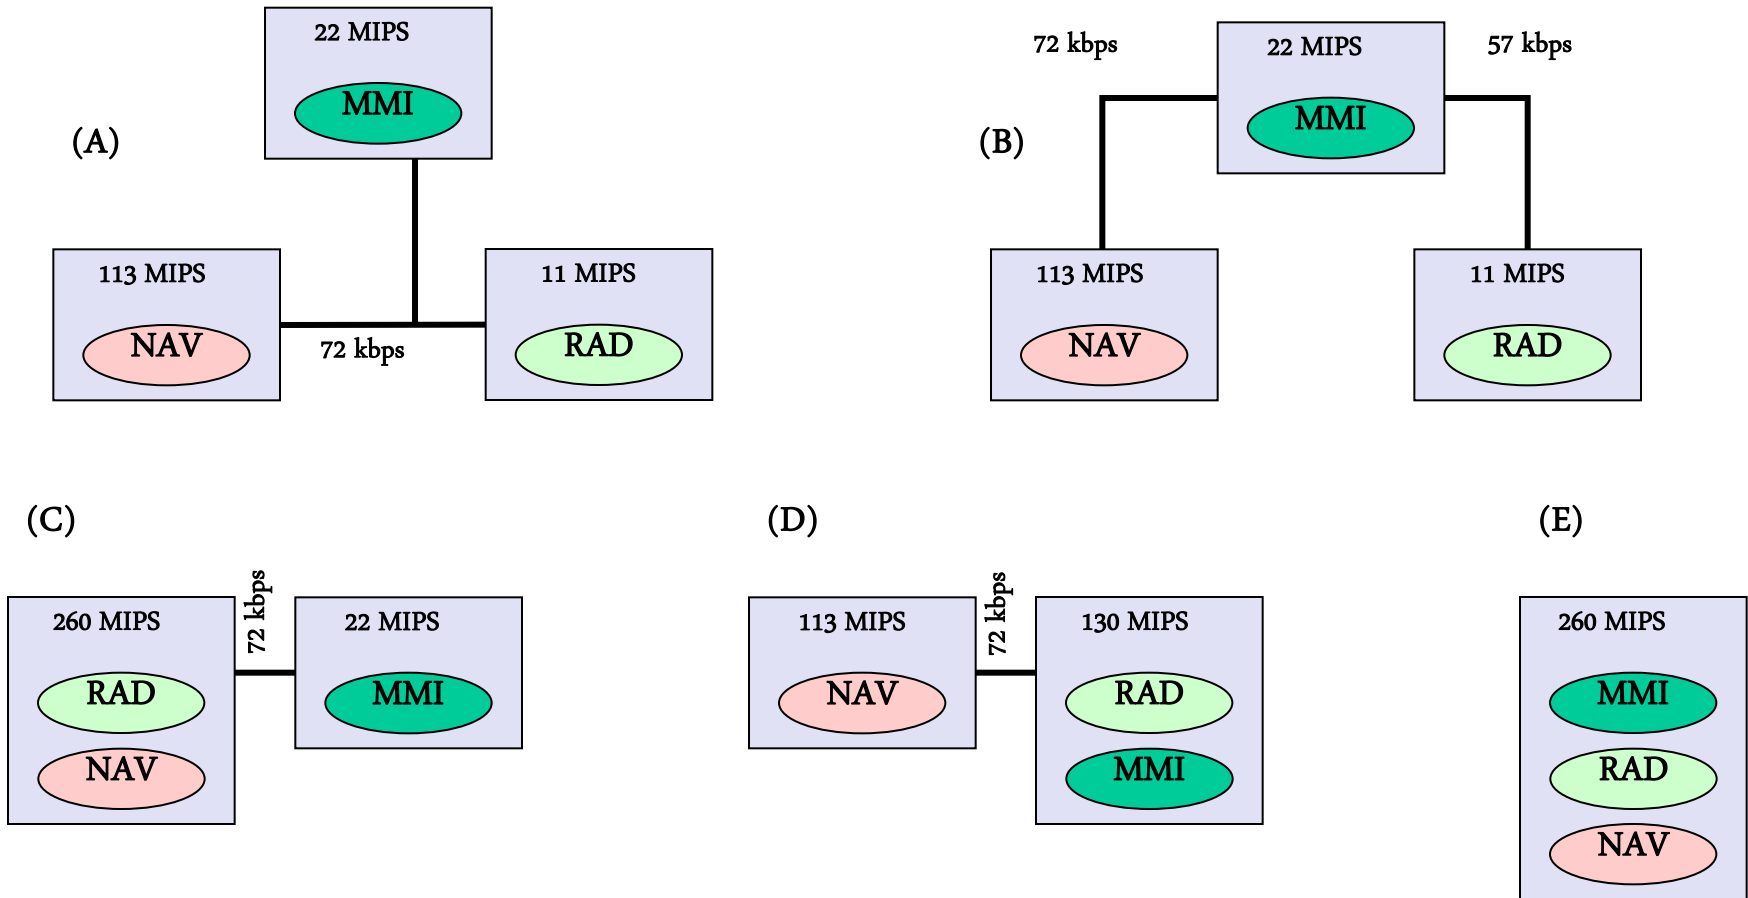
Use case 3: Receive TMC Messages



Use case 3: Receive TMC Messages




Proposed Architecture Alternatives



Step 1: Environment (Event Steams)

Event Stream Model


e.g. Address Lookup
(1 event / sec)



Step 2: Architectural Elements


Event Stream Model

e.g. Address Lookup
(1 event / sec)



Resource Model

e.g. unloaded RISC CPU
(113 MIPS)

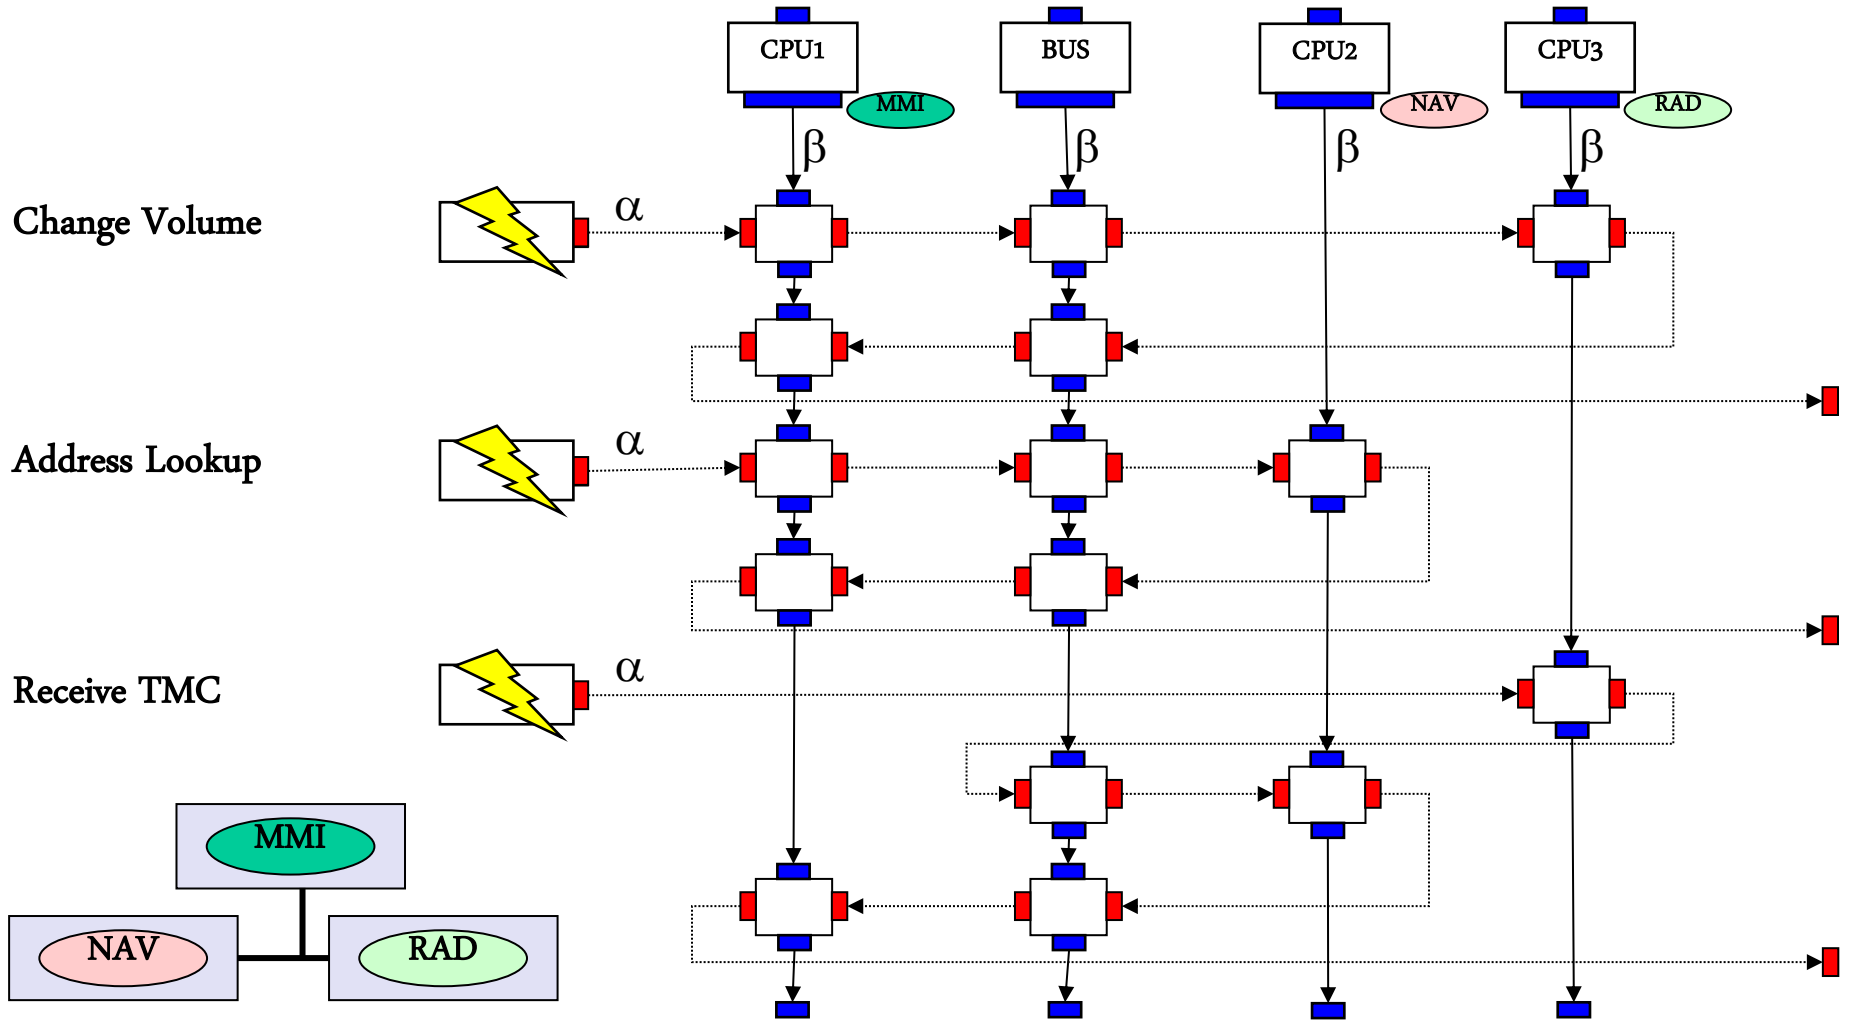


Step 3: Mapping / Scheduling

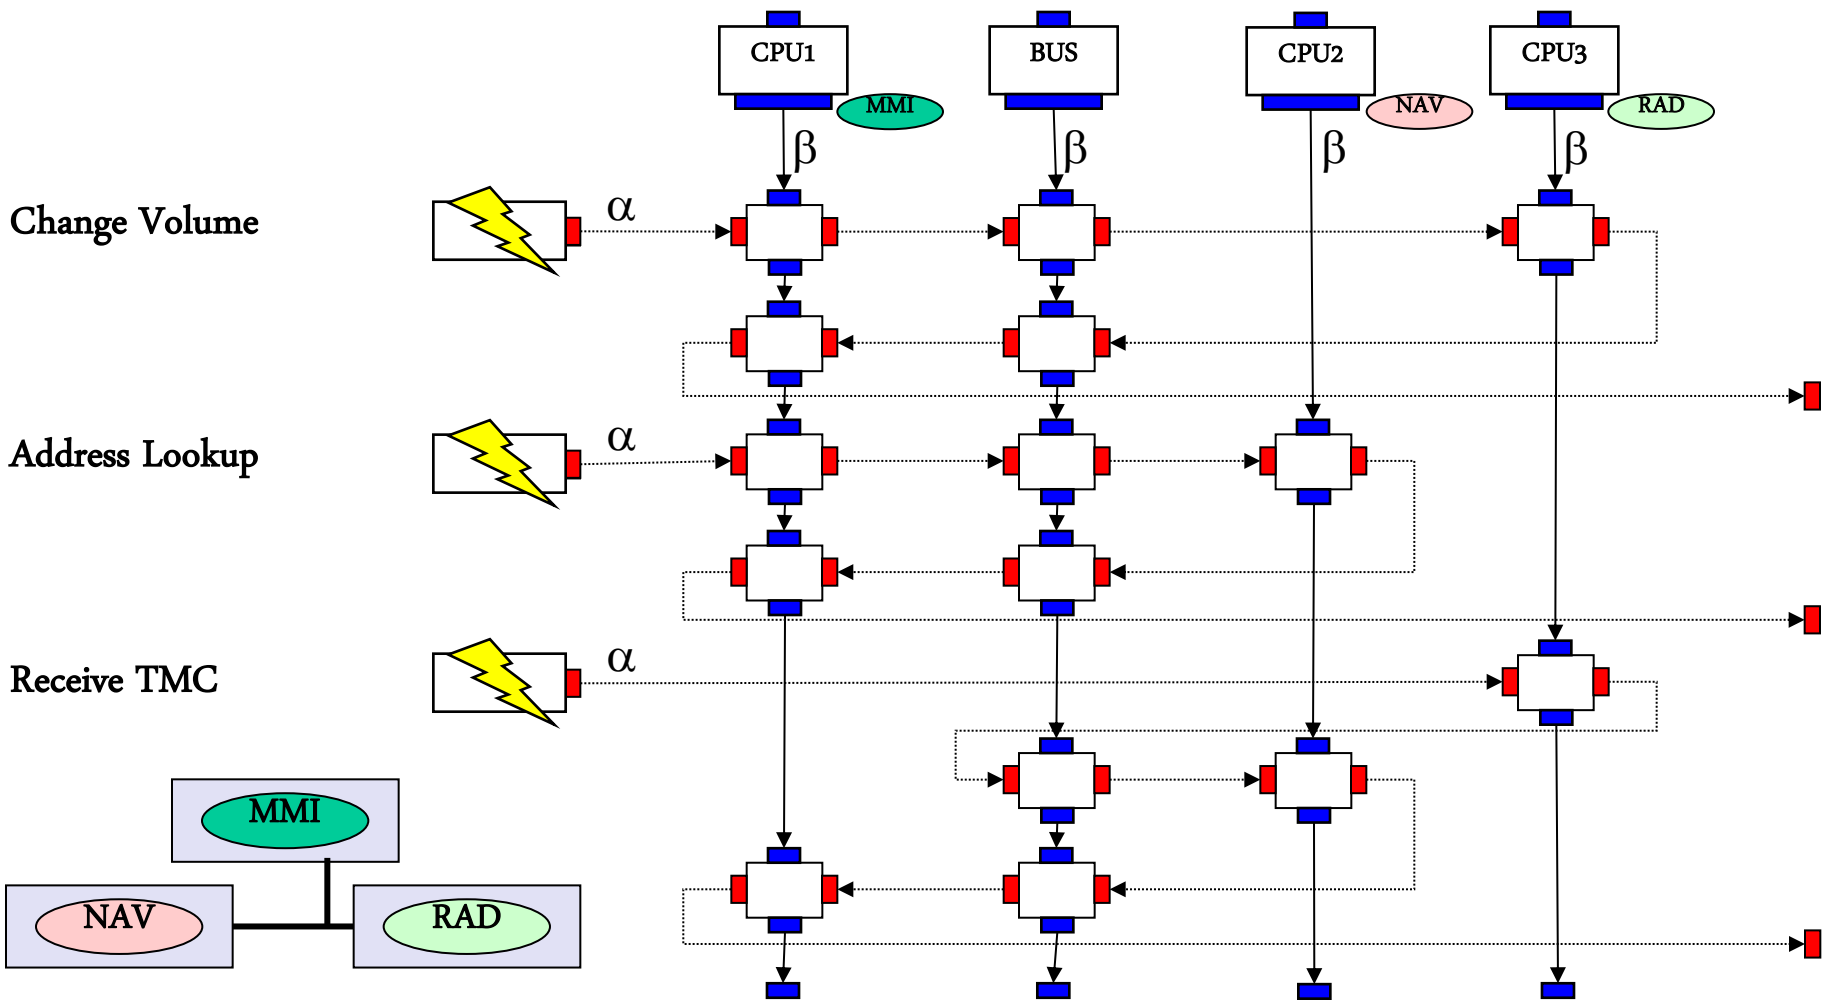
Rate Monotonic Scheduling
(Pre-emptive fixed priority scheduling):

- Priority 1: Change Volume (p=1/32 s)
- Priority 2: Address Lookup (p=1 s)
- Priority 3: Receive TMC (p=6 s)

Step 4: Performance Model



Step 5: Analysis








Analysis – Design Question 1

How do the proposed system architectures compare in respect to end-to-end delays?

Analysis – Design Question 1

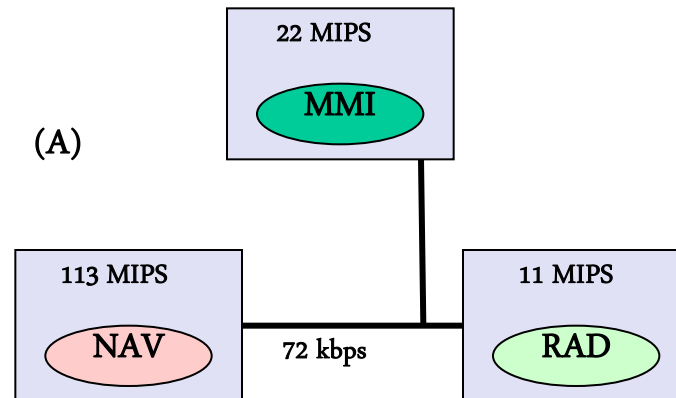
End-to-end delays:



Analysis – Design Question 2

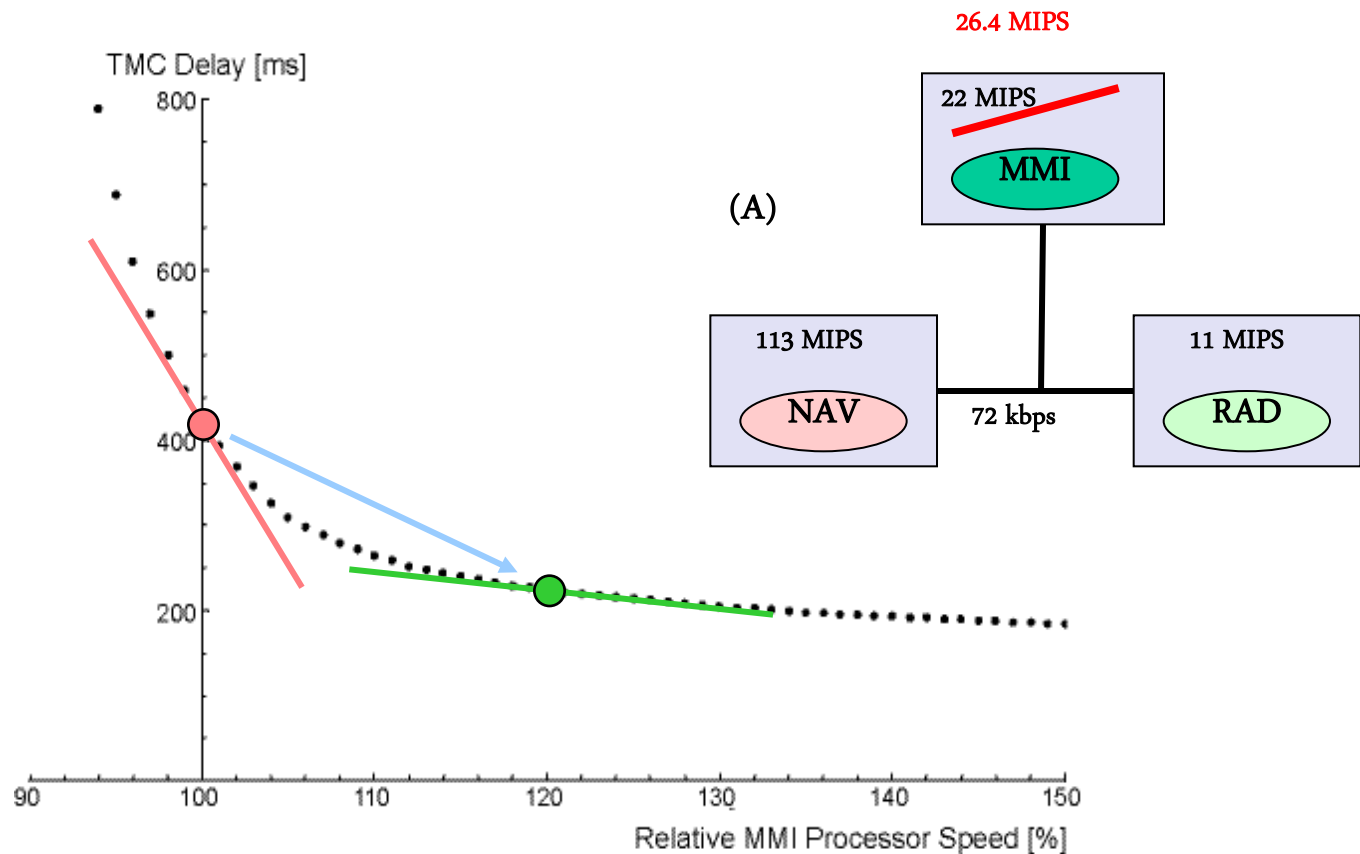
How robust is architecture A?

Where is the bottleneck of this architecture?



Analysis – Design Question 2

TMC delay vs. MMI processor speed:



Conclusions

- Easy to construct models (~ half day)
- Evaluation speed is fast and linear to model complexity (~ 1s per evaluation)
- Needs little information to construct early models (Fits early design cycle very well)
- Even though involved mathematics is very complex, the method is easy to use (Language of engineers)

How to evaluate designs according to multiple criteria?

Many different criteria are relevant for evaluating designs:

- Average & worst case delay
- power/energy consumption
- thermal behavior
- reliability, safety, security
- cost, size
- weight
- EMC characteristics
- radiation hardness, environmental friendliness, ..



How to compare different designs?
(Some designs are “better” than others)

Average vs. worst case energy consumption

- The **average energy consumption** E_{AV} is based on the consumption for selected sets of input data (which?)
- The **worst case energy consumption** E_{WC} is a safe upper bound on the energy consumption
- The worst case usage pattern for the battery is \neq from the worst case for the overall energy consumption
- In general, the pattern for worst case energy consumption is \neq from the worst case thermal pattern


Evaluation of energy consumption: Challenges

- Energy consumption hardly predictable from the source code, due to difficult to predict impact of compiler & linker
- Small variations of the code can lead to large variations of the energy consumption
 - ex. notorious examples
 - Example: shifting code in memory by one byte
- Energy consumption must be predicted from executable code (like the *WCET*)
- The energy consumption might even depend very much on which instance of the hardware is used



Energy models


- Measurements: (potentially) precise, fixed architecture
- Models: flexible architecture, less precise
- Combined models



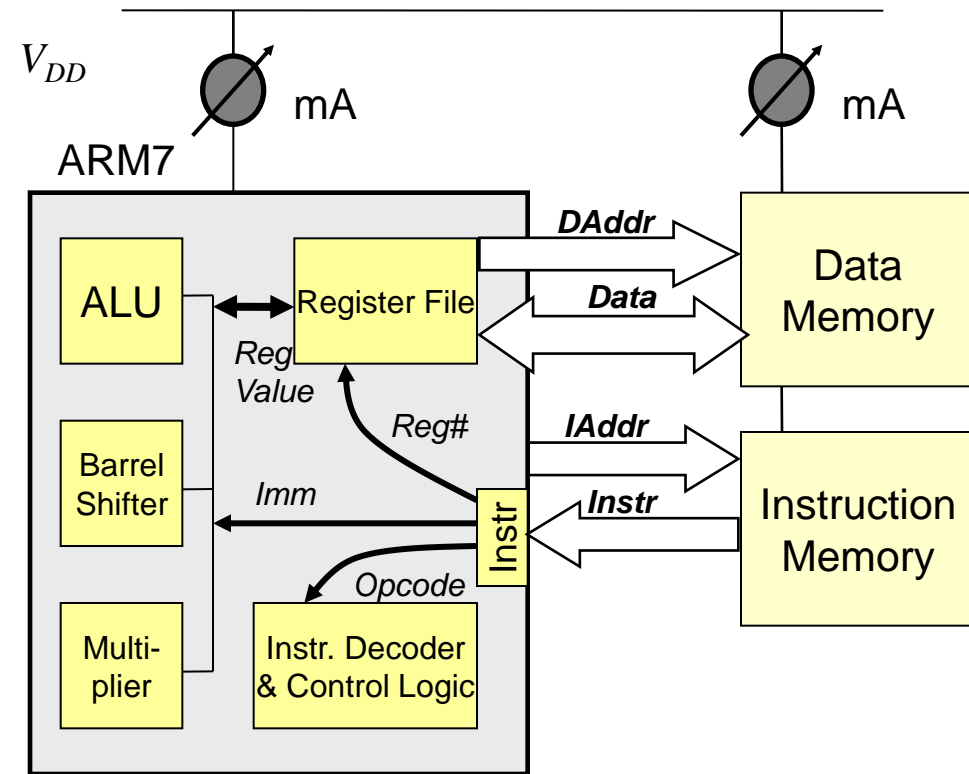
In general, accuracy remains a problem

- Currents difficult to measure

Steinke's model



E.g.: ATMELEVAL board with ARM7TDMI and ext. SRAM



$$E_{total} = E_{cpu_instr} + E_{cpu_data} + E_{mem_instr} + E_{mem_data}$$

S. Steinke, M. Knauer, L. Wehmeyer, P. Marwedel: An Accurate and Fine Grain Instruction-Level Energy Model Supporting Software Optimizations, Int. Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS), 2001

Example: Instruction dependent costs in the CPU

Cost for a sequence of m instructions

$$\begin{aligned} E_{cpu_instr} = & \sum MinCostCPU(\mathbf{Opcode}_i) + FUCost(\mathbf{Instr}_{i-1}, \mathbf{Instr}_i) + \\ & \alpha_1 * \sum w(\mathbf{Imm}_{i,j}) \quad + \beta_1 * \sum h(\mathbf{Imm}_{i-1,j}, \mathbf{Imm}_{i,j}) + \\ & \alpha_2 * \sum w(\mathbf{Reg}_{i,k}) \quad + \beta_2 * \sum h(\mathbf{Reg}_{i-1,k}, \mathbf{Reg}_{i,k}) + \\ & \alpha_3 * \sum w(\mathbf{RegVal}_{i,k}) + \beta_3 * \sum h(\mathbf{RegVal}_{i-1,k}, \mathbf{RegVal}_{i,k}) + \\ & \alpha_4 * \sum w(\mathbf{IAddr}_i) \quad + \beta_4 * \sum h(\mathbf{IAddr}_{i-1}, \mathbf{IAddr}_i) \end{aligned}$$


w : number of ones;

h : Hamming distance;


$FUCost$: cost of switching functional units;

α, β : determined through experiments.

Hamming Distance between adjacent addresses is playing major role



Energy-efficient execution on graphics processor (GPU)




current clamp

Energy per frame CPU	3.26 J	5.84 J	10.52 J	Reduced to
Energy per frame GPU	0.93 J	1.56 J	2.76 J	avg 27%

C. Timm, A. Gelenberg, P. Marwedel, F. Weichert: Energy Considerations within the Integration of General Purpose GPUs in Embedded Systems. Intern. Conf. on Advances in Distributed and Parallel Computing, 2010


C. Timm, F. Weichert, P. Marwedel, H. Müller: Design Space Exploration Towards a Realtime and Energy-Aware GPGPU-based Analysis of Biosensor Data. Computer Science - Research and Development, ENA-HPC, 2011

CACTI model



Cache model used

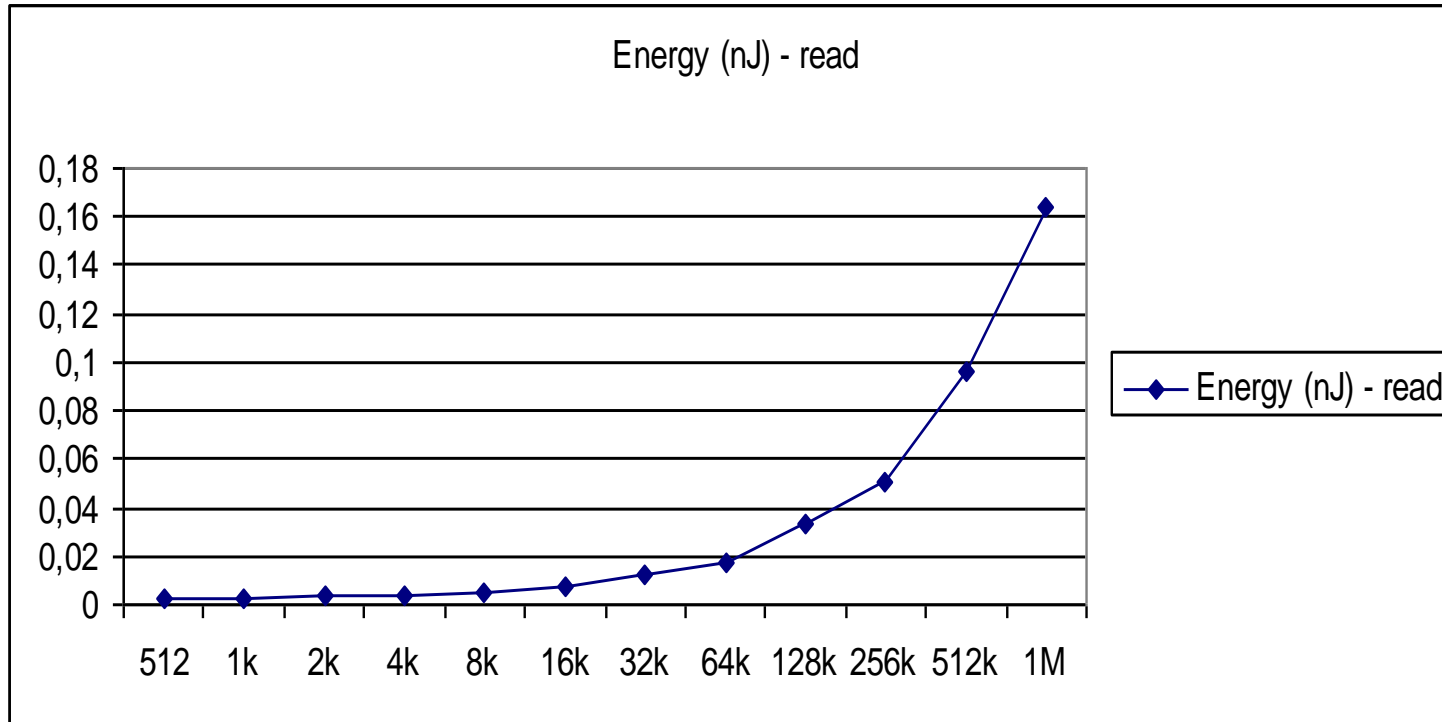
Comparison with SPICE



<http://research.compaq.com/wrl/people/jouppi/CACTI.html>

Energy consumption of memories

Example: CACTI / high performance Scratchpad (SRAM):




16 bit read; size in bytes; 65 nm technology

Source: Olivera Jovanovic, TU Dortmund, 2011

Energy consumption of memories (2)

Example CACTI: Scratchpad (SRAM) vs. DRAM (DDR2):



16 bit read; size in bytes;
65 nm for SRAM, 80 nm for DRAM

Source: Olivera Jovanovic,
TU Dortmund, 2011

DRAM power

Complex DRAM models:

- <http://www.micron.com/products/support/power-calc>
- T. Vogelsang: Understanding the Energy Consumption of Dynamic Random Access Memories, Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture, pp. 363—374, <http://dx.doi.org/10.1109/MICRO.2010.42>


Steinke's "combined" model

- Measured values for the processor
- Model-based values for memories
(validated against existing measurements)

Examples of energy models

- Measurements:
 - Tiwari (1994): Energy consumption within processors
 - Russell, Jacome (1998): Measurements for 2 fixed configurations
 - Simunic (1999): Values from data sheets. Not very precise.
 - Timm: measurements for graphics card
- Models:
 - CACTI [Jouppi, 1996]: Predicted energy consumption of caches
 - Wattch [Brooks, 2000]: Power estimation at the architectural level, without circuit or layout, known to be imprecise
- Combined models
 - Steinke et al., TU Dortmund (2001): mixed model

Worst case energy consumption via worst case computing time?



- Computing the E_{WC} using WCET estimations

$$E_{WC} = \int_0^{WCET_{EST}} P(t) dt$$

- Tight bounds if $P(t)$ has small variations & $WCET_{EST}$ is tight
- Little value if $P(t)$ varies too much.

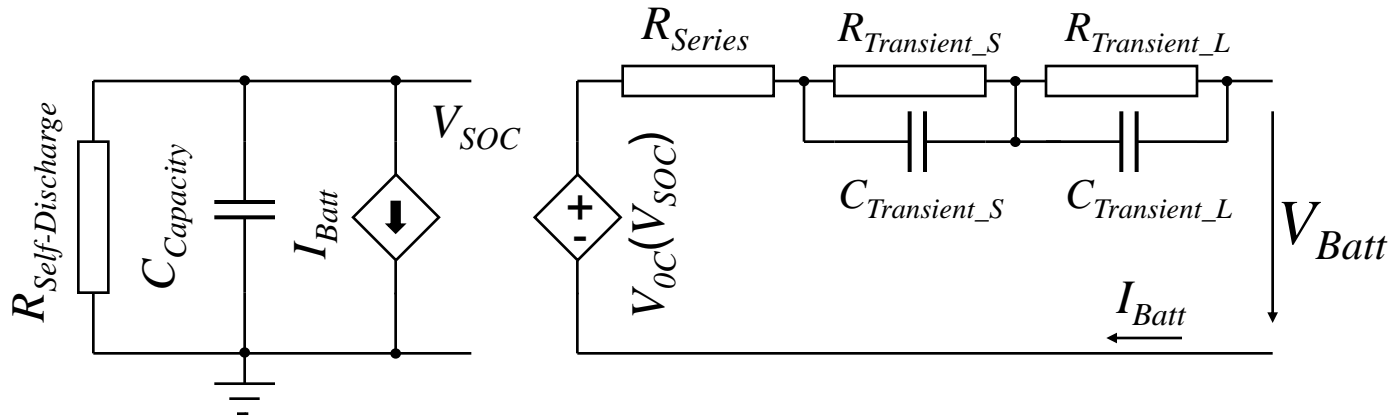
Battery models

- (Chemical) & physical models
e.g. concentrated solution theory, partial differential eq.s
many, frequently unknown parameters (50+); xy hours simulation time
- Empirical models
Simple equations, inaccurate
 - Peukert's law: lifetime = C/I^α , with empirical α
 - Weibull fit
- Abstract models
 - • Electrical circuit models
 - Discrete time model (e.g. in VHDL)
 - Stochastic models (e.g. Markov processes)
- Mixed models
e.g. electrical models with physical explanation

frequently with fitting

Sources include: R. Rao, S. Vrudhula, D.N. Rakhmatov: Battery Modeling for Energy-Aware System Design, IEEE Computer, 2003, pp. 77

Model proposed by Chen and Rincón-Mora

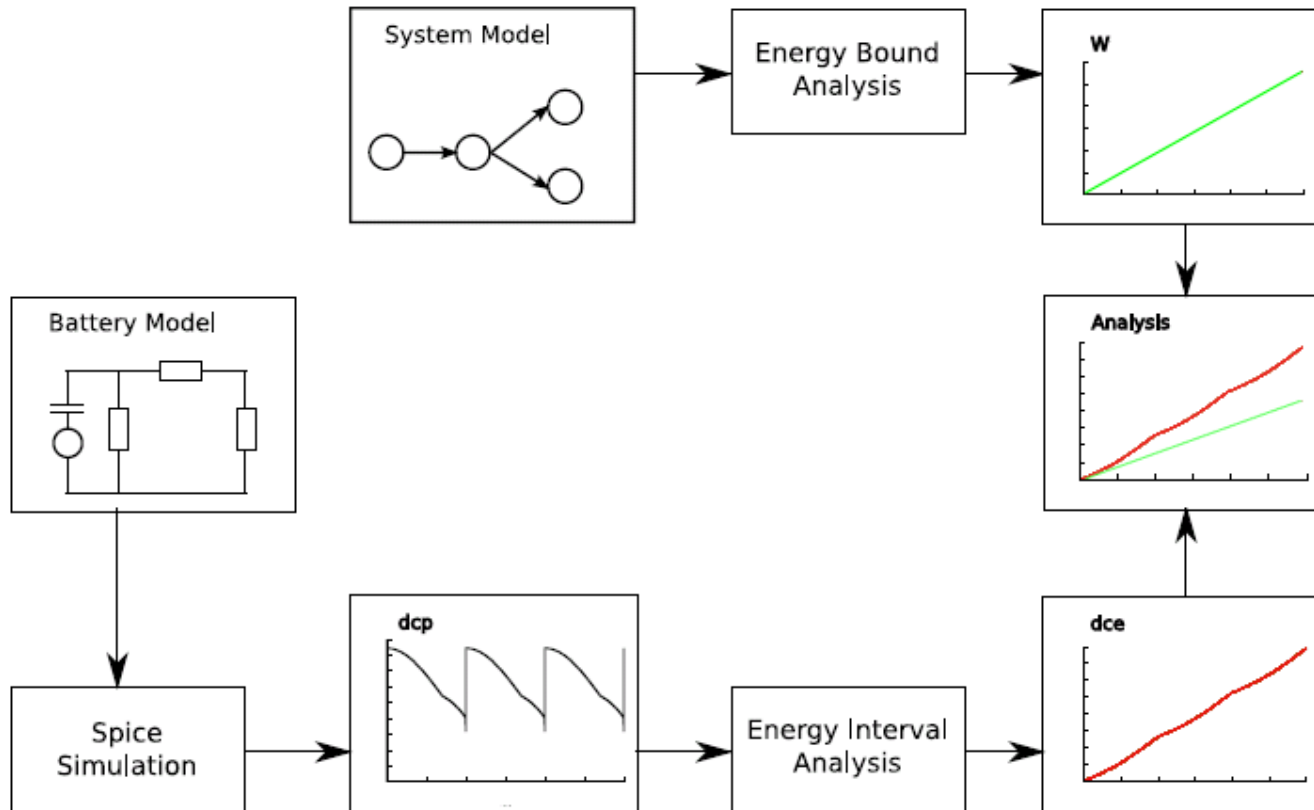


Source: M. Chen, G. A. Rincón-Mora: Accurate Electrical Battery Model Capable of Predicting Runtime and I - V Performance, *IEEE Trans. on Energy Conversion*, 2006, pp. 504

- Full charge capacitor: $C_{Capacity} = 3600 \cdot Capacity \cdot f_1(\text{cycle}) \cdot f_2(\text{Temp})$
- Self-discharge resistor: $R_{Self-Discharge}$ (might depend on parameters)
- Current dependency of V_{Batt} : modeled by $R_{series} + R_{Transient_S} + R_{Transient_L}$
- I_{Batt} charges and discharges $C_{Capacity}$
- Voltage controlled voltage source V_{OC} captures nonlinear dependency between the state of charge and V_{OC} (measurement can take days)
- R_{Series} : models immediate voltage drop at load change

Battery capacity sufficient?

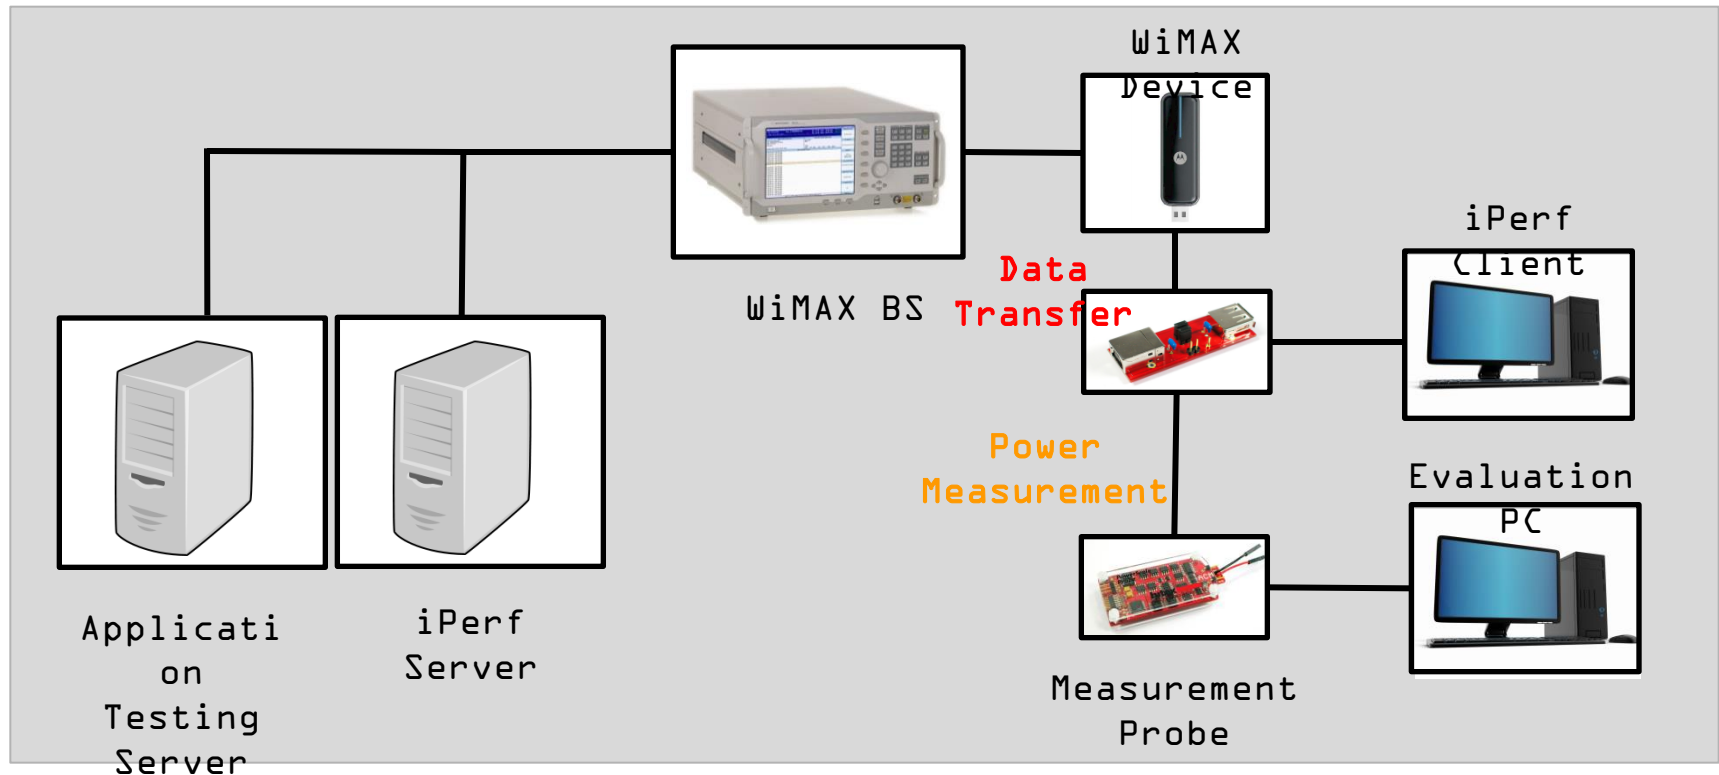
Question can be solved with adapted real-time calculus



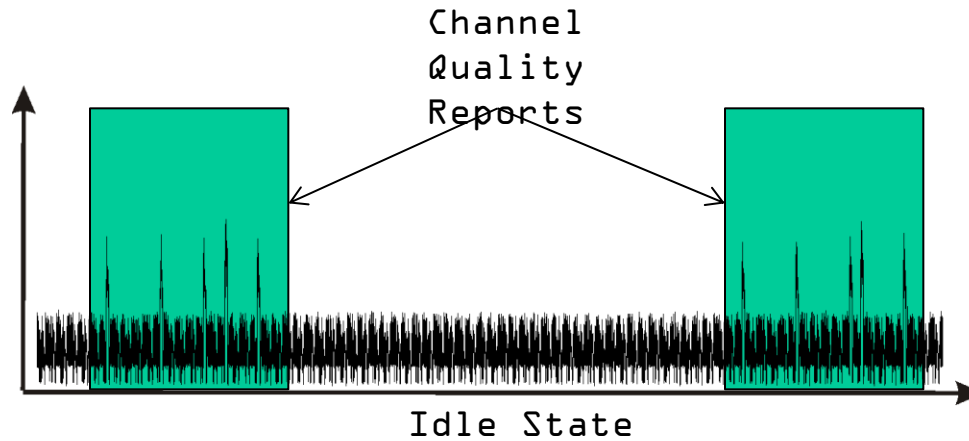
Lipskoch, H., Albers, K. and Slomka, F.: Battery discharge aware energy feasibility analysis, Proceedings of the 4th international Conference on Hardware/Software Codesign and System Synthesis, CODES+ISSS '06, pp. 22-27, 2006.

Energy models for communication: An Energy Model for Mobile WiMAX Devices


- How does the application data rate influence the energy efficiency?
- What is the impact of very small amounts of data on the efficiency?
- Relationship between submitted power and consumed energy?




Traffic Dependent Energy Consumption



- Average Power in Idle State: 880 mW
- Channel Quality Reports every 300 ms lead to increased average power of 930 mW
- Transmission is costing significantly increased energy consumption
- Reception is not increasing the power consumption compared with idle state



Modeling the Impact of T_x -Power Variation

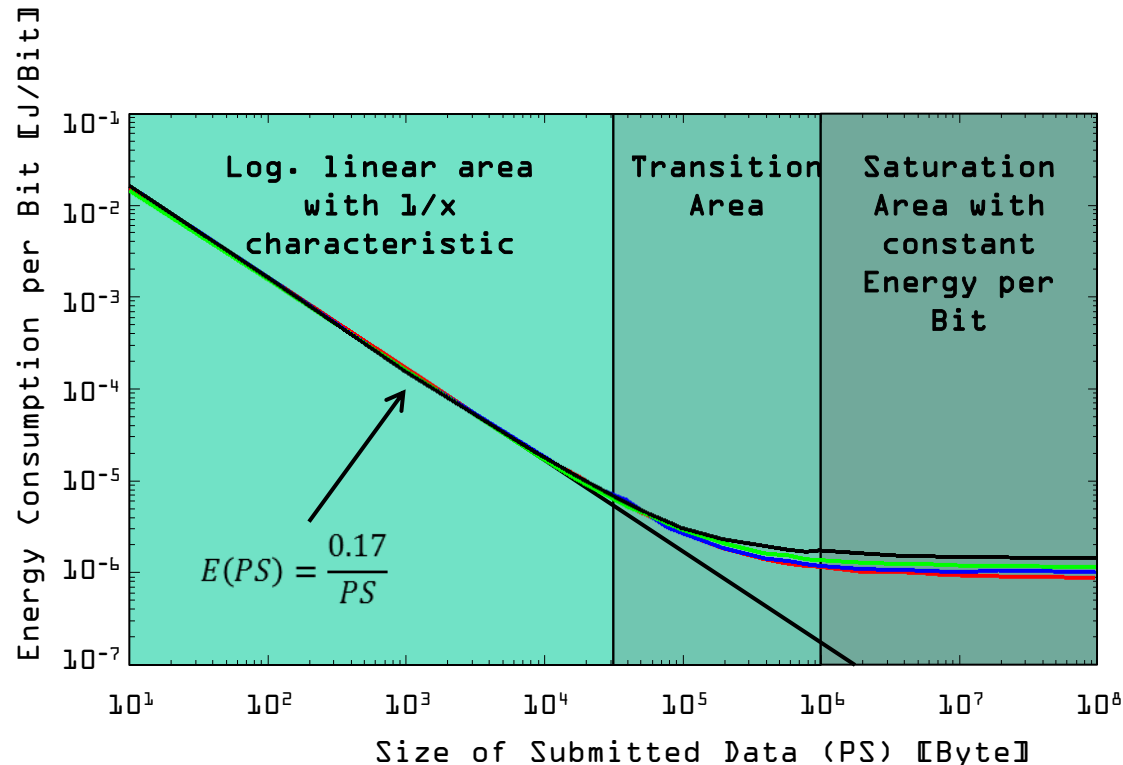


- Energy per Bit is constant for low T_x power (below -25 dBm)
- For higher T_x power, the consumed energy can be approximated by a 2nd degree polynomial
- Significant energy savings can be achieved by using as low power as possible

$$E(P_{Tx}) = \begin{cases} C(\Psi) & , P_{Tx} < P_{Tx,min} \\ \alpha \cdot P_{Tx}^2 + \beta \cdot P_{Tx} + \gamma + C(\Psi) & , P_{Tx} > P_{Tx,min} \end{cases}$$

α	β	γ	$C(\Psi=35/12)$	$C(\Psi=32/15)$	$C(\Psi=29/18)$	$C(\Psi=26/21)$	$P_{Tx,min}$
1.0325e-6	4.6e-5	5.225e-4	1.1084e-3	8.3530e-4	6.916e-4	5.994e-4	-25 dBm

Modeling the Impact of Different File Sizes



The Energy Model for different file sizes can be divided into three parts

Log Linear Area:

Rapidly decreasing energy consumption per Bit for packet sizes below 20 kByte

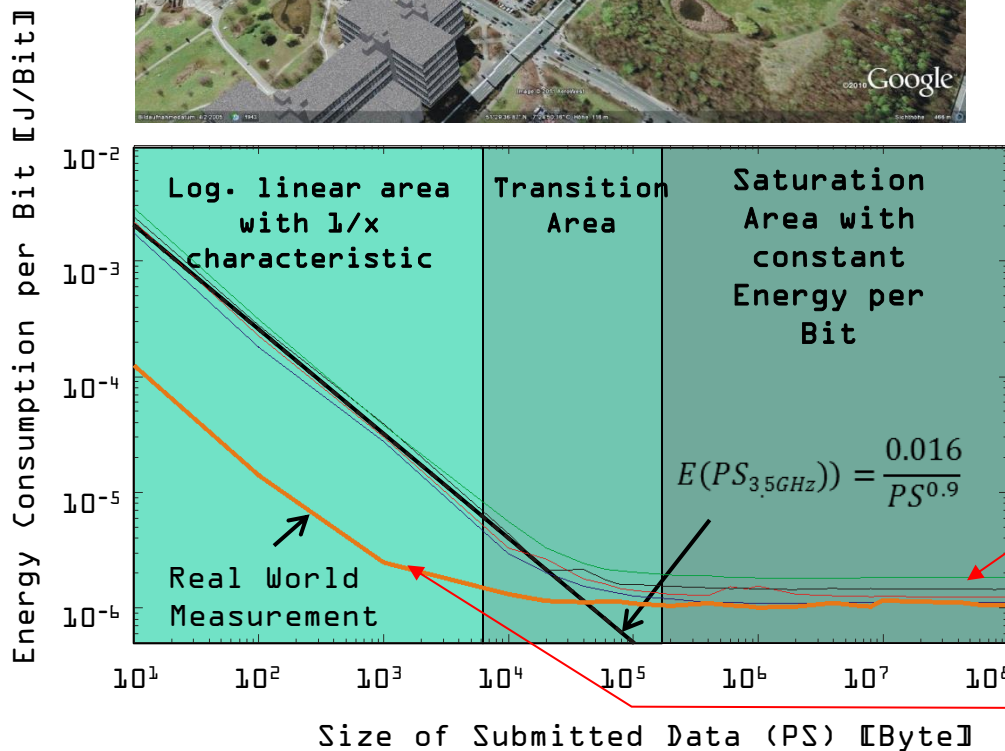

Transition Area:

Transition to constant energy consumption

Saturation Area

For packet sizes above 900 kByte, collecting more data does not make sense from an energy efficiency perspective

Validation for Different Devices




- The overall model is valid for different devices, and different frequency bands
- Chipset specific offsets have to be applied for the log linear area
- Lab equipment consumes some energy after bits have been sent.
- Real (more modern) BS instead of lab equipment is covered by the model (offset has to be applied for low PS)

How to evaluate designs according to multiple criteria?

Many different criteria are relevant for evaluating designs:

- Average & worst case delay
- power/energy consumption
- thermal behavior
- reliability, safety, security
- cost, size
- weight
- EMC characteristics
- radiation hardness, environmental friendliness, ..



How to compare different designs?
(Some designs are “better” than others)

Thermal models


Thermal models becoming increasingly important

- since temperatures become more relevant due to increased performance, and
- since temperatures affect
 - usability and
 - dependability.



Thermal conductivity

$$P_{th} = \kappa \frac{\Delta T \cdot A}{l} \quad (1)$$



Where

- P_{th} : thermal power transferred
- κ : thermal conductivity
- ΔT : temperature difference
- A : area
- l : length

Thermal conductivity κ reflects the amount of thermal energy per unit of time transferred through a plate made of some material of area A and thickness l when the temperatures at the opposite sides differ by one temperature unit (e.g. Kelvin)

Examples of thermal conductivity

Material	Thermal conductivity [W/(m K)]
Copper	240-401
Aluminum (95.5%)	236
Silicon	148
Wood (perpendicular to fibre)	0.09-0.19
Concrete	0.08-0.25
Air (21% oxygen)	0.0262



<http://de.wikipedia.org/wiki/Wärmeleitfähigkeit>

Thermal conductance & resistance

- **Thermal conductance** = amount of thermal energy which passes through a plate per unit of time if the temperatures at the two ends of the plate differ by one unit of temperature (e.g. Kelvin).

$$P_{th} = \kappa \frac{\Delta T \cdot A}{l} \quad (1) \quad \rightarrow \quad \frac{P_{th}}{\Delta T} = \kappa \frac{A}{l} \quad (2)$$

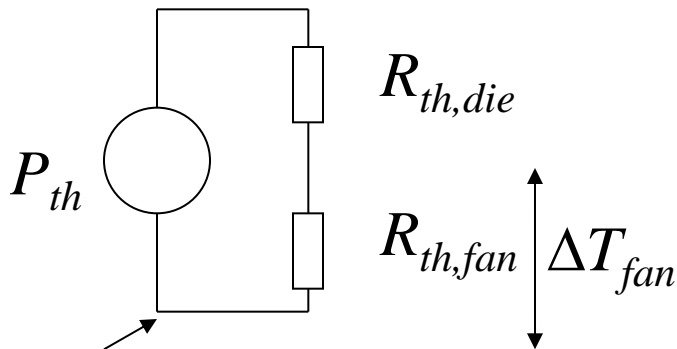
- The reciprocal of thermal conductance is called **thermal resistance** R_{th} .

$$R_{th} = \frac{\Delta T}{P_{th}} = \frac{l}{\kappa \cdot A} \quad (3)$$

Equivalent thermal circuits

- Thermal resistances add up like electrical resistances
- 👉 Thermal modeling mapped to circuit modeling

e.g.: microprocessor:



Ground \approx Reference temperature

$$\Delta T = R_{th} \cdot P_{th} \quad (3)$$

$$R_{th} = R_{th,die} + R_{th,fan} \quad (4)$$

For $R_{th,die} = 0.4$ [W/K],

$R_{th,fan} = 0.3$ [W/K],

$P_{th} = 10$ [W]:

👉 $\Delta T = 7$ [K],

$\Delta T_{fan} = 3$ [K]

So far, we have just considered the steady state.


Dynamic thermal properties

In general, transients and thermal capacity to be considered:

$$C_{th} = m \cdot c$$

where C_{th} : thermal capacity, m : mass, c : specific heat

☞ Networks comprising resistances and capacities



Extra voltage source can make reference temperature explicit


<http://www.infineon.com/dgdl/smdpack.PDF?folderId=db3a304412b407950112b417b3e623f4&fileId=db3a304412b407950112b417b42923f5>

Equivalences

Electrical model		Thermal model	
Current	I	Thermal flow, flow of “power”	$P_{th} = \dot{Q}$
Total charge	$Q = \int I dt$	Thermal energy	$E_{th} = \int P_{th} dt$
Resistance	R	Thermal resistance	R_{th}
Potential	φ	Temperature	T
Voltage = potential difference	$U = \Delta\varphi$	Temperature difference	ΔT
Capacitance	C	Thermal capacitance	C_{th}
Ohm’s law	$U = R I$	Δ Temperature at R_{th}	$\Delta T = R_{th} \cdot P_{th}$

Examples of thermal resistance of P-TO263-7-3

Component	Value & Dimension
Thermal resistance of chip	0.48 [K/W]
Thermal time constant of chip	≈1.5 ms
Thermal capacity of chip	≈3 [mWs/K]
Thermal resistance of heat slug	0.24 [K/W]
Thermal capacity of heat slug	310 [mWs/K]
Thermal time constant of heat slug	70 [ms]




<http://www.infineon.com/dgdl/smdpack.PDF?folderId=db3a304412b407950112b417b3e623f4&fileId=db3a304412b407950112b417b42923f5>

Hotspot –

A popular thermal simulator for processors

- Localized heating much faster than chip-wide (millisec time scale)
- Chip-wide treatment is inaccurate (neglects hot spots)
- Temperature is sensitive to chip layout (floorplan)
- 👉 Fine-grained, dynamic model of temperature
- Authors say: Validated against FEM models




(2D model, 2.5 D exists)

<http://lava.cs.virginia.edu/HotSpot/documentation.htm>
Including PowerPoint slides from ISCA 2003

Results of simulations based on thermal models (1)


Encapsulated cryptographic coprocessor:




Source: http://www.coolingzone.com/Guest/News/NL_JUN_2001/Campi/Jun_Campi_2001.html

Results of simulations based on thermal models (2)

Microprocessor



▲ Flomerics image showing the thermal solution with a metal lid.



▲ Flomerics image showing the thermal solution without a metal lid.

Source: http://www.flotherm.com/applications/app141/hot_chip.pdf

Summary

- Thiele's real-time calculus (RTC)/MPA
 - Using bounds on the number of events in input streams
 - Using bounds on available processing capability
 - Derives bounds on the number of events in output streams
 - Derives bound on remaining processing capability, buffer sizes, ...
 - Examples demonstrate design procedure based on RTC
- Energy and power consumption
 - Measurements
 - Models (with calibration)
 - Mixed approaches
 - Energy for computation, storage and communication
- Thermal behavior
 - Mapping to thermal circuit model

Reserve