Year 4 Review
Dresden, March 16th, 2012

Cluster

Achievements and Perspectives:

Operating Systems and Networks

leader: Giorgio Buttazzo

Scuola Superiore Sant'Anna, Pisa, Italy

Outline of the Presentation

- **Objectives**
- **Partners and activities**
- **Achievements**
- **Future plan**

Embedded Systems are becoming more complex, and characterized by dynamic behavior and distributed organization

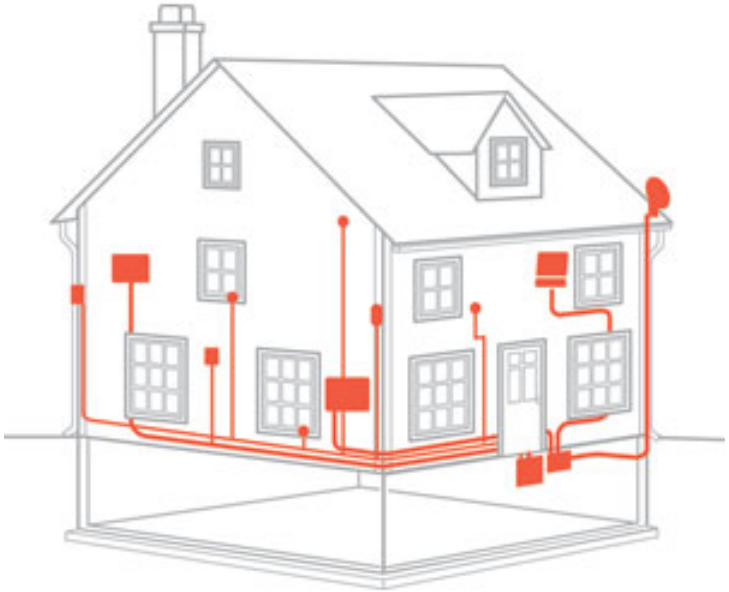

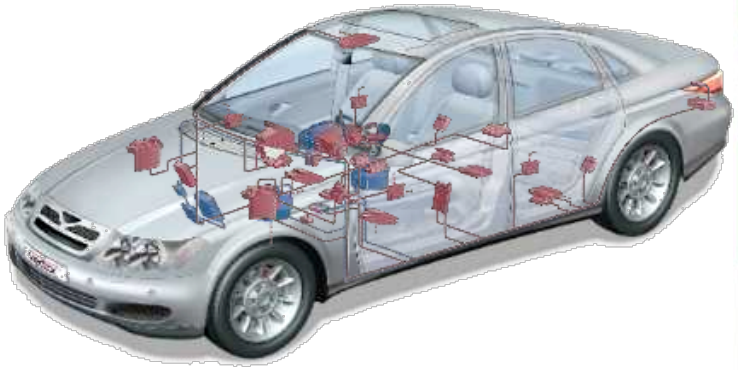

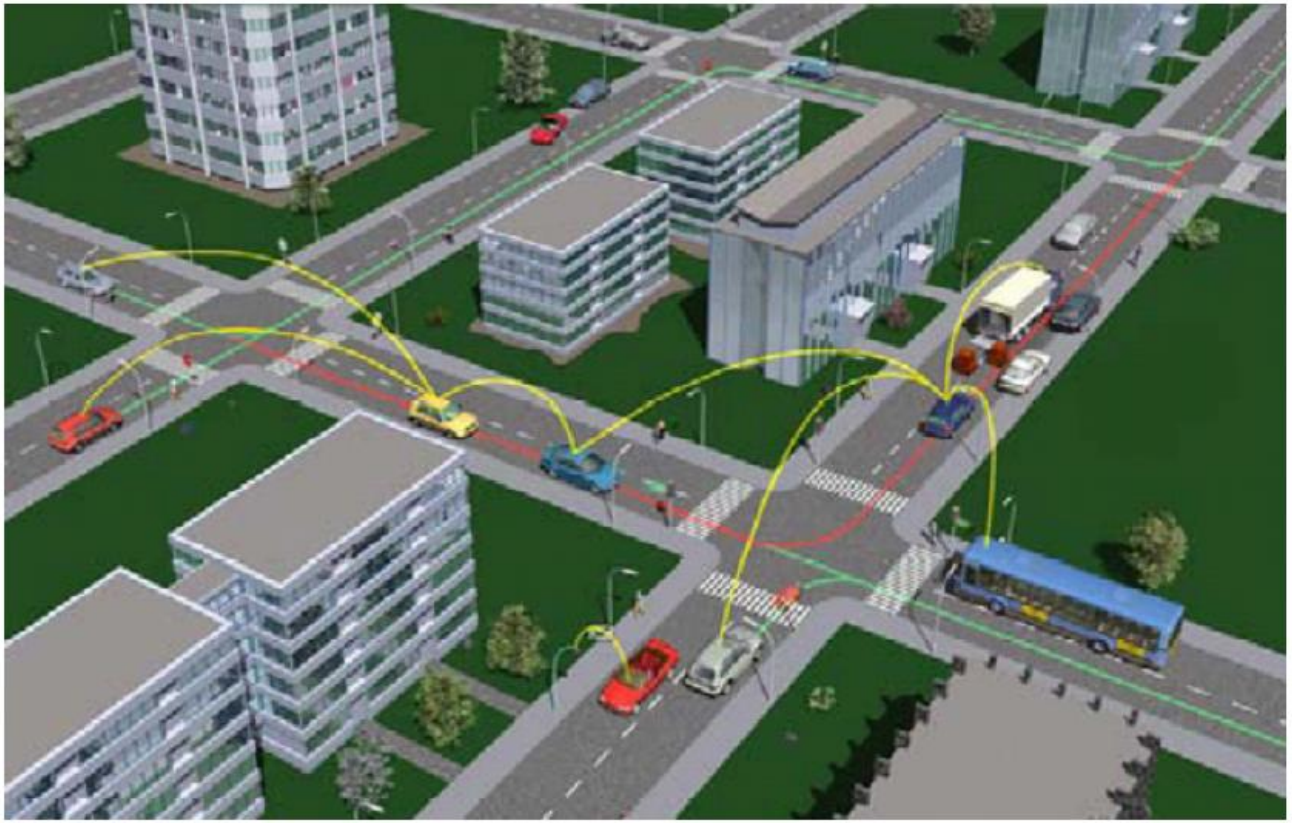

Overall High-Level Objectives and Vision

Provide a more efficient and predictable support (at the OS and Network level) to the development of future embedded systems. In particular:

- Allow simple and flexible resource management to cope with the growing complexity;
- Take advantage of multi-core platforms;
- Support distributed computing to deal with the ubiquitous nature of the computing infrastructure;
- Increase system adaptivity to react to environmental changes.

Partners	Role
Core Partners	
SSSA, Pisa (leader)	RT scheduling and RT kernels
Univ. of Porto	networking, distributed applications
TU Kaiserslautern	video streaming, off-line scheduling
Univ. of Cantabria	schedulability analysis and OS standards
Univ. of York	fixed priority scheduling
Univ. of Lund	real-time and control issues
PI Porto	heterogeneous networks
IMEC	Low power and memory management
Affiliated Partners	
Windriver	RT operating systems
Microchip Technology	Embedded applications
NXP	QoS management in multimedia systems
Evidence	kernels and tools for RT systems
UP Catalonia	control methodologies for RT systems
Univ. of Catania	distributed systems
Univ. of Dresden	Microkernel architectures
Univ. of Madrid	QoS and resource management
Univ. of Pavia	RT applications and kernels
Univ. of Valencia	Real-time and control issues
EPFL Lausanne	Network protocols and distributed systems
Univ. of Balearic Islands	Dependable RT distributed systems

Geographical distribution of the cluster partners

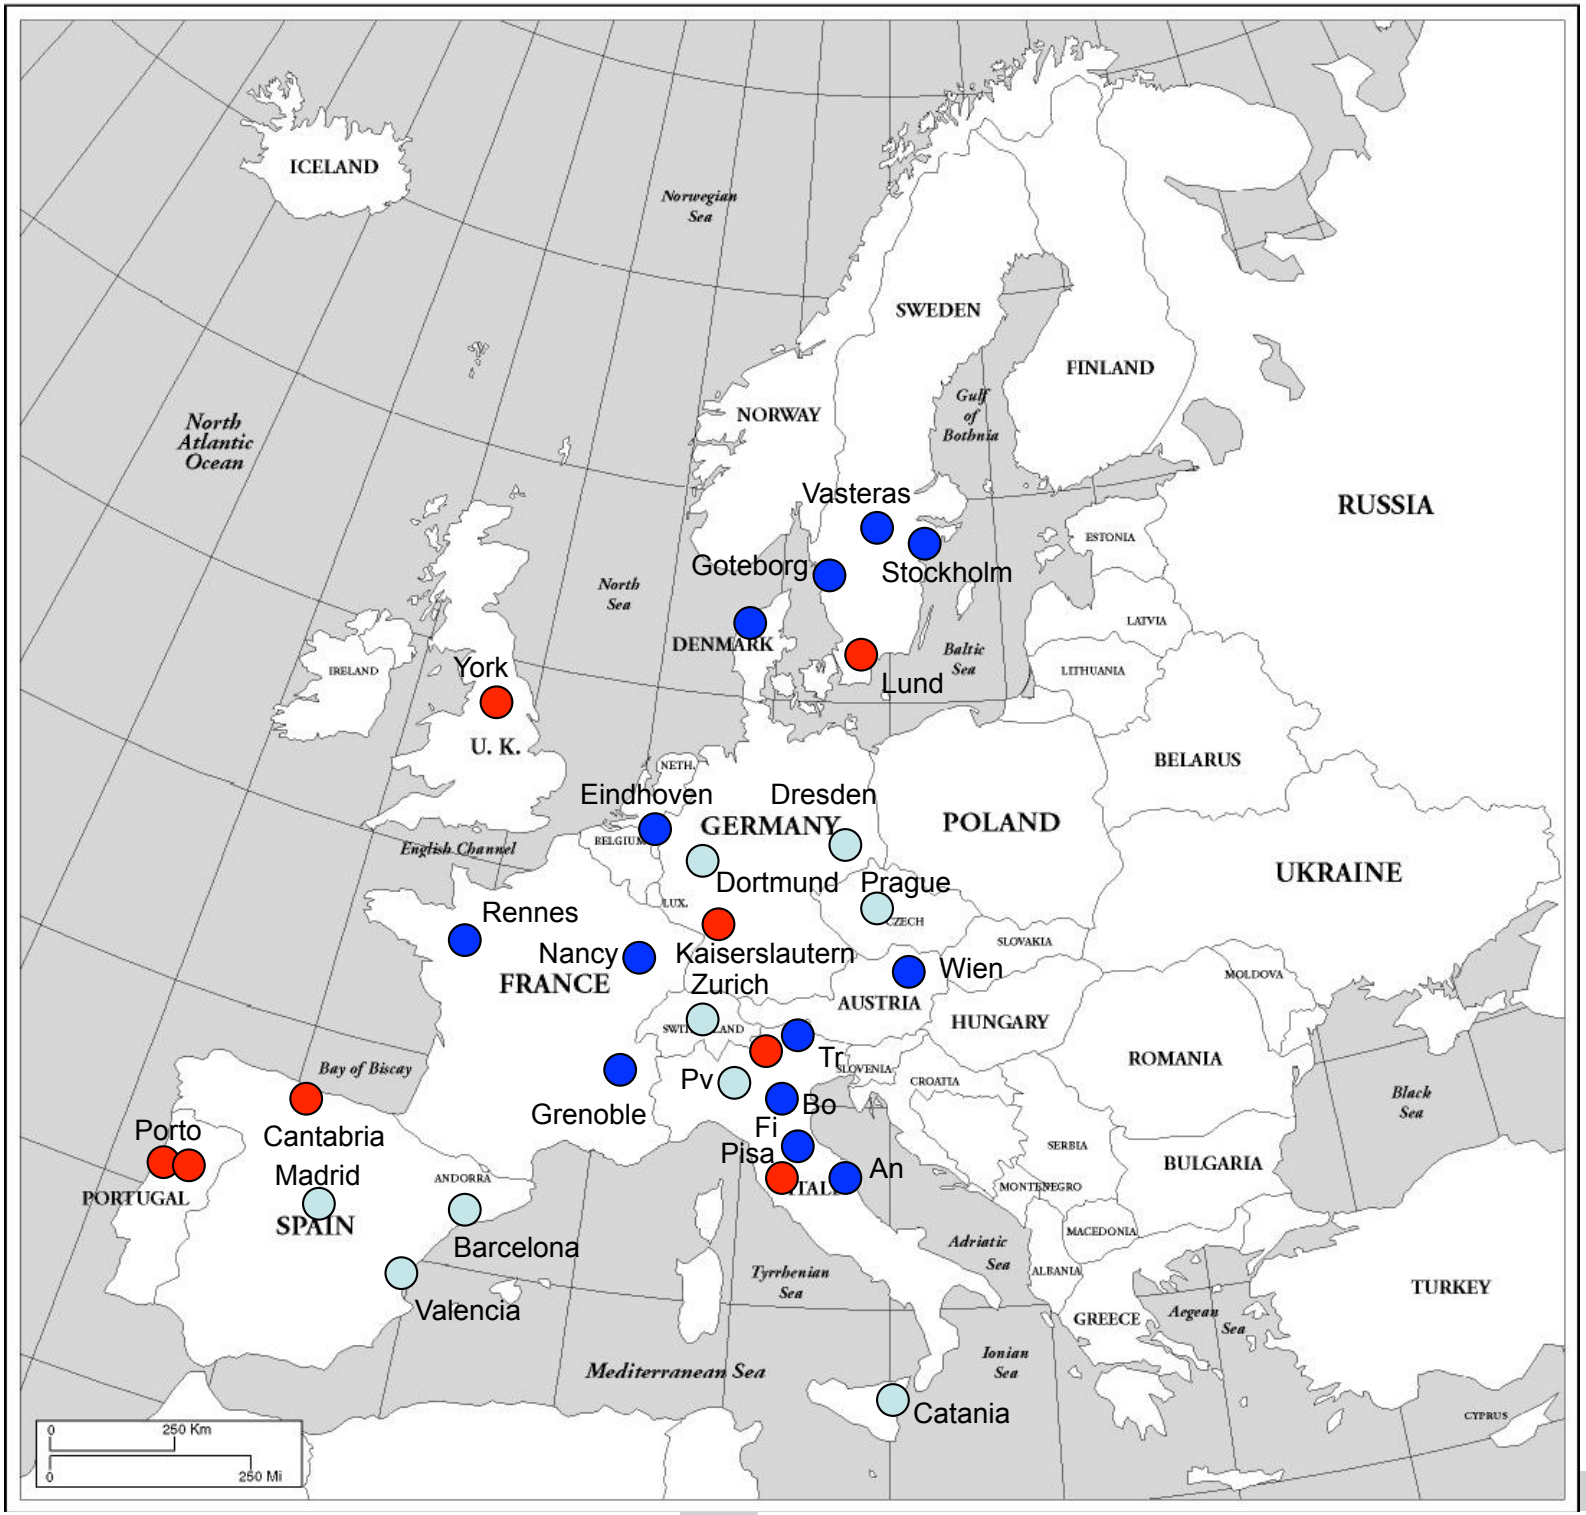

- Core partners
- Affiliates
- Others


Philadelphia
Pittsburgh
UIUC
Berkeley
Virginia
UNC
Austin
FSU

Additional International Partners

- University of Illinois at Urbana-Champaign
  - Contacts: Lui Sha, Tarek Abdelzaher, Marco Caccamo
  - Topics: Sensor networks, RT scheduling and control
- University of Virginia
  - Contacts: John Stankovic, Sang Son
  - Topics: Sensor networks and RT data management
- University of North Carolina at Chapel Hill
  - Contacts: Sanjoy Baruah, James Anderson, Kevin Jeffay
  - Topics: Multi-processor scheduling, multimedia systems
- Carnegie Mellon University
  - Contacts: Ragunathan Rajkumar, John Lehoczky
  - Topics: Resource reservations, QoS management
- University of Pittsburgh
  - Contacts: Daniel Mossé,
  - Topics: Energy-Aware Scheduling

Cluster activities

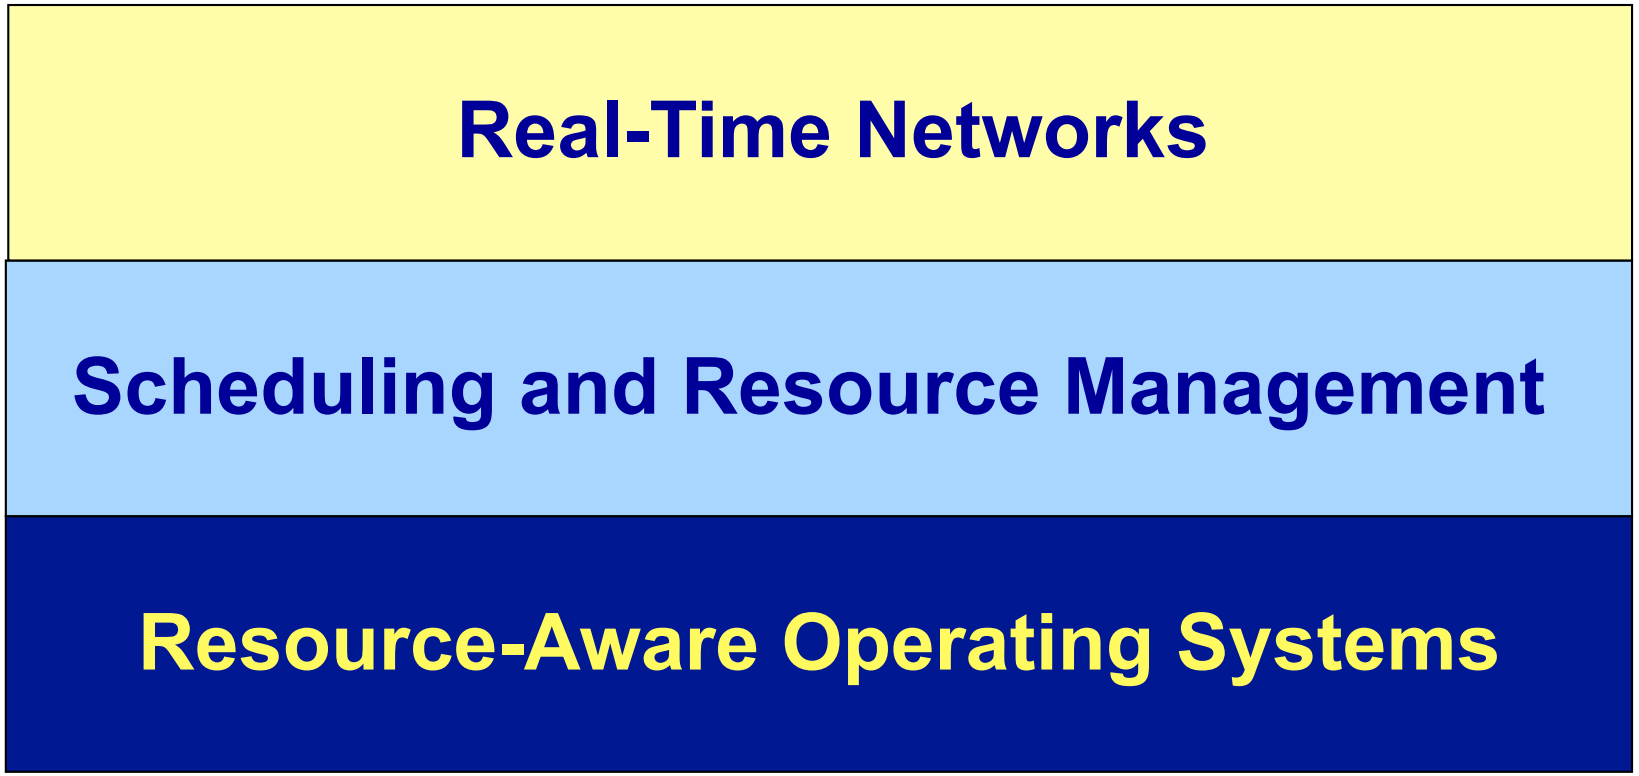

Cluster outcomes – Year 4

- **152** **publications** from individual groups **(**A1: **38**, A2: **87**, A3: **27)**
- **70** **joint publications (**A1: **16**, A2: **16**, A3: **21)**
- **14** **Keynote speeches**
- **26** **Conference/Workshops/Special Tracks/**
- **6** **Educational activities** (summer schools and graduate courses)
- **7** **involvements in standardization activities**
- **7** **tools and platforms integration activities**
- **7** **student exchanges**
- **6** **joint projects** (ACTORS, IRMOS, PREDATOR, FRESCOR, WASP, ILAND)

Cluster outcomes in total

	Y1	Y2	Y3	Y4	Total
Individual publications	60	102	141	152	**455**
Joint publlications	26	47	53	70	**219**
Keynote speeches	5	8	7	14	**41**
Workshops & Tutorials	14	13	14	26	**67**
Graduate Courses	1	1	6	6	**14**

Workshops and Conferences

Events initiated and steered by the cluster

- **OSPERT**: Workshop on Op. Sys. Platforms for Emb. RT applications
- **RTN**: Int. Workshop on RT Networks
- **APRES**: Int. Workshop on Adaptive and Reconfigurable Systems

Events strategically steered by the cluster

- **ECRTS**: Euromicro Conference on Real-Time Systems
- **RTSS**: IEEE Real-Time Systems Symposium
- **RTAS**: IEEE RT and Embedded Tech. and Appl. Symposium
- **ETFA**: IEEE Int. Conf. on Emerging Tech. and Fact. Automation
- **HSCC**: ACM Int. Conf. on Hybrid Systems: Computation and Control
- **CRTS**: Int. Wks. on Compos. Theory and Tech. for RT Emb. Systems

Major conferences have Special Issues on the International Journal of Real-Time Systems (Springer)

Special issues organized by the cluster

1. "Real-Time and Network Systems," S. Baruah and Y. Sorel, Eds., Real-Time Systems, 48(1), 2012.
2. "Real Time Systems Resource Management," J. Anderson Ed., Real-Time Systems, 47(5), 2011.
3. "Embedded and RT Computing Systems and Appls," E. Tovar, Ed., Real-Time Systems, 47(3), 2011.
4. "Energy Aware Real-Time Systems," D. Mossé, J. Leite, D. Kusic, Eds., Real-Time Systems, 47(2), 2011.
5. "Real-Time and Network Systems," M. Chetto and M. Sjödin, Eds., Real-Time Systems, 46(3), 2010.
6. "Euromicro Conference on Real-Time Systems," I. Puaut, Ed., Real-Time Systems, 46(1), 2010.
7. "Real-Time Appl.s and Tools Design," G. Buttazzo and T.-W. Kuo, Eds., IEEE Trans. on Ind. Inf., 6(4), 2010.
8. "Real-Time Networking Systems," P. Minet, Ed., Real-Time Systems, 43(2), 2009.
9. "Euromicro Conference on Real-Time Systems," A.Burns, Ed., Real-Time Systems, 43(1), 2009.
10. "Real-Time Systems - Part II," G. Buttazzo and T.-W. Kuo, Eds., IEEE Trans. on Ind. Inf., 5(1), 2009.
11. "Euromicro Conference on Real-Time Systems," M. González, Ed., Real-Time Systems, 40(3), 2008.
12. "Euromicro Conference on Real-Time Systems," G. Lipari, Ed., Real-Time Systems, 39(1-3), 2008.
13. "Real-Time Systems - Part I," G. Buttazzo and T.-W. Kuo, Eds., IEEE Trans. on Ind. Inf., 4(4), 2008.

Graduate Courses and Seminars

- **Graduate Course on Combinatorial Optimization**
  Scuola Superiore Sant'Anna, Pisa, Italy – October-November 2011.
- **ARTIST Graduate Course on Real-Time Kernels for Microcontrollers**
  Scuola Superiore Sant'Anna, Pisa, Italy – June 13-17, 2011.
- **Graduate Course on Android Framework**
  Scuola Superiore Sant'Anna, Pisa, Italy – November-December 2011.
- **Workshop on Real-Time System Models for Schedulability Analysis**
  University of Cantabria, Santander, Spain - February 7-8, 2011.
- **15th International Real-Time Ada Workshop (IRTAW-15)**
  Liébana (Cantabria), Spain, September 2011.
- **Course on Real-Time Communication for Embedded Systems**
  ENSIAS, Rabat, Morocco, 19-21 December 2011.

Participation in Standards

- **UML Profile QoS and Fault Tolerance**
  Member: Miguel A. de Miguel, UP Madrid
- **ADA**
  Member: Alan Burns, Univ. of York
- **POSIX 1003**
  Member: Michael Gonzalez Harbour, Univ. of Cantabria
- **MPEG Multimedia Middleware (M3W)**
  Member: Alejandro Alonso, UP Madrid
- **TinyOS 15.4 and ZigBee Working Groups**
  CISTER-ISEP at Polytechnic Institute of Porto
- **Working Group 16-17: Wireless Industrial communication networks**
  Member: Lucia Lo Bello, Univ. of Catania
- **Linux development team**
  Scuola Superiore Sant'Anna, Pisa

Cluster outcomes - overall

Industrial collaborations

Avionics:	*Airbus, Embraer*
Automotive:	*Bosch, Magneti Marelli, TTTech*
Railway systems:	*Ansaldo*
Robotics:	*ENSA (Equipos Nucleares)*
Video surveillance:	*Visual Tools*
Healthcare:	*Trialog, CAEN*
Consumer electronics:	*Ericsson, Philips, NXP*
RTOS producers:	*Wind River, Evidence, CISCO*
Platforms:	*ARM, Microchip Technology*
Embedded systems:	*Intecs, Embedded Solutions*

Industrial impact

- EDF with Resource Reservation in **Linux 2.6.x**
- Resource reservation is used by **Ericsson** (ACTORS project)
- Cache aware-Limited preemption scheduling for automotive/avionic applications considered by **Airbus** and **Bosch** (PREDATOR project)
- ERIKA kernel used by
  - **Merloni** for washing machines
  - **Magneti Marelli** for automotive applications
  - **Aprilia** for motorbike engine control
- FTT-SE Ethernet protocol is used by **Trialog** (iLAND project)
- MARTE OS is used by **Equipos Nucleares**

Tools and platforms

- **Erika+MPARM**: Real-Time Emulator for Multicore ARM7 platforms
- **PartiCore** – A Partitioning Tool for Multi-core Reservations
- **YaoSim**: Yet Another Operating system SIMulator
- **Marte OS**: an efficient operating system for embedded applications
- **HaRTES**: a predictable and reconfigurable Ethernet switch for hard real-time communications
- **iLAND**: a middleware architecture for supporting reconfiguration in distributed soft real-time systems
- **Simulation environment** for Multimedia Sensor Networks
- **Framework for adding Flexibility** to the Time-Triggered paradigm

Year 4 Review
Dresden, March 16th, 2012

Achievements and Perspectives:

Resource-Aware Operating Systems

Activity leader: *Giorgio Buttazzo*
Scuola Superiore Sant'Anna
Pisa, Italy

Objectives

Exploit the expertise in the NoE to make operating systems more

- predictable (in terms of timing behavior)
- efficient (in terms of resource usage)
- robust (to tolerate overload conditions)
- easy to use (to simplify user interface)

Approach

- Investigate novel kernel **methodologies**
  - Resource reservation
  - Contract-based scheduling
  - Limited preemptive scheduling
  - Energy-aware policies
- Implement these techniques in **existing RTOSes**
- Provide **appropriate tools**

Emphasis on multi-core platforms

Existing expertise

Industrial partners

- Airbus FR
- Bosch FR
- Thales FR
- Ericsson SE
- ENEA SE
- Absint DE
- Visual Tools ES
- Evidence IT
- Rapita Systems UK

Academic partners

- Sant'Anna, Pisa IT
- Univ. of Bologna IT
- Univ. of York UK
- Univ. of Cantabria ES
- Univ. of Valencia ES
- Univ. of Prague CZ
- Univ. of Lund SE
- U. Kaiserslautern DE
- Univ. of Saarland DE
- Univ. of Dortmund DE
- EPFL Lausanne CH
- ETH Zurich CH

Target RTOSes

ERIKA ENTERPRISE

EVIDENCE

http://erika.tuxfamily.org/

- Small platforms (1-2 Kbytes)
- OSEK compliance

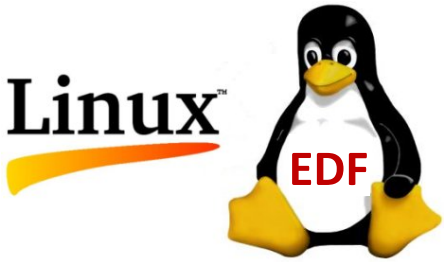

+ SCHED_DEADLINE

http://www.evidence.eu.com/sched_deadline.html

- PC-like platforms
- POSIX compliance

Both support

- Multi-core platforms
- Resource reservation
- Deadline-based scheduling

Integrating EDF in Linux

Deadline Scheduling on Linux (*Pisa, Evidence*)

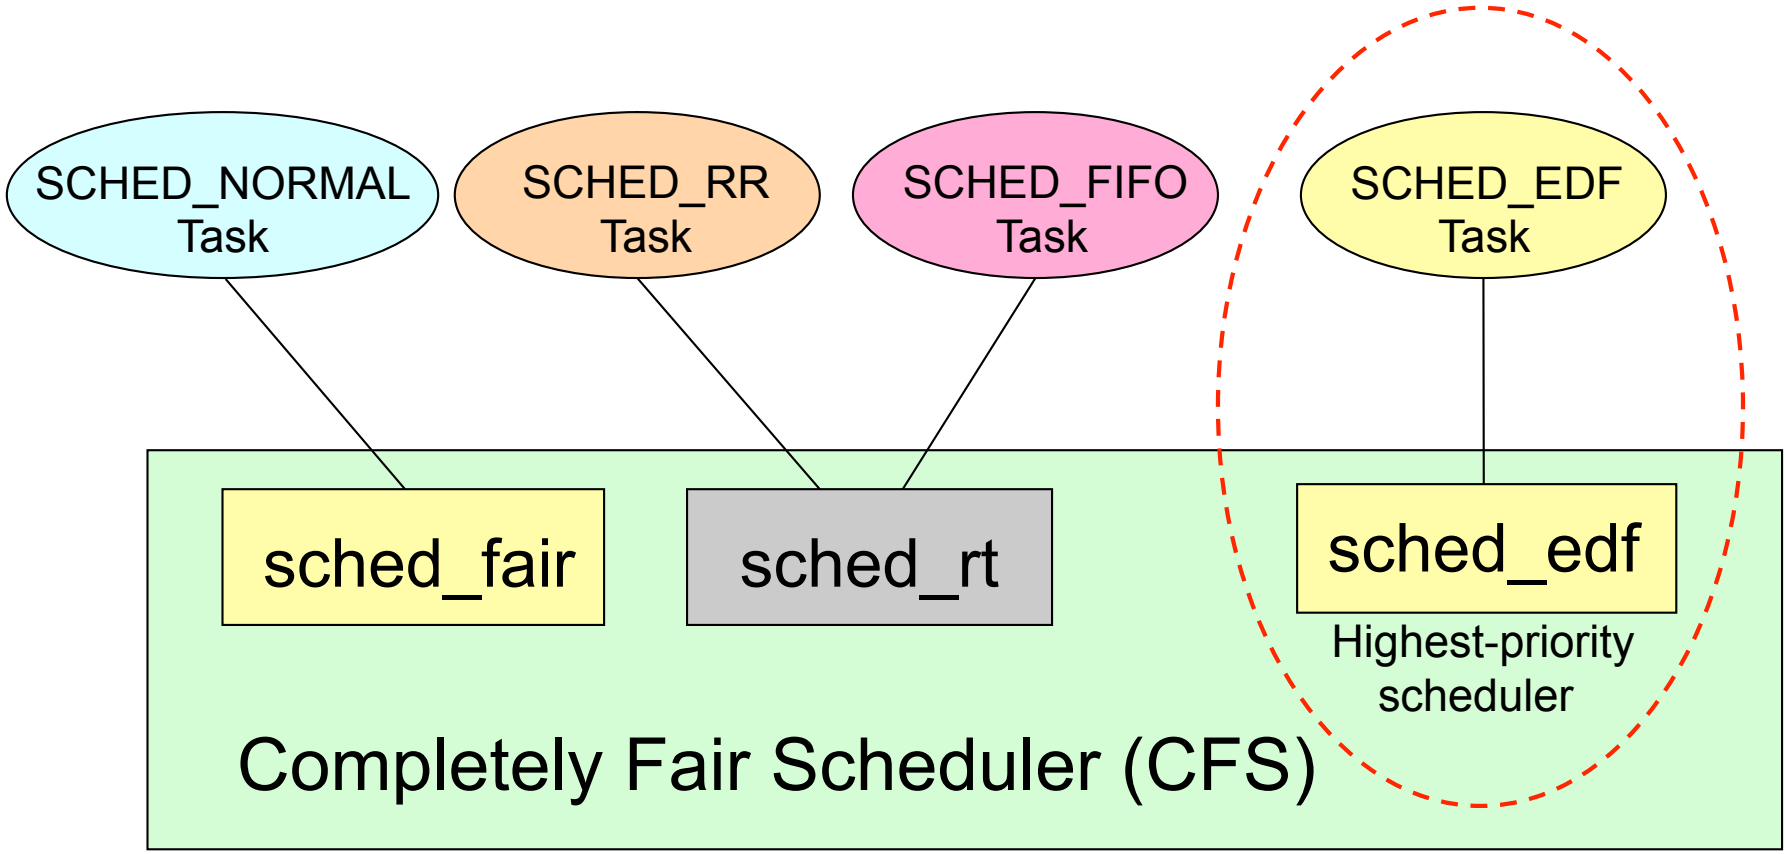

Interrupt handling in MaRTE OS [Cantabria]

- Some systems have high utilization from interrupts.
- Common implementations charge interrupt times arbitrarily, causing large scheduling errors.
- A new implementation in **MaRTE OS** introduces:
  - a specific clock for interrupt time
  - timers triggered with high accuracy
  - acceptable overhead
    (192 ns for a minimal 3μs interrupt)

Example system: GPS-based attitude sensor

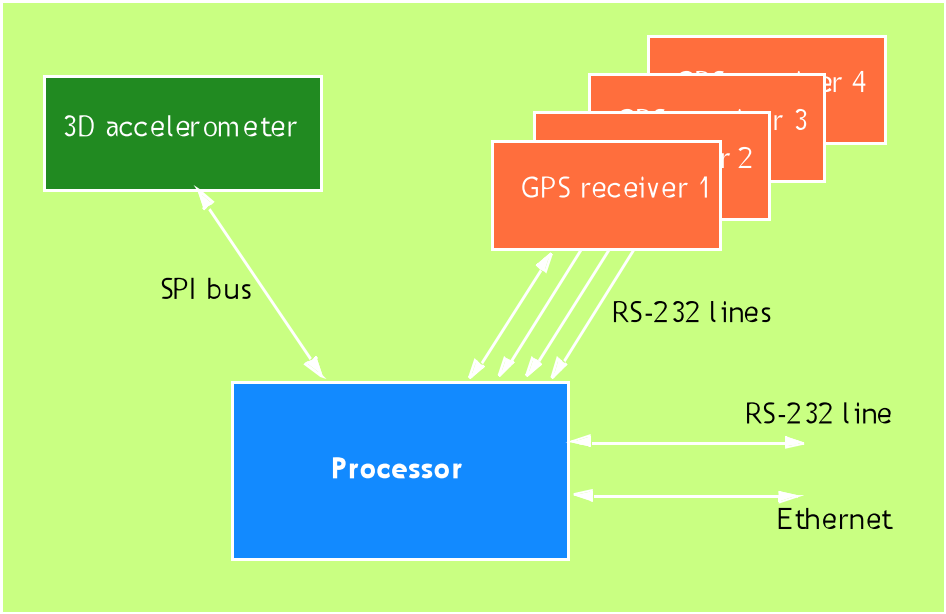

Result: correct resource reservations in systems with high interrupt utilizations

Tools

RTSIM	Real-time scheduling simulator
YAOSIM	Simulator for limited preemptive scheduling
RT-DRUID	Design and Analysis tool for Erika Enterprise
PARTICORE	Partitioning Tool for Multi-core Reservations

Main Technical Achievements

Partitioning applications on multi-core platforms

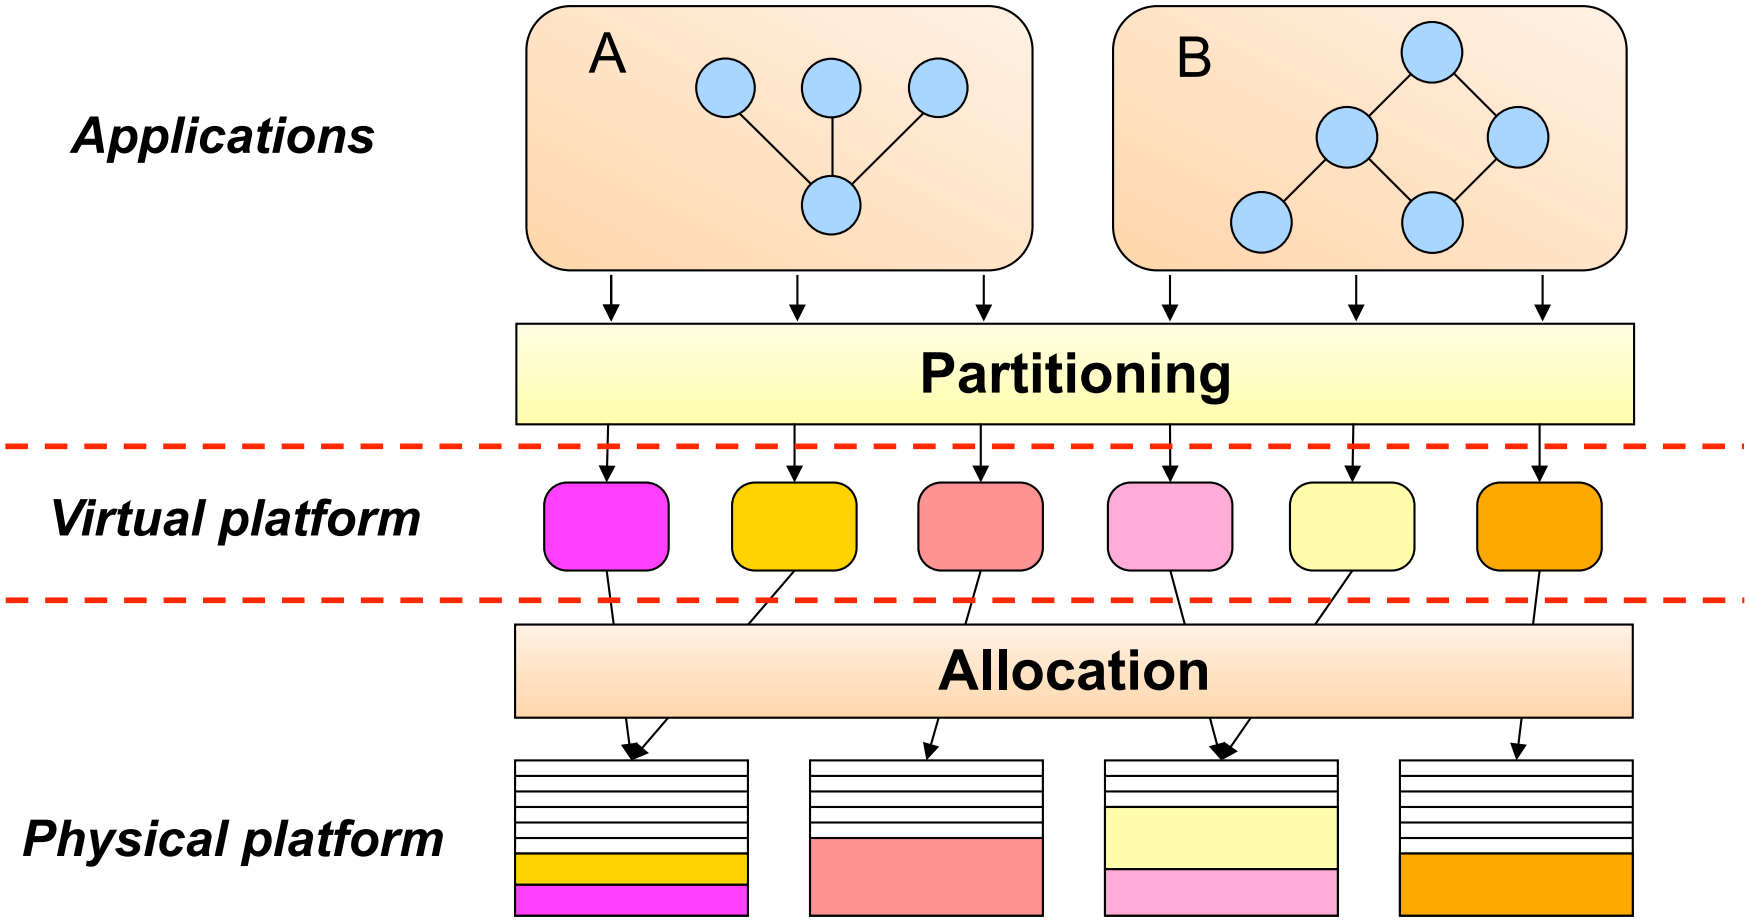

The PartiCORE Tool

URL: http://particore.sssup.it/

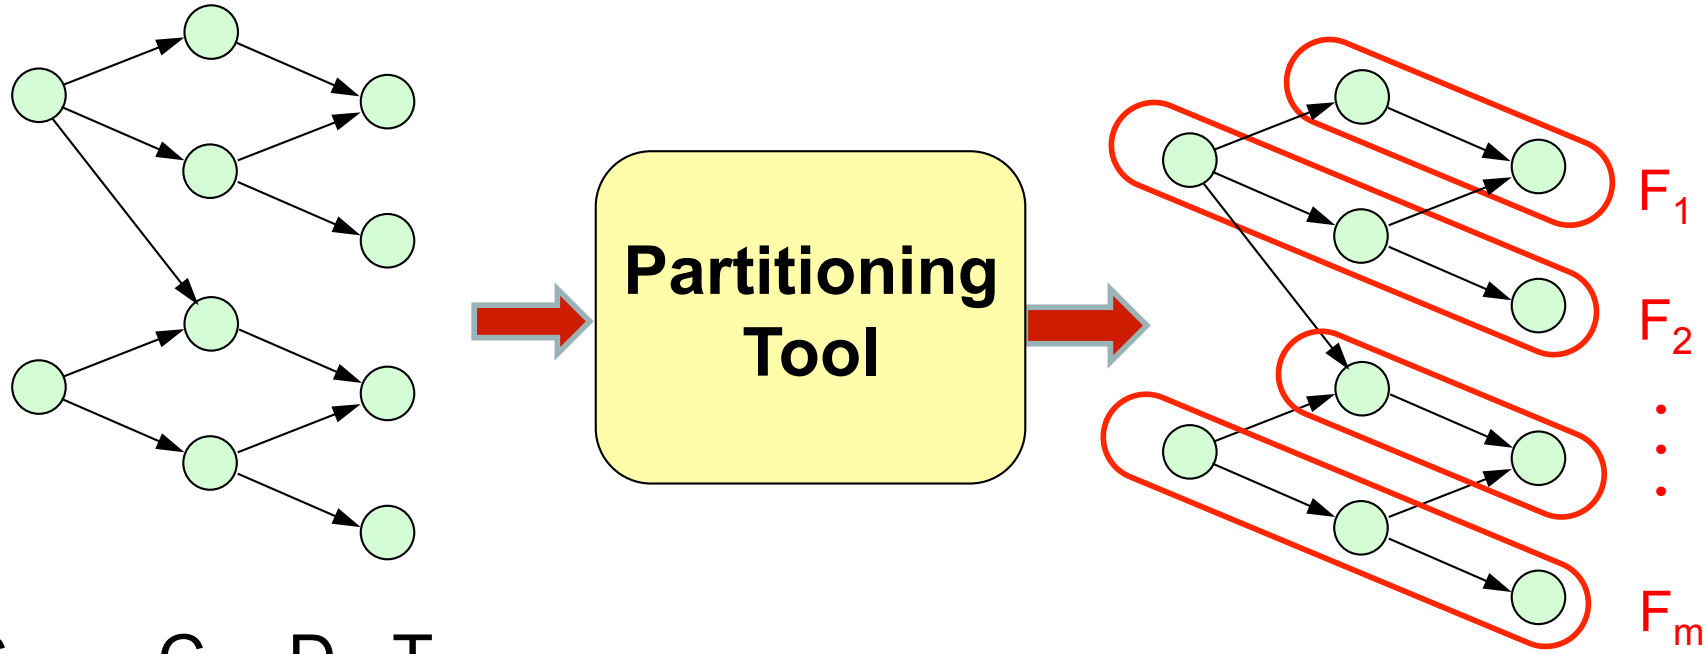

$C_1, \ldots C_n, D, T$

ctx-sw. overhead σ

$B_1, \ldots, B_m$

Such that $B = \sum B_k$ is minimum

Partitioning Tool for Multi-core Reservations

If the graph does not show. Please download svg support plugin.
*Browsers natively support svg: Firefox 1.5+, Opera 8.5+, Safari 3.0+, Chrome 1.0+

Precedence Graph

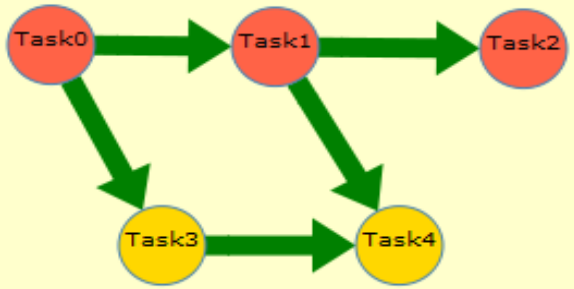

Timeline Representation/Parellel Number Function

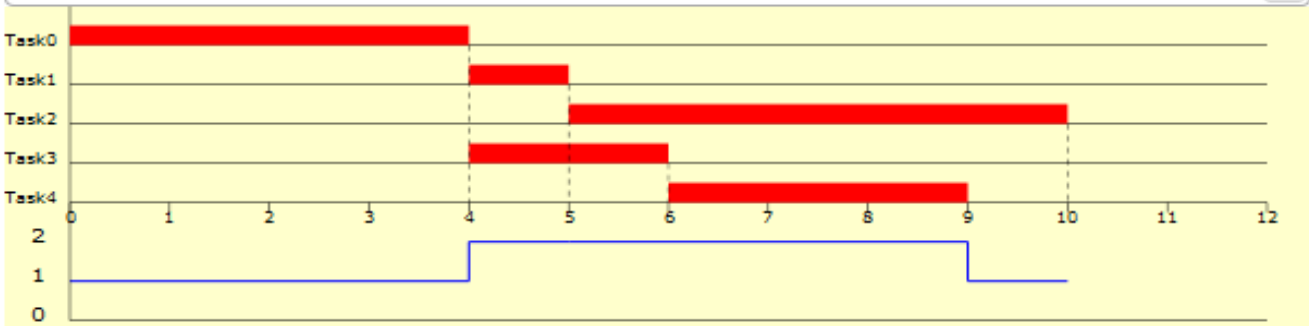

Demand Bound Function/Alpha-Delta Server

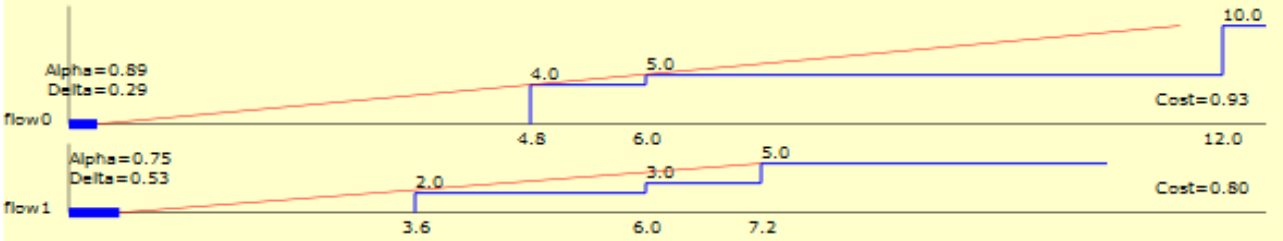

Application | Tasks | Partitions | Log

Application parameters

Arrival time: 0 Period: 12

Deadline: 12 sigma: 0.1

sequentialC: 15 parallelC: 10

max Parallel Number: 2

Set Params

Configuration

Deadline Assignment Method

(•) Chetto* ( ) Chetto

Application

Input Format 1 2

Create Save

Partitioning Tool for Multi-core Reservations

If the graph does not show. Please download svg support plugin.
*Browsers natively support svg: Firefox 1.5+, Opera 8.5+, Safari 3.0+, Chrome 1.0+

Precedence Graph Application 1

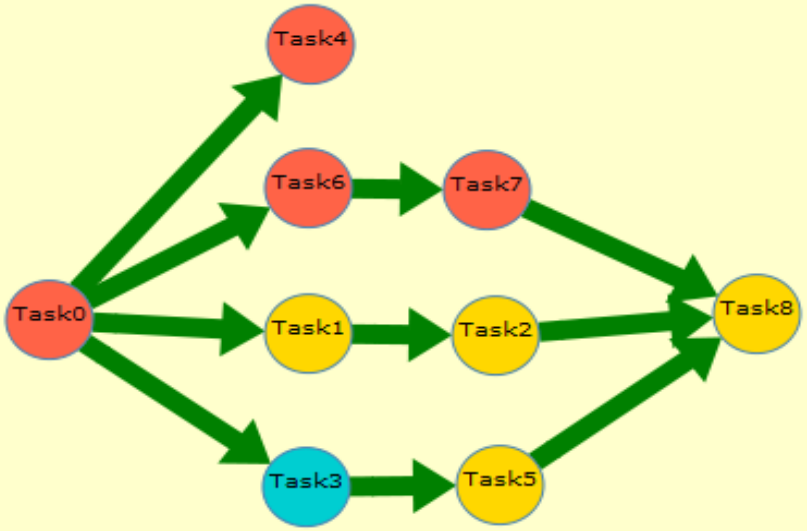

Timeline Representation/Parellel Number Function

Demand Bound Function/Alpha-Delta Server

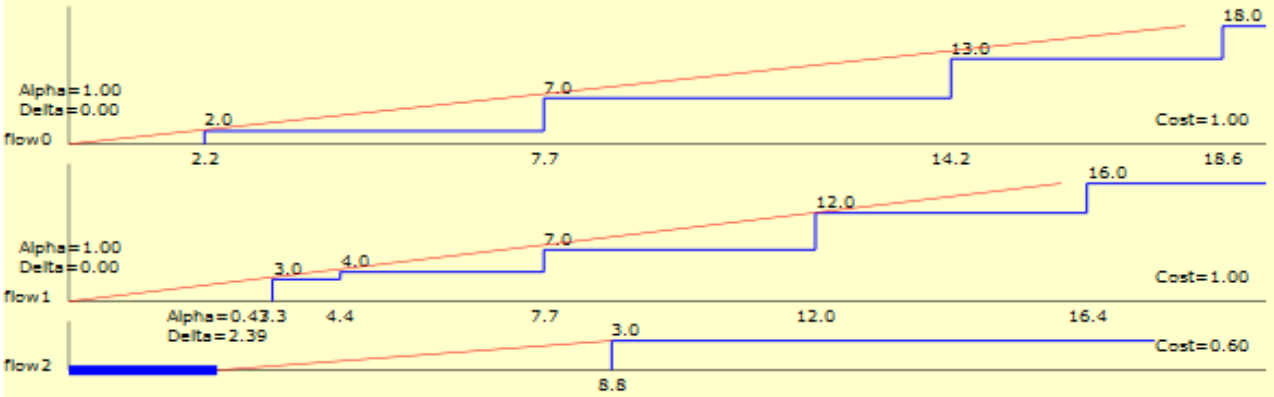

Application | Tasks | Partitions | Log

Optimization Objective

- Total Bandwidth 2.60416
- Fragmentation 2.60416

Automatic Partitioning

Search | Heuristic

Result: partition0

```
[[0,6,7,4],[1,2,5,8],[3]]
[[6,7,4],[0,1,2,5,8],[3]]
```

Manual Partitioning

--Select--

Create | Remove

Partitioning Tool for Multi-core Reservations

If the graph does not show. Please download svg support plugin.
*Browsers natively support svg: Firefox 1.5+, Opera 8.5+, Safari 3.0+, Chrome 1.0+

Precedence Graph

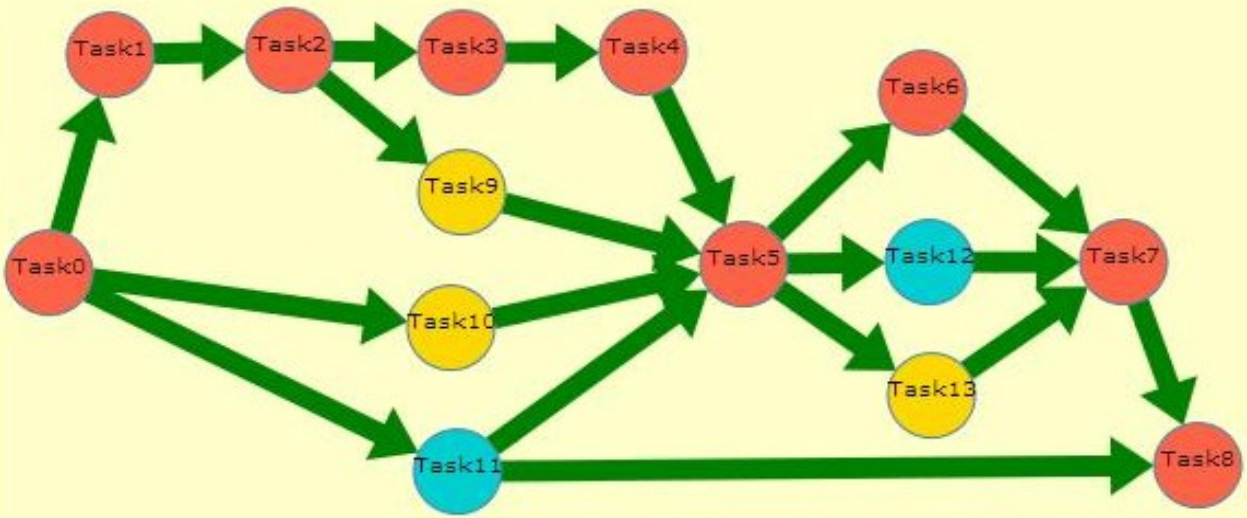

Timeline Representation/Parellel Number Function

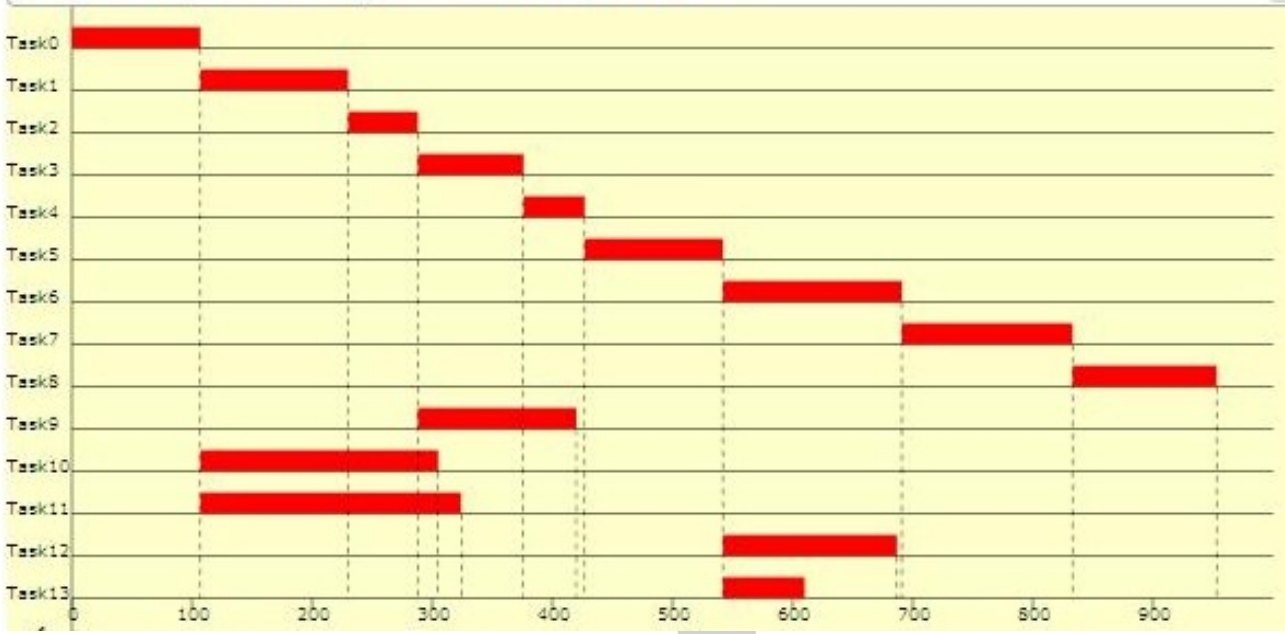

Application | Tasks | Partitions | Log

Optimization Objective

- Total Bandwidth: 2.8632
- Fragmentation: 2.91336

Automatic Partitioning

Search | Heuristic

Result: partition0

[[0,1,2,3,4,5,6,7,8],[10,9,13],[11,12]]

Manual Partitioning

--Select--

Create | Remove

Summer School

Graduate Course on Embedded Control Systems: Theory and Practice

Scuola Superiore Sant'Anna, Pisa, Italy – June 13-17, 2011

- Giorgio Buttazzo - Scuola Superiore Sant'Anna, Italy
- Pau Martì – University of Catalonia, Spain
- Manel Velasco – University of Catalonia, Spain

Educational Platform for Embedded Control Systems

Pisa, Pavia, Evidence, Microchip Technology, Embedded Solutions

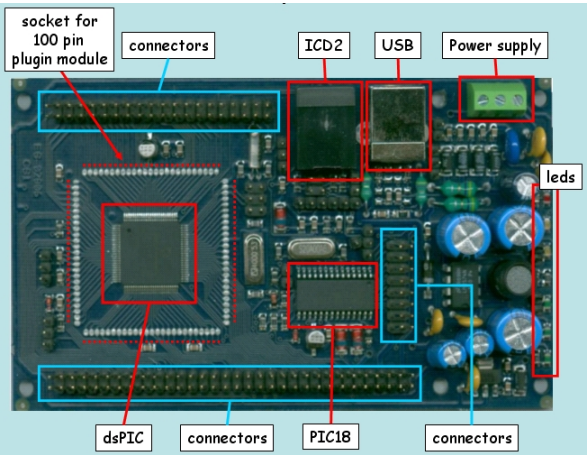

- Fast prototyping RT control applications
- Facing RT scheduling issues
- Handling resource management
- Used in graduate courses

Sample applications implemented by students

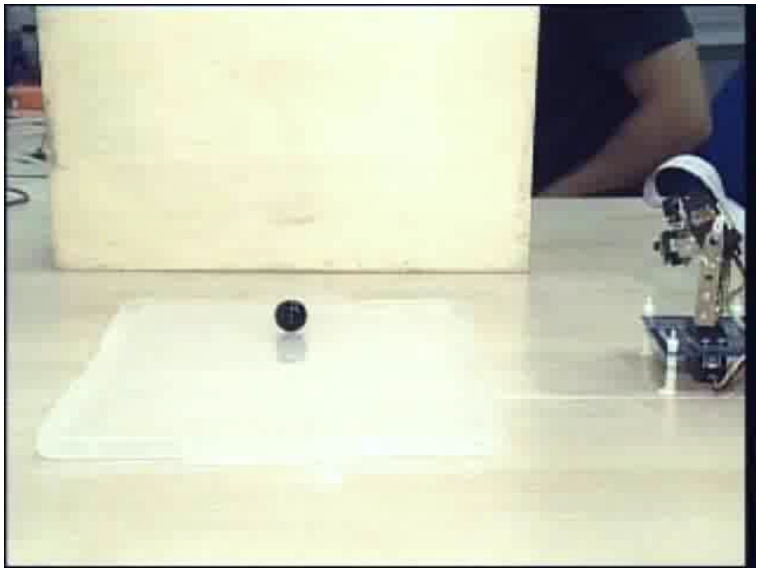

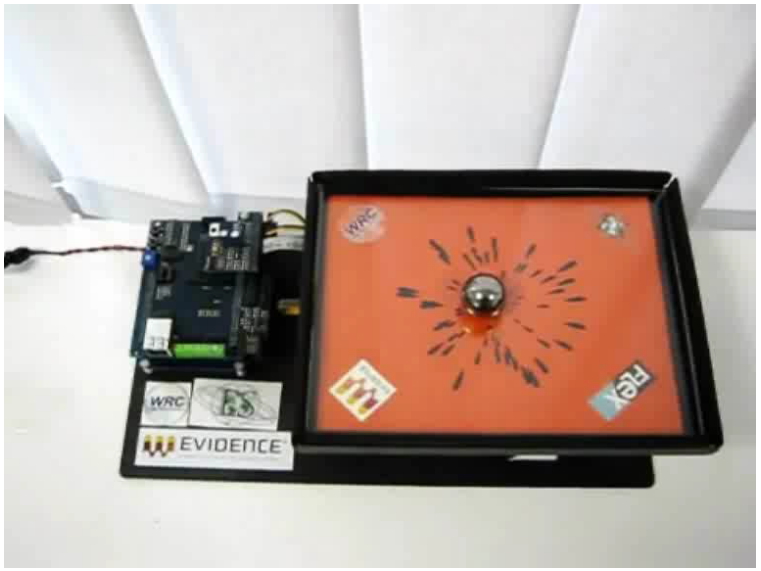

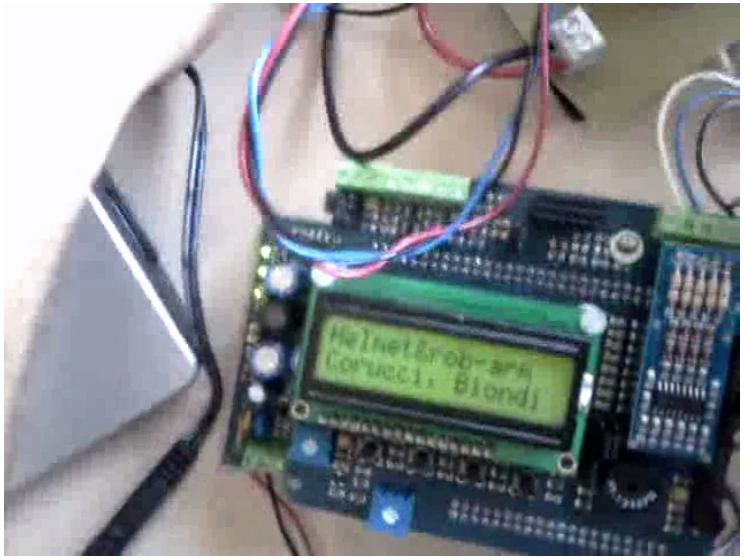

Resource-Aware Operating Systems

43 publications from individual groups

28 joint publications

9 Keynote speeches

14 Workshops/Tutorials

5 Educational activities

3 Tool-integration activities

5 student exchanges

Year 4 Review
March 16th, 2012

Scheduling and Resource Management

Activity leader: *Alan Burns*
University of York
York, UK

Objectives

- **Provide Policies**
  - For effective resource usage
- **Provide Analysis**
  - For predicting system behaviour
  - Simulation, scheduling analysis, measurement, model checking
- **Provide Models**
  - For composing systems
  - Time triggered and event-triggered work flow
  - For static and dynamic usage patterns

Challenges

- To move from single processor platforms to multiprocessor, multi-core, FPGA, etc.
- To integrate various resources and abstract views of the overall system
  - Integrate policies
  - Integrate analysis
  - Integrate models
- Static and Dynamic, peer-to-peer and hierarchical

Problems Tackled in 2011/12

- Continued work on scheduling and placement algorithms for multiprocessor systems
- Developed new theoretical results in terms of energy and power
- Continued work on contract based scheduling
- Considered resource sharing in distributed systems
- Adaptive (dynamic) resource sharing
- Resource-aware communication middleware
- Resource management software in programming languages
- Mixed criticality systems

Overall Assessment and Vision

- Remains a very active area in EU
- 30 technical achievements in deliverables
- 82 papers – many joint authored or result of collaborations
- 400+ papers over the 4 years of ArtistDesign

A highlight – mixed criticality

- Becoming a key aspect of many systems (eg AUTOSAR initiative)
- System partitioned on criticality are much less resource efficient than those in which components with different criticality levels are interleaved
  - For example, single processor, 5 crit levels, speed-up factor could be as high at 300%
- Work on this topic at a number of project sites
  - For example, Braunschweig, Kaiserslautern, Porto, Uppsala and York
  - And with Professor Sanjoy Baruah (UNC, US)

Applications of mixed criticality

high demands
strict timing behavior
safety critical

hard real-time applications
temporal correctness

flexible real-time applications
not completely known

non real-time activities
don't disturb the RT part

Lecture planning

- core topics I want to cover — A B C D
- relevant topics, not necessary — 1 2 3 4
- 90 mins lecture, no interaction — A 1 B 2 3 C D 4

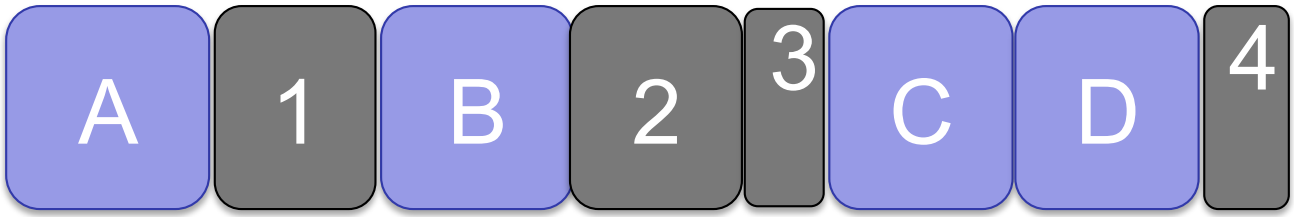

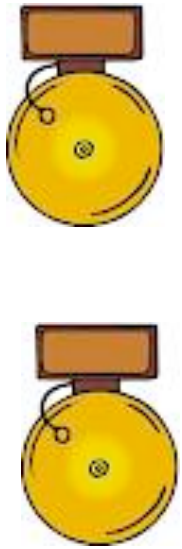

what if topic B becomes interactive?

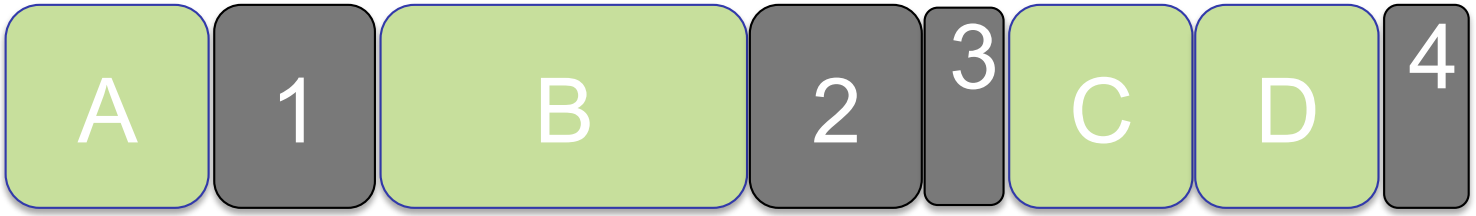

Lecture planning

- make two versions of core topic
  - no questions
    - interactive

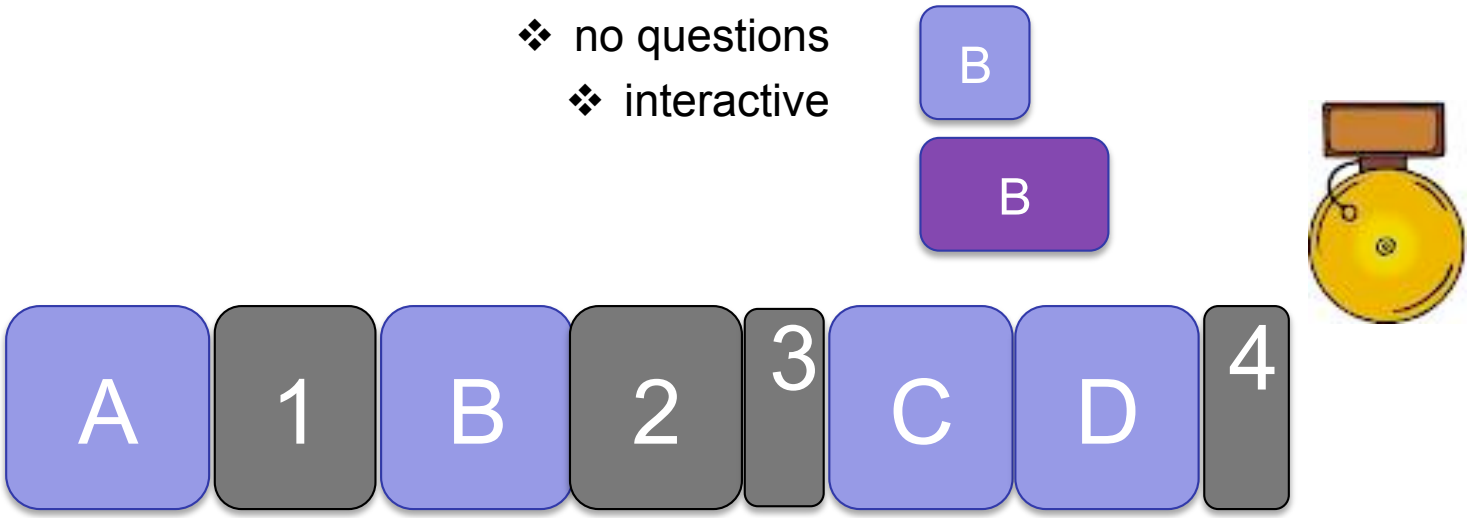

Interactive lecture

Lecture planning

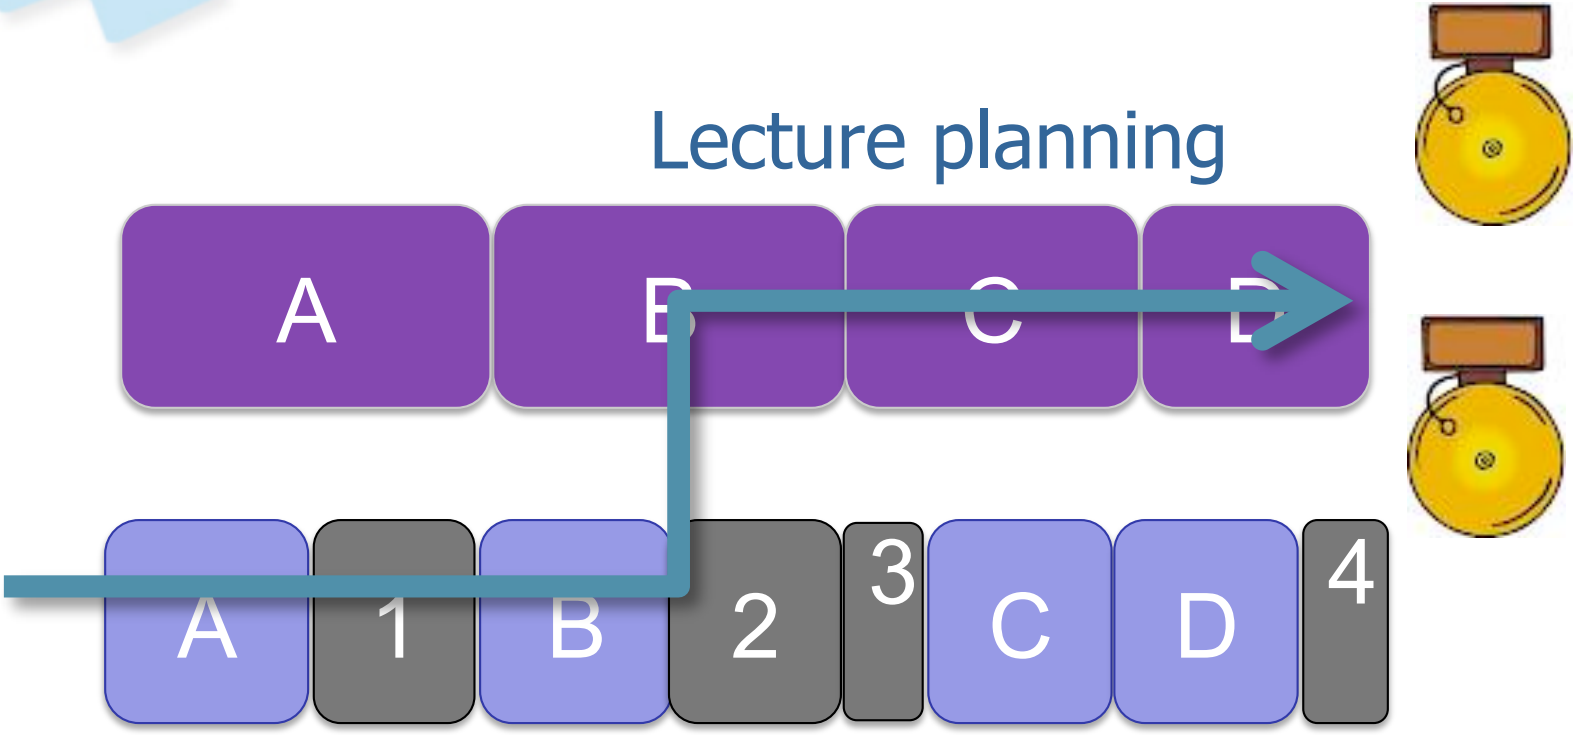

what if topic B becomes interactive?

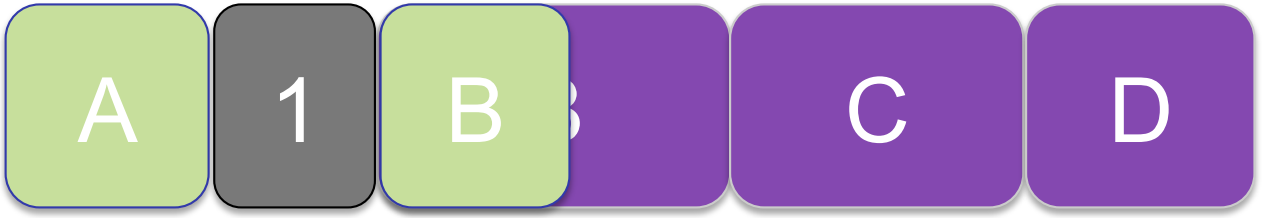

Current Work

- Example papers
  - at RTSS one of mixed criticality and table driven scheduling

    Involving the mode switch just illustrated
  - Another paper at RTSS on mixed criticality and fixed priority scheduling

    Adapted mode change analysis to construct an improved scheduling test for the general sporadic task model
- Others working on different aspects of the issues raised by mixed critically systems

Conclusions

- Over the full 4 years of ArtistDesign, research in the topic of scheduling and resource management has been
  - Active
  - Diverse
  - Effective
  - Useful
  - and rewarding

Year 4, Final Review
Dresden, March 16th, 2012

Real-Time Networks

Activity leader: *Luis Almeida*
University of Porto
Porto, Portugal

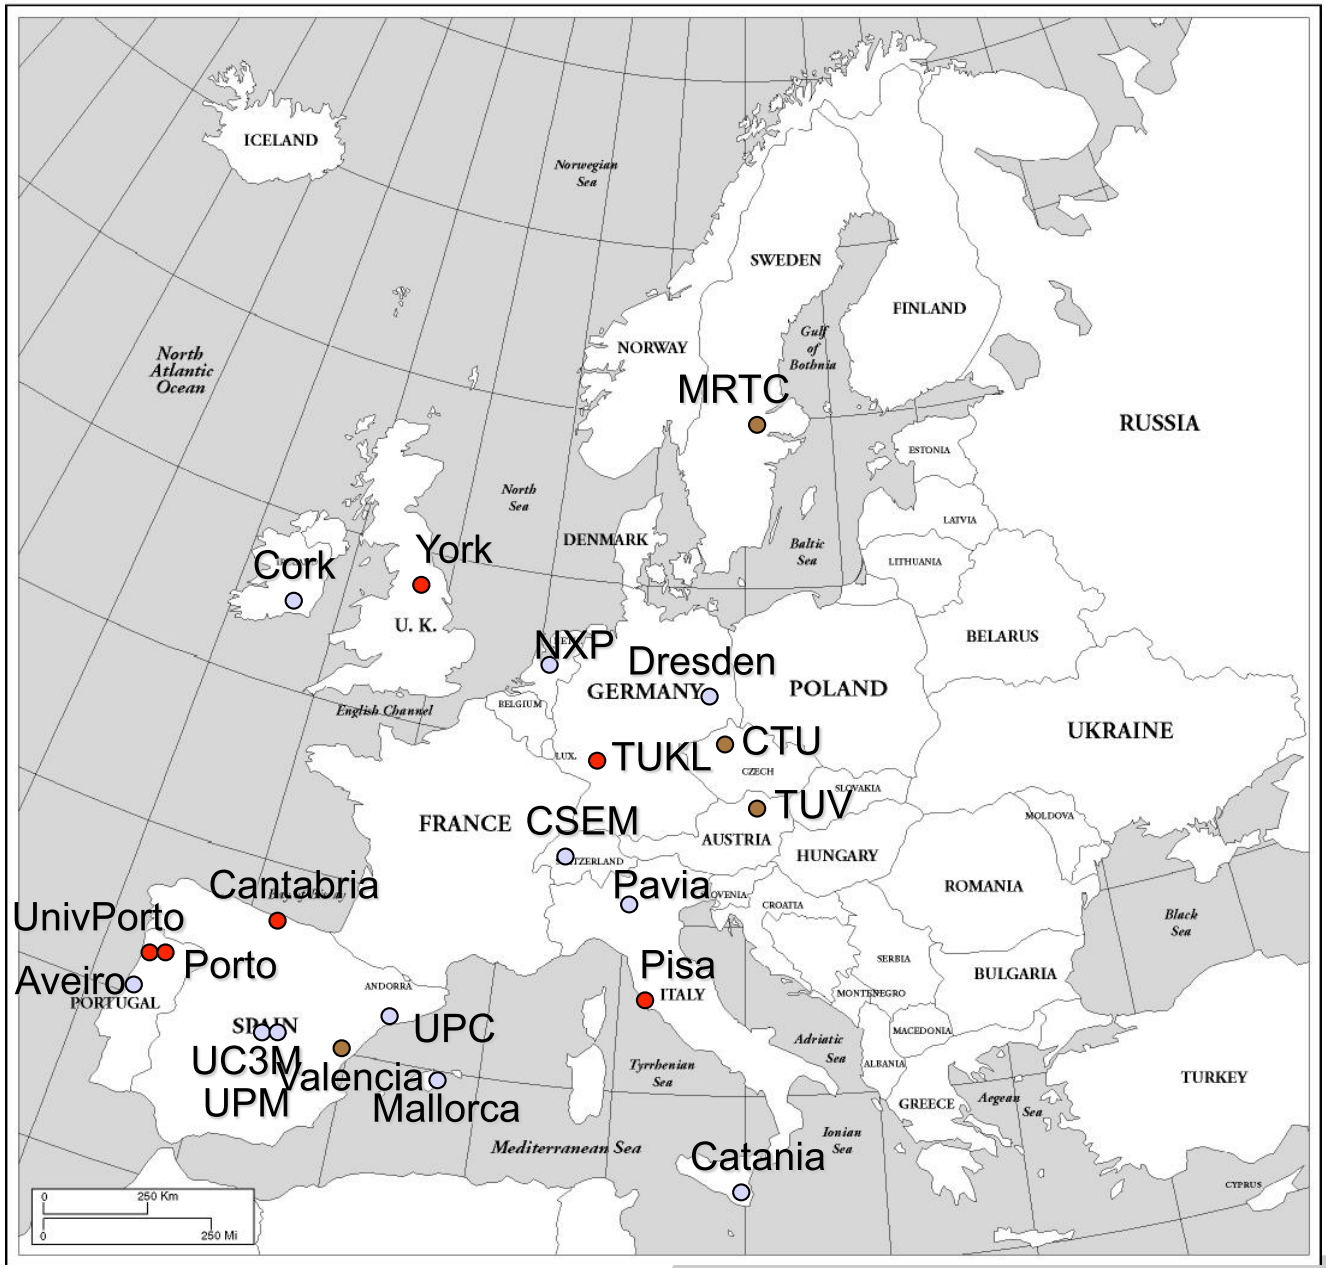

Real-Time Networks Activity

- 6 Core partners
- 9 Affiliated partners
- 4 Other Core partners

Objectives

- **Managing complexity in networked embedded systems**
  - QoS adaptation and graceful degradation
  - higher integration with protection

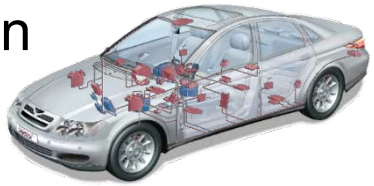

- **Towards *(real-time) wireless everywhere***
  - WSN, MANETs, cooperating embedded systems
  - Reduce communication–related energy consumption

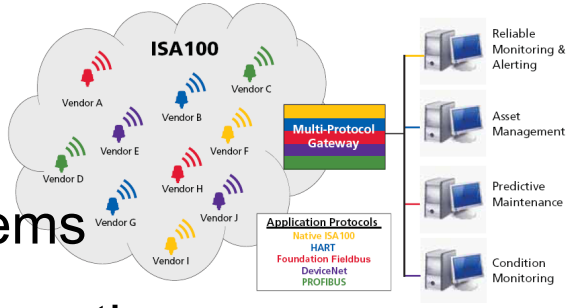

- **Networking technology outreach**
  - courses, seminars, schools, standards, joint R&D projects

Challenges

- **(Real-time) wireless everywhere (WSN)**
  - **Synchronizati**
    - – H
  - **Manag**
    - – Long lifetime, high scalability and data aggregation
- **Managing complexity (NES)**
  - **QoS ad and graceful degradation**
    - – R ations, topology changes, a other reconfigurations
  - **Networking suppo**
    - – Provide real-time and com
      - » Efficient temporal partitioning and dynamic, end-to-end resource reservation

Timeliness and energy in Wireless (Sensor) Networks

Flexibility, robustness and efficiency in NES

Highlights from Year 4

❖Timeliness and energy in W(S)N

- **Engineering WSN** (**integrated toolset** to assist deployment planning, worst-case analysis/dimensioning, simulation and nodes testing/programming using **Open-ZB** and the **Z monitor**)
- **Distributed visual-based tracking and localization in WSN** (for ambient intelligence, intelligent transportation, teams of robots…)

❖Flexibility, robustness and efficiency in NES

- Improved timing analysis and tools for **CAN** (impact of **device drivers**, topology optimization, **fault injectors**, **stars**…)
- New analysis for **Ethernet/AFDX** switches (**MAST suite**), adaptive virtual channels (**HaRTES, FTT-SE**), use in the automotive domain
- **DDS** analysis, real-time **service composition** and **reconfiguration**
- Middleware and design approach for **high-integrity systems**

Summary of outcomes from Year 4

- **46 publications** from individual groups (11 in journals)
- **30 joint publications** (4 journals, 2 book chapters, 16 groups beyond CP+AP)
- **3 workshops** (RTN, SOCNE, NeRES)
- **4 special sessions/tracks** (ETFA, ISIE, INDIN, IECON)
- **5 tutorials/seminars** (3 in schools, 2 in courses)
- **1 special section** (IEEE Trans. on Ind. Informatics)
- **5 joint international projects** FP6/7-STREP (**FlexWARE, MADES**), ARTEMIS (**iLAND, EMMON, INDEXYS**), all with strong industrial participation plus **3 joint national** (**HARTES-P, CANbids-E, IPERMOB-I**)
- **Industrial collaborations** with

  Critical Software (P), Visual Tools (E), Trialog (F), Magneti-Marelli (I), Evidence (I), STMicroelectronics (I), NXP (NL), EADS (D), Softeam (F), X/Open (UK), TXT E-solutions (I),

Summary of global outcomes

Publications produced

- **121** publications from **individual groups**
- **89** publications from **joint work**

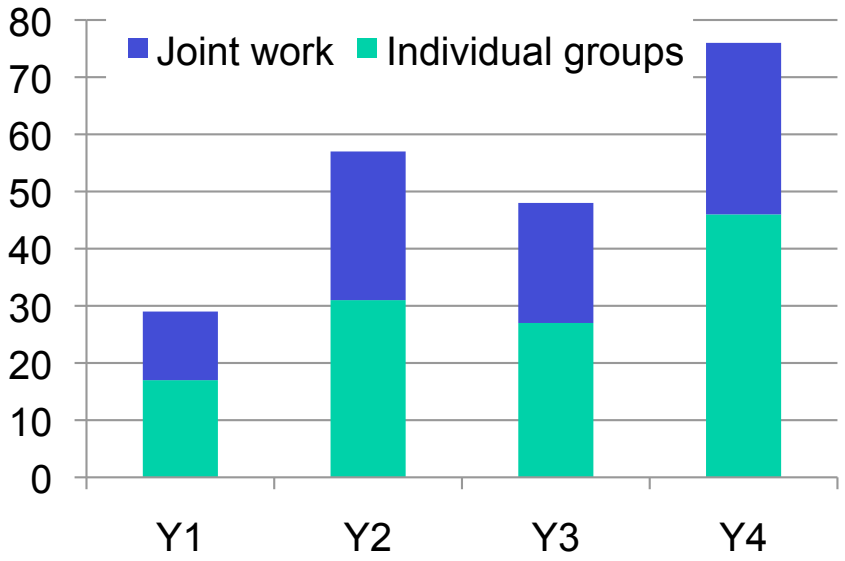

Other activities

- **16 workshops** (RTN, APRES, SOCNE…)
- **12 special sessions/tracks** (ETFA, INDIN…)
- **27 tutorials/seminars** (~all Artist Summer Schools….)
- **10 joint projects** (FP6/7-STREP, ARTEMIS, nationals)
- Continued participation in the **TinyOS Net2 Working Group** (*OpenZB stack*)
- **16 new collaborations** beyond Core Partners and Affiliated Partners

Global highlights

- **Protocols, tools and analysis for wireless networks**
  - ***WSN**: **toolset** to design, analyze, configure and deploy dense networks*
    - **OpenZB** protocol stack, **Z** monitor, **TinyOS** Net2 Working Group
    - **Visual tracking** and **localization** for ITS
  - ***MANETS**: RTDB middleware and protocol for **collaborating robots***
  - ***Industrial systems**: Real-time WiFi, WirelessHART (ISA 100)*
- **Related projects**

WASP** - Wireless Accessible Sensor Populations. **Contact: TUKL
EMMON** - EMbedded MONitoring. **Contact: ISEP-Porto
CONET** - Cooperating Objects Network of Excellence. **Contact: ISEP-Porto
FLEXWARE** - Flexible Wireless Autom. in Real-Time Env. **Contact: Catania
IPERMOB** - Perv. Hetero. Infrast. to Control Urban Mobility in Real-time. **Contact: Pisa,

Evidence

Global highlights

- **Protocols and middleware for robust and flexible real-time communication**
  - *Modeling and analysis suites: for distributed embedded systems (**MAST**)*
  - ***Ethernet**: new analysis (**AFDX, AVBs**), new tools (**FTT-SE / HaRTES**)*
  - ***CAN**: new analysis, **(Re)CANcentrate**, topology optimization*
  - ***RT middleware**: analysis, new middleware (**iLAND**, **HI** systems)*
- **Related projects**

iLAND** - mIddLewAre for deterministic dynamically reconfigurable NetworkeD embedded systems. **Contact: Madrid-UC3M, UnivPorto
HaRTES** - Hard Real-Time Ethernet Sw. **Contact: Aveiro, UnivPorto, Mallorca
CANbids** - CAN infra. for dependable systems. **Contact: Mallorca, UnivPorto
MADES** - UML / MARTE based model-driven approach. **Contact: York
INDEXYS** - INDustrial EXploitation of the genesYS cross-domain architecture. **Contact: TUKL and NXP

The end ?

- **Steady on-going collaborations**

  *Involving **27 groups** across the world*

- Several **projects starting** or continuing
- Integration in many **complementary communities**
- …

Check our wiki

http://twiki.fe.up.pt/bin/view/ArtistDesign

Towards a real-time connected world