Highly Scalable Aggregate Computations in Cyber-Physical Systems:
Physical Environment Meets Communication Protocols

Eduardo Tovar
CISTER/IPP-HURRAY Research Group (http://www.cister.isep.ipp.pt)
School of Engineering (ISEP) of the Polytechnic of Porto (IPP)
Porto, Portugal
emt@dei.isep.ipp.pt

Credits to: Nuno Pereira, Mário Alves and Björn Andersson

ISEP
INSTITUTO SUPERIOR
DE ENGENHARIA DO PORTO
POLITÉCNICO DO PORTO

CISTER
Research Centre in
Real-Time Computing Systems
FCT Research Unit 608

Outline

- What are Cyber Physical Systems?
- The challenge: large-scale, dense sensor networks
- A co-design approach for efficient data processing
- Dominance in wireless

What are Cyber-Physical Systems? (1)

- it looks like Cyber-Physical Systems (CPS) is the latest "hot topic" in the ICT domain (at least in the US ;-))
  - main initial event:
    - (Oct. 2006): NSF Workshop on Cyber-Physical Systems, Austin, TX
      - a couple of EU researchers attended
  - since then (list not complete):
    - (2006-2007): a number of sister NSF workshops
    - CPS track in RTSS 2007, RTSS 2008
    - EU-US Workshop on WSN, CPS and **Beyond**, Edinburgh, July 2007
    - CPS track in ICDCS 2008; CPS-related Satellite Workshop to ICDCS
    - CPS-CA'08 (DCOSS'08 satellite workshop)
    - CPS track in RTAS 2008
    - CPSWeek (RTAS, IPSN, HSCC), 2008 and 2009.
      - ArtistDesign (Berkeley group though ;-)) involvement in a Workshop (From Embedded Systems to Cyber-Physical Systems: a Review of the State-of-the-Art and Research Needs)
    - etc., etc.

What are Cyber-Physical Systems? (2)

- then, it looks like that
  - the Real-Time (RT) systems community is the driving force behind the eagerness
    - although, whatever CPS are, a number of people claims that the differentiating element interdisciplinary (CS; Control; SE; Signal Processing; Physics; Networking; etc.)
      - why mainly RT people?
  - a number of CPS research groups have been popping / spinning / recasting out of RT research groups / centers, e.g.:
    - UIUC, Uva, Vanderbilt, UPenn, Porto, etc.
  - a number of big initiatives have been put in place to tackle CPS, e.g,:
    - Center for Sensed Critical Infrastructure Research (CenSCIR), CMU

What are Cyber-Physical Systems? (3)

- why is that?
  - new funding opportunities?
    - yes!
  - better marketing the "old" RT and embedded computing stuff?
    - yes!
  - opportunities for appearing in scientific service (keynote talks, conference chair, workshop chair)?
    - yes!

What are Cyber-Physical Systems? (4)

- then, what the hell are CPS?
  - wikipedia (http://en.wikipedia.org/wiki/Cyber-physical_system; as of 14-Oct-2008):

"Cyber-physical systems (CPS) are computing systems that interact with physical processes. The tight integration between the computation and the physical system is what differentiates CPS from other forms of computing, making CPS a kind of embedded system. However, unlike more traditional embedded systems, CPS are typically designed as networks of interacting elements instead of as standalone devices.

The US National Science Foundation (NSF) has identified cyber-physical systems as a key area of research. Starting in late 2006, the NSF and NITRD sponsored several workshops on cyber-physical systems."

Is it that different?

Examples (also from wikipedia):

"*Common applications of CPS typically fall under sensor-based systems and autonomous systems. For example, many wireless sensor networks monitor some aspect of the environment and relay the processed information to a central node. Other types of CPS include autonomous automotive systems, medical monitoring, process control systems, distributed robotics, and automatic pilot avionics."*

What are Cyber-Physical Systems? (5)

- so, CPS seems to be really nothing new…
  - embedded systems have been in place for a long time and these systems often combine physical processes with computing ☹
  - physical interactions require timing requirements to be taken into account ☹
  - networks of embedded systems have been in place since a long time ago ☹

What are Cyber-Physical Systems? (6)

- …or is it the case that CPS (whatever it is) is really something new?
  - probably differences result from **the massive use** of networked embedded computing devices (**large scale** and **dense deployments**) ✓
    - ⇒ ubiquity; high resolution of physical perception

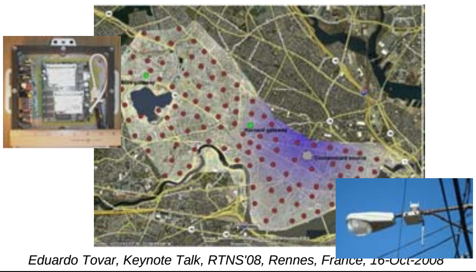

What are Cyber-Physical Systems? (7)

- …or is it the case that CPS (whatever it is) is really something new? (ctnd.)
  - and probably differences result from the confluence of networking, personal computing and embedded systems ✓

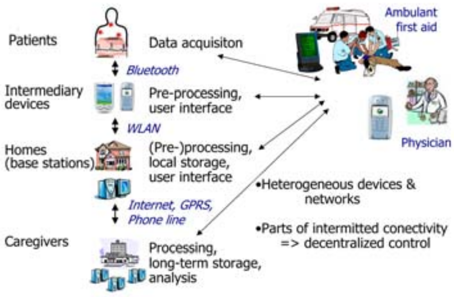

What are Cyber-Physical Systems? (8)

- prospective applications (1)
  - environmental monitoring
  - ambience intelligence / smart places (homes, public spaces, stores)
  - critical physical infrastructures (e.g. bridges, tunnels) monitoring
  - homeland security
  - utilities transportation systems (e.g. electrical, gas, water, oil)
  - factory automation and process control in large plants
  - domotics (home/building automation)
  - park/forest hazard monitoring/management
  - sports/religious/cultural events
  - disaster management (search & rescue in buildings/mines)
  - health care monitoring/management (e.g. hospitals)
  - intelligent transportation systems (e.g. highways, trains, metro)
  - games
  - participatory sensing
  - social networks
  - etc.

What are Cyber-Physical Systems? (9)

- more on applications (2)
  - environmental

  - physical infrastructures

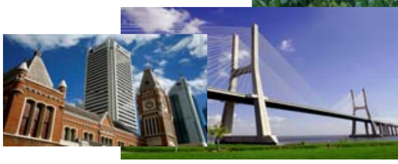

  - precision agriculture

What are Cyber-Physical Systems? (10)

- more on applications (3)
  - industrial plants

  - logistics

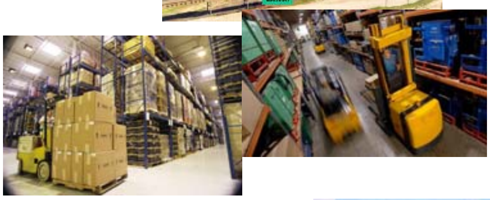

  - target tracking

What are Cyber-Physical Systems? (11)

- more on applications (4)
  - assisted living & smart healthcare
    - ALARM-Net project (UVa)
    - CodeBlue project (Harvard)

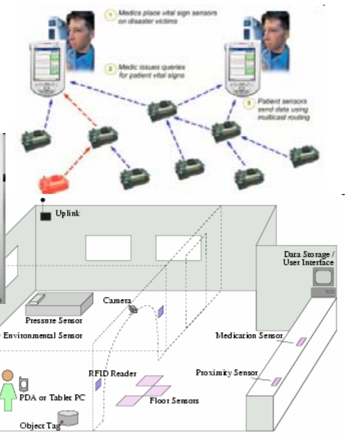

What are Cyber-Physical Systems? (12)

- more on applications (5)
  - structural monitoring
    - the BP preventive maintenance tanker initiative

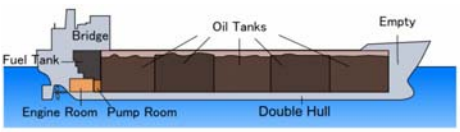

What are Cyber-Physical Systems? (13)

- other terms to designate CPS
  - embedded systems
  - deeply embedded systems
  - the Internet of Things
    - (not to be mistaken with the Internet of Dummy Things)
  - Cooperating Objects
    - FP7 NoE CONET: www.cooperating-objects.org

What are Cyber-Physical Systems? (14)

- CPS-CA'08 Definition of CPS
  - a (strong) Cyber-Physical System is a (computer) system which interacts with its physical environment and which is aware (in terms of data structures, abstractions, …) of its physical environment at run-time
    - what does awareness mean?
  - a (weak) cyber-physical system is a system that interacts with its physical environment

What are Cyber-Physical Systems? (15)

- CPS challenges
  - so many but, importantly
    - re-defining computing foundations (some examples)
      - Edward A. Lee, “Cyber-Physical Systems - Are Computing Foundations Adequate?”, 2006.
        - » (distributed) programming paradigms (languages still lacking temporal semantics, suitable concurrency models and hardware abstractions)
        - » networking protocols with timeliness as a structuring concern
        - » systems theory that combines "physical concerns" (control systems, signal processing, etc.) and "computational concerns" (complexity, schedulability, computability, etc.)
      - Alan Burns, “Time for Cyber-Physical Systems”, 2006/2008
        - » the notion of time-bands
        - » computer systems able to express: “these two events must occur simultaneously”; “this door must be closed all times”

Outline

- What are Cyber Physical Systems?
- **The challenge: Large-scale, dense sensor networks**
- A co-design approach for efficient data processing
- Dominance in wireless

The Challenge: Large-Scale, Dense Sensor Networks (1)

- Moore's law
  - cost (and size) of a single embedded computer node with sensing, processing and (wireless) communication capabilities drops towards zero
    - economically feasible to deploy very large and dense computer networks of such nodes
      - to take very large number of sensor readings from the physical world
      - to compute quantities and take decisions out of those sensor readings
    - the trend is to connect embedded computers through communication networks in order to collaboratively infer and control the state of the physical processes

The Challenge: Large-Scale, Dense Sensor Networks (2)

- large scale, dense sensor deployments
  - can cover a large area
  - can offer a better resolution
    - higher quality of sensing/control (e.g., capability of detecting the occurrence of an event)
  - but typically, applications are not interested in all sensor readings, but in computing a function based on sensor readings
    - e.g., MIN or AVERAGE
    - more complex functions
      - e.g., finding the most likely location of an object based on sensor readings

The Challenge: Large-Scale, Dense Sensor Networks (3)

- but, these networked embedded computers are
  - resource-constrained
    - typically battery-operated
    - with reduced computing and communication capabilities
    - therefore energy-efficient operation is important
  - and, because of the physical interaction, it is often necessary that the delay from sensing until actuation (decision) is low and bounded

The Challenge: Large-Scale, Dense Sensor Networks (4)

- … the challenge is then how
  - to perform **scalable** and **efficient** information processing in such large-scale, dense cyber-physical systems
    - with:
      - (i) low delay
      - (ii) low resource usage
  - what do we mean by scalability and efficiency?
    - **"efficient information processing"**
      - the desired computation is performed while consuming very little resources (energy, communication links, memory, processor)
    - "**scalable**"
      - consumption of resources increases slowly or not at all as the number of sensor readings to be processed and/or the number of embedded computer nodes increases

The Challenge: Large-Scale, Dense Sensor Networks (5)

- three approaches to tackle the challenge
  - approach A1
    - design **distributed algorithms** for sensor data processing such that they request to use resources and the utilization of resources is low
  - approach A2
    - design underlying **networked distributed computing systems with corresponding resource management schemes** such that given request to use the resources, the utilization of resources is low
  - approach A3
    - **co-design** distributed algorithms for sensor data processing and underlying networked distributed systems with corresponding resource management schemes such that the utilization of resources is low

The Challenge: Large-Scale, Dense Sensor Networks (6)

- another way to describe these approaches or…
  - how these approaches relate to CPS?
    - A1 deals with **Physical** systems
    - A2 deals with information processing (**Cyber**) systems
    - and A3 deals with the tight interaction between the physical and information processing (**Cyber-Physical**) system

The Challenge: Large-Scale, Dense Sensor Networks (7)

- A1 efforts
  - distributed computing
    - to achieve consensus
    - leader election
    - coloring
  - gossiping, rumor, infection protocols
  - down-sampling, pre-processing, selective reading
  - etc.
- A2 efforts
  - resource-aware operating systems
  - TDMA-based comm. protocols which reduce energy consumption and ensure that collisions will not occur
  - topologies (e.g., logical hexagonal) for fast convergecast without data aggregation
  - etc.

The Challenge: Large-Scale, Dense Sensor Networks (8)

- A3 efforts
  - is the focus of the rest of this talk…

Outline

- What are Cyber Physical Systems?
- The challenge: Large-scale, dense sensor networks
- **A co-design approach for efficient data processing**
- Dominance in wireless

A co-design approach for efficient data processing (1)

- the problem of performing scalable and efficient information processing in large-scale CPS must be solved
  - otherwise the usefulness of large scale, dense deployments is reduced significantly
- we believe that it is important to take a "clean-slate" approach (as approach A3 does)
  - in order to attain the best possible performance for systems in the long term

A co-design approach for efficient data processing (2)

- consider the (simple) problem of computing a **simple** aggregate quantity such as MIN:
  - the minimum (MIN) sensed temperature (or other physical quantity) among the nodes at a given moment
  - assume the following as a large and dense deployment
    - 10 nodes (just for the sake of exemplification)

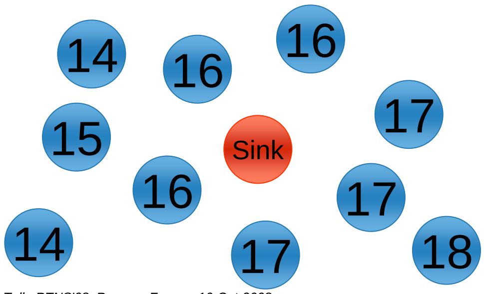

Sorry, it does not seem that large and dense ☹

Eduardo Tovar, Keynote Talk, RTNS'08, Rennes, France, 16-Oct-2008 29

A co-design approach for efficient data processing (3)

- MIN is trivial, but for systems with large and dense deployment of nodes (such the one in Fig. below 😐)
  - time-complexity as a function of number of nodes (no scalability)
    - this is true even if in-network data aggregation (convergecast trees) is used
      - since density reduces opportunities for parallel transmission
  - 10 msgs

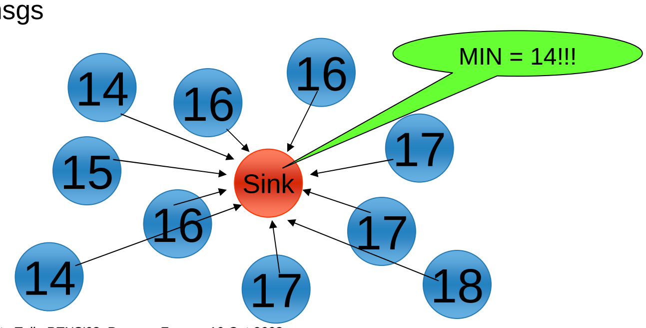

Eduardo Tovar, Keynote Talk, RTNS'08, Rennes, France, 16-Oct-2008 30

A co-design approach for efficient data processing (4)

- we have an ambition, though:
  - compute MIN with a time-complexity that is independent of the number of nodes
    - in fact, with a time-complexity that is equivalent to the time of transmitting a single message
      - only possible if all send at the same time…

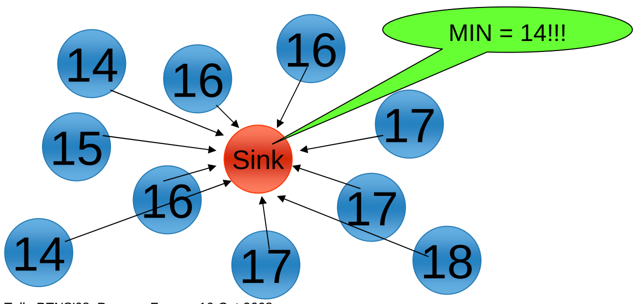

A co-design approach for efficient data processing (5)

- is such a medium access control (MAC) possible?
  - CAN (Controller Area Network) uses a dominance/binary-countdown protocol
    - developed by Robert Bosch GmbH
    - originally for the automotive industry
    - widely used in many other areas
      - building automation, industrial control, monitoring, ...
    - millions of nodes and systems deployed
  - characteristics of CAN
    - designed for a wired bus
    - each node (message) has a unique identifier (=priority)
      - lower values for priority mean higher priority
    - resolve bus contention using a bitwise arbitration (non-destructive collision)
      - if a node sends a '1' but hears a '0', he loses
    - notion of recessive and dominant bits
      - 0 is dominant; 1 is recessive
    - bus implements a wired-AND

A co-design approach for efficient data processing (6)

- example CAN network
  - three CAN nodes, with priorities (IDs) 4, 2 and 1

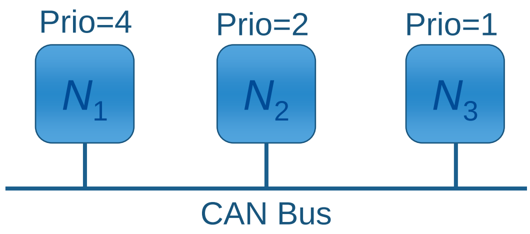

A co-design approach for efficient data processing (7)

- example (ctnd.)

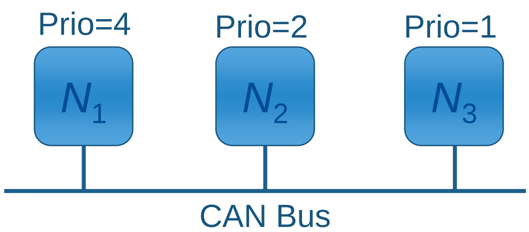

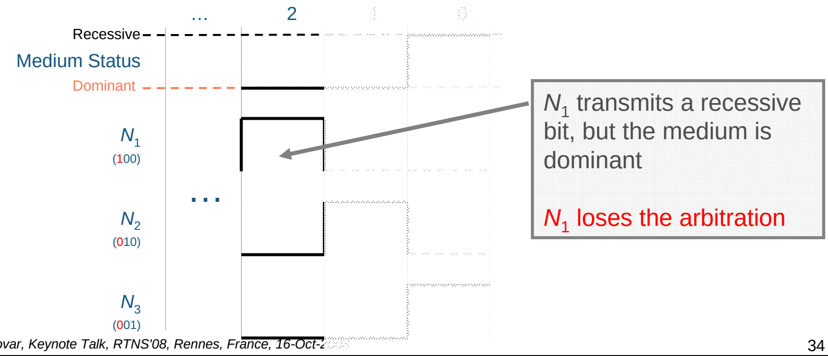

A co-design approach for efficient data processing (8)

- example (ctnd.)

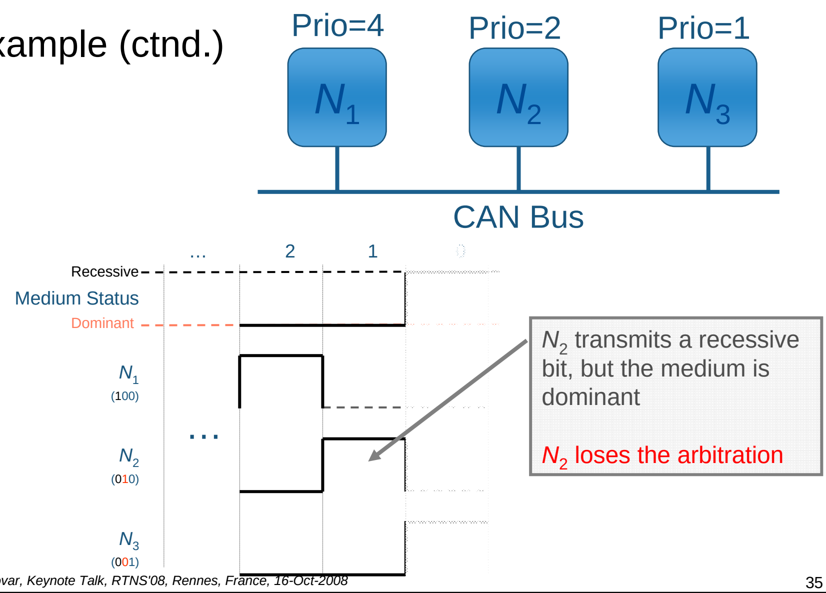

A co-design approach for efficient data processing (9)

- example (ctnd.)

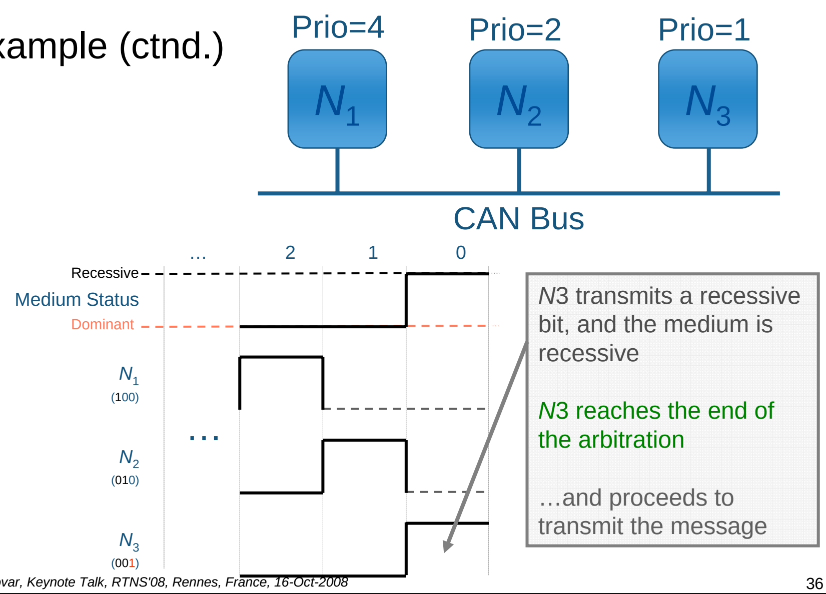

A co-design approach for efficient data processing (10)

- example (ctnd.)

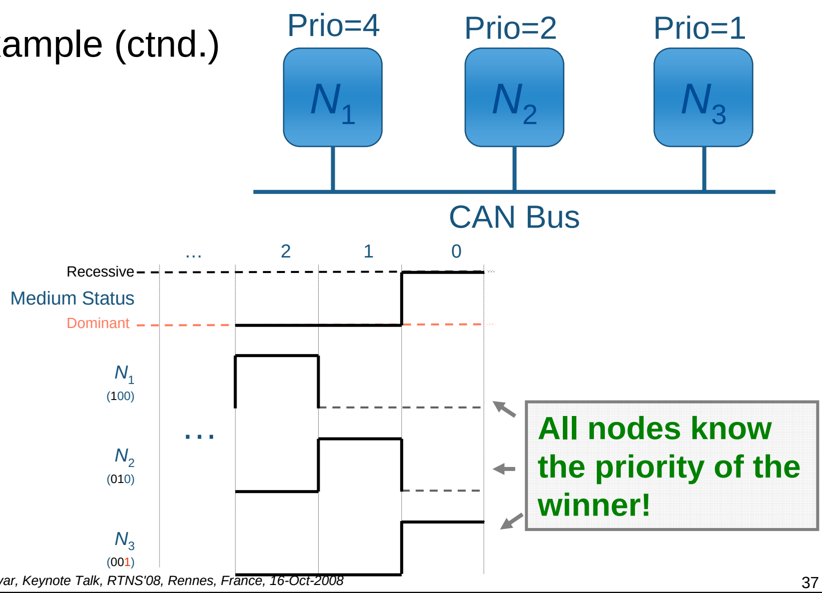

A co-design approach for efficient data processing (11)

- we propose to use the contention field differently of CAN
  - during runtime, the contention (or priority) field is computed **as a function of the physical quantity** (or characteristic) of interest
    - it is a **Physical Dynamic Priority Dominance ((PD)^2) protocol**
      - this will be an important building block for computing aggregate quantities with a low time-complexity
      - the (PD)^2 protocol is an example **where communication and computation is tightly coupled with the physical environment**
        - » a clear co-design feature

A co-design approach for efficient data processing (12)

- by using this principle then
  - a panoply of other aggregate quantities / functionalities can be computed offering low-time complexity
    - MAX (obvious)
    - estimation of COUNT
    - estimation of MEDIAN
    - approximate interpolations
    - localizing objects
    - etc.

Let's see how these work

A co-design approach for efficient data processing (13)

- estimation of COUNT (1)
  - intuition behind it
    - if the contention field is a nonnegative random number obtained at runtime, then the probability that the minimum value of the contention field is 0 approaches 1 as the number of nodes get very large
    - however, if there are only a few nodes, then it is highly unlikely that the minimum among the random values is zero
    - it is then possible to estimate the number of nodes by computing the MIN of the random numbers
      - this can with *k* iterations and using Maximum Likelihood estimation
      - MIN is not a function of a sensed physical quantity, instead it is a function of a physical reality

A co-design approach for efficient data processing (14)

- estimation of COUNT (2)

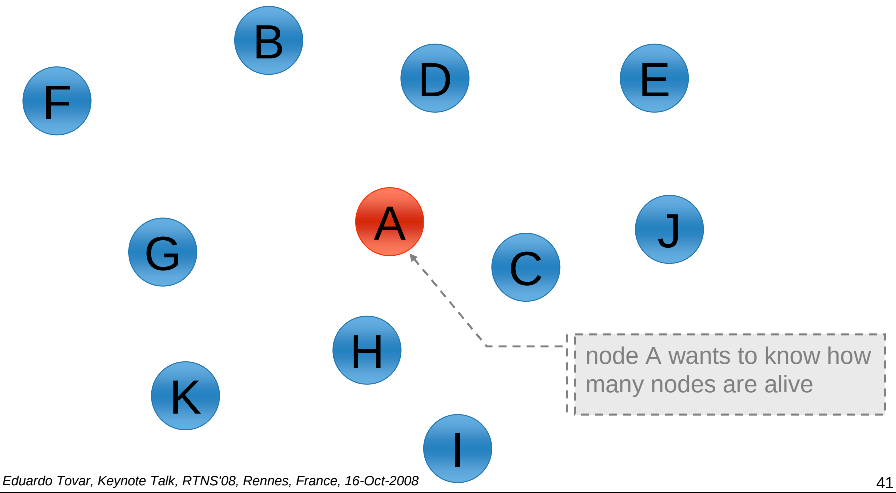

A co-design approach for efficient data processing (15)

- estimation of COUNT (3)

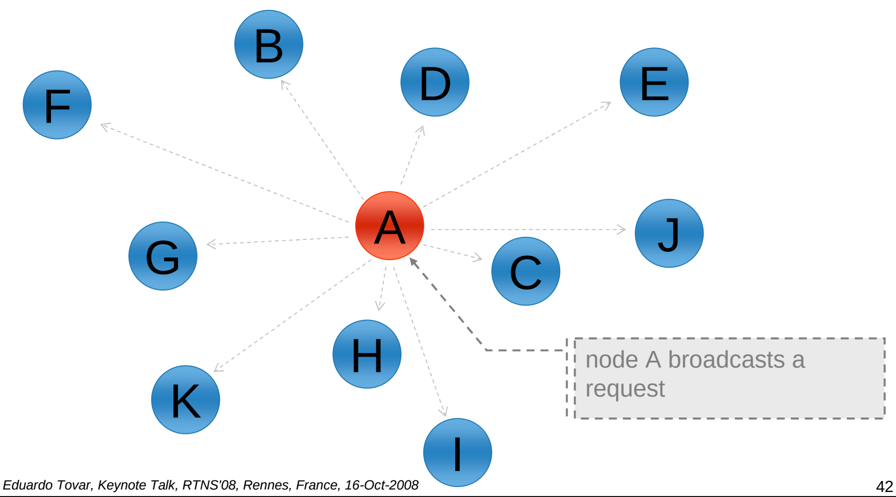

A co-design approach for efficient data processing (16)

- estimation of COUNT (4)

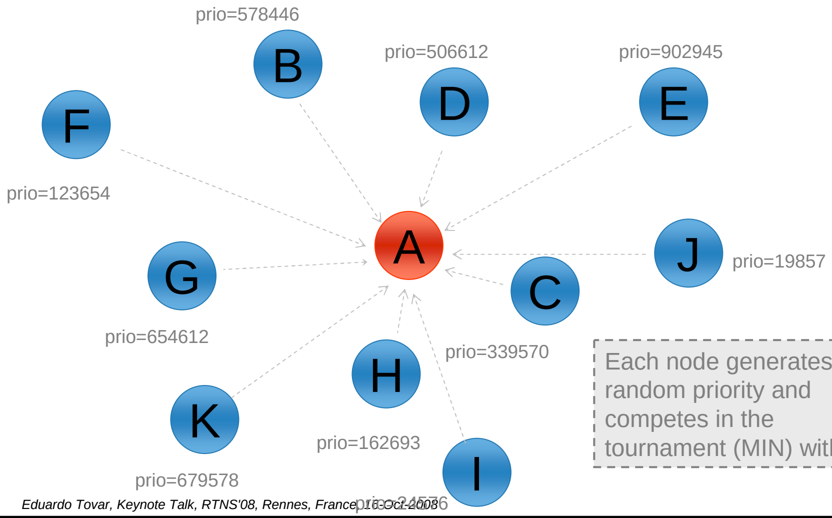

Each node generates a random priority and competes in the tournament (MIN) with it…

A co-design approach for efficient data processing (17)

- estimation of COUNT (5)

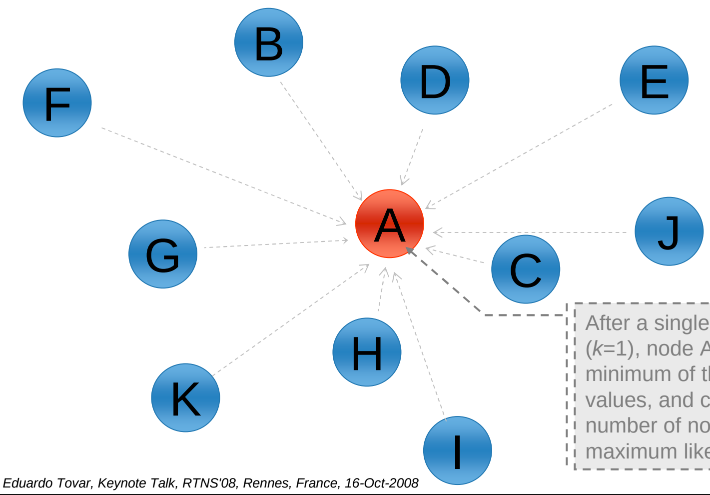

After a single tournament (*k*=1), node A knows the minimum of the random values, and computes the number of nodes based on a maximum likelihood estimator

A co-design approach for efficient data processing (18)

- estimation of COUNT (6)

Algorithm 1 Estimating COUNT (the number of nodes)

Require: All nodes start Algorithm 1 simultaneously.
Input: *active* - a global boolean variable indicating if the node is considered in the COUNT

1: **function** $nnodes(j$: integer, x : array$[1..k]$ of integer$)$ **return** a real
2: r : array$[1..k]$ of integer
3: x : array$[1..k]$ of integer
4: q : integer
5: **for** $q \leftarrow 1$ **to** k
6: **if** $(active = \text{TRUE})$ **then**
7: $r[q] \leftarrow random(0, MAXV)$
8: **else**
9: $r[q] \leftarrow MAXV$
10: **end if**
11: $x[q] \leftarrow$ send_empty$(r[q])$
12: **end for**
13: **if** $(\exists q : x[q] = MAXV)$ **then**
14: $est\_nodes \leftarrow 1$
15: **else**
16: $est\_nodes \leftarrow ML\_estimation(x[1], x[2], ..., x[k])$
17: **end if**
18: **return** est_nodes *// the estimation of COUNT*

Algorithm 2 Function ML_estimation

Require: The division of two integers (as is done in line 6) returns a real number.

1: **function** $ML\_estimation(x$: array$[1..k]$ of integer$)$ **return** an integer
2: v : array$[1..k]$ of real
3: $sumv, q$: integer
4: $sumv \leftarrow 0$
5: **for** $q \leftarrow 1$ **to** k
6: $v[q] \leftarrow ln\left(\frac{1}{1-\frac{x[q]}{MAXV}}\right)$
7: $sumv \leftarrow sumv + v[q]$
8: **end for**
9: **return** $\lceil \frac{k}{sumv} \rceil$
10: **end function**

A co-design approach for efficient data processing (19)

- estimation of COUNT (7)
  - $k = 5$

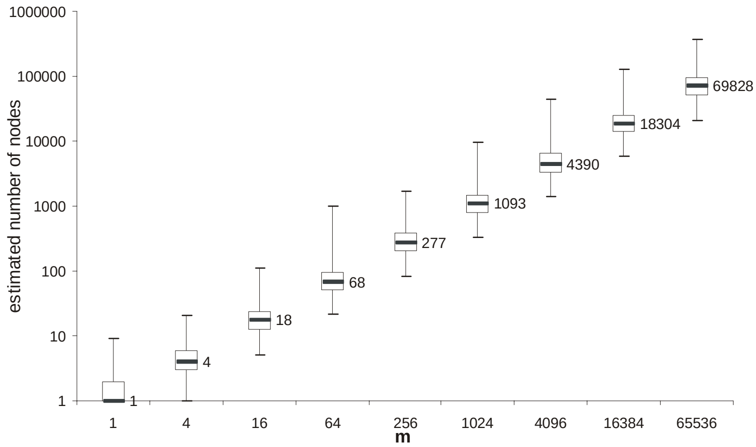

A co-design approach for efficient data processing (20)

- estimation of COUNT (8)
  - $k = 20$

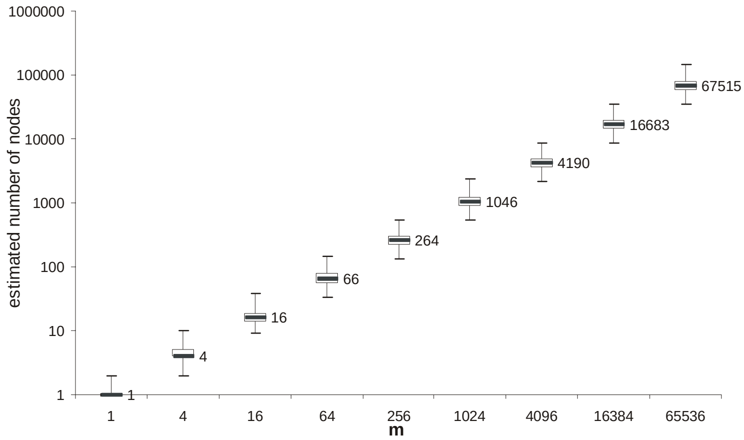

A co-design approach for efficient data processing (21)

- approximate interpolations (1)
  - consider a signal (say concentration of a hazardous gas) that varies with location (*x*,*y*)

how can this signal be obtained?

(assume a deployment of a dense network of sensing nodes)

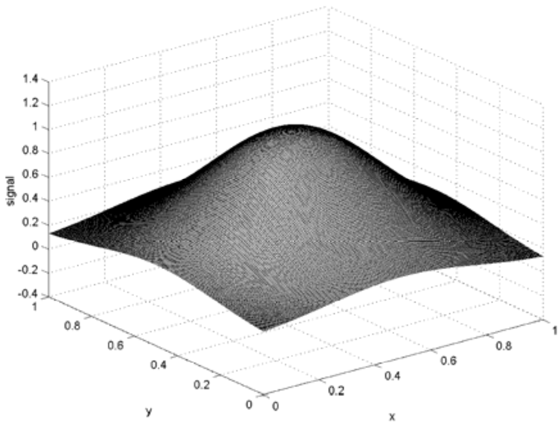

A co-design approach for efficient data processing (22)

- approximate interpolations (2)
  - consider a signal that varies with location (*x*,*y*)

the signal changes with time… and one wants to get the signal a few times per second…

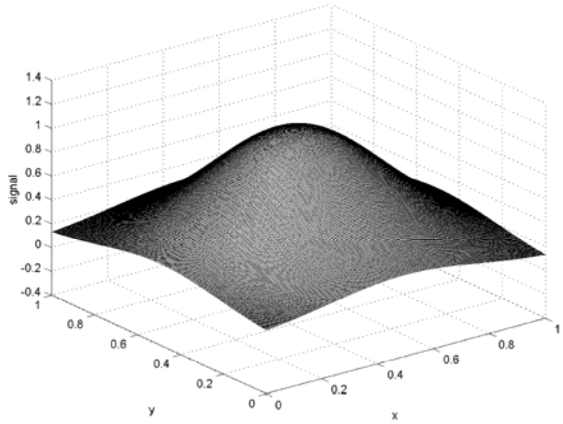

A co-design approach for efficient data processing (23)

- approximate interpolations (3)
  - consider a signal that varies with location (*x*,*y*)

how can this signal be obtained?

one can get sensor readings from **all** the nodes and then perform curve fitting… but it would be slow: O(*m*)

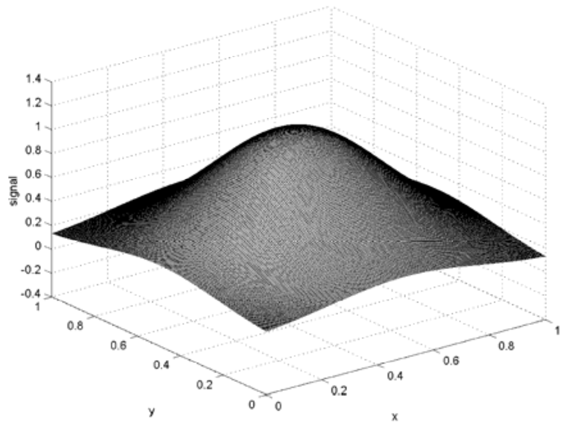

A co-design approach for efficient data processing (24)

- approximate interpolations (4)
  - consider a signal that varies with location (*x*,*y*)

how can this signal be obtained?

one can get sensor readings from **some of** the nodes and then perform curve fitting… but from which sensor nodes?

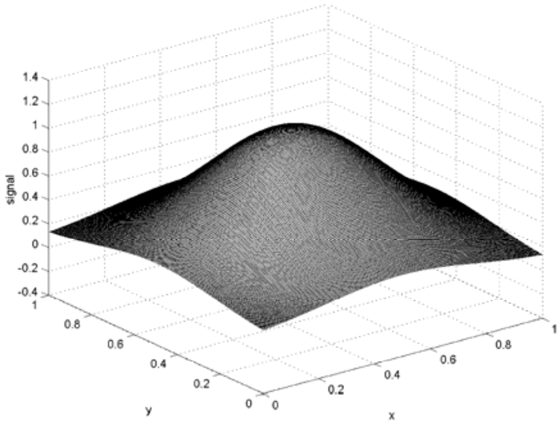

A co-design approach for efficient data processing (25)

- approximate interpolations (5)
  - consider a signal that varies with location (*x*,*y*)

how can this signal be obtained?

one can get sensor readings from **some of** the nodes and then perform curve fitting… but from which sensor nodes?

Select nodes randomly?

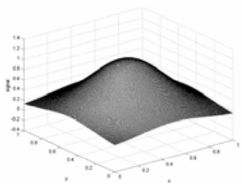

original signal

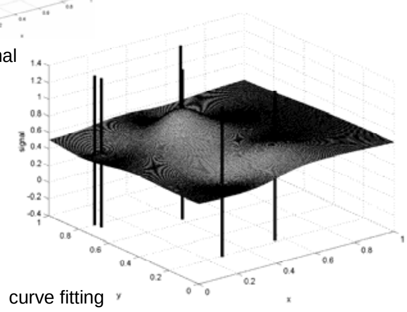

curve fitting

A co-design approach for efficient data processing (26)

- approximate interpolations (6)
  - consider a signal that varies with location (*x*,*y*)

how can this signal be obtained?

one can get sensor readings from **some of** the nodes and then perform curve fitting… but from which sensor nodes?

Select nodes randomly?

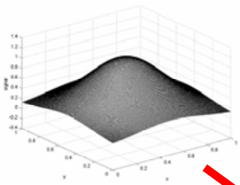

original signal

bad match

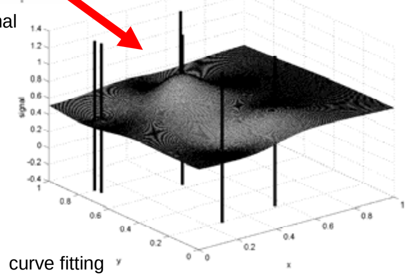

curve fitting

A co-design approach for efficient data processing (27)

- approximate interpolations (7)
  - consider a signal that varies with location (*x*,*y*)

how can this signal be obtained?

one can get sensor readings from **some of** the nodes and then perform curve fitting… but from which sensor nodes?

Carefully selected nodes?

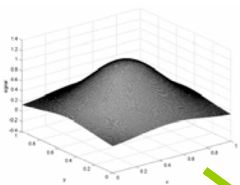

original signal

good match

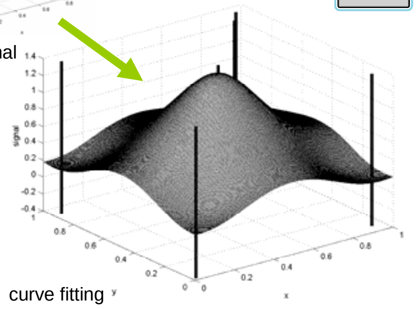

curve fitting

A co-design approach for efficient data processing (28)

- approximate interpolations (8)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

start with a flat surface no data points (nodes) selected

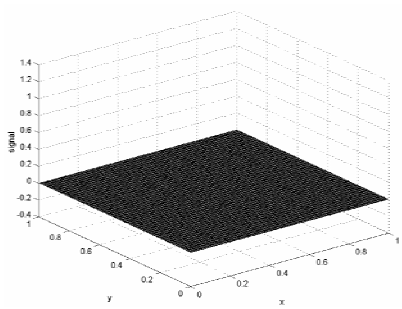

curve fitting

A co-design approach for efficient data processing (29)

- approximate interpolations (9)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

1 data point (node) selected

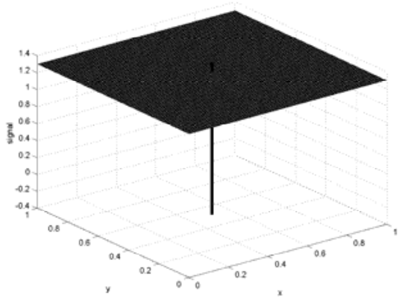

curve fitting

A co-design approach for efficient data processing (30)

- approximate interpolations (10)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

2 data points (nodes) selected

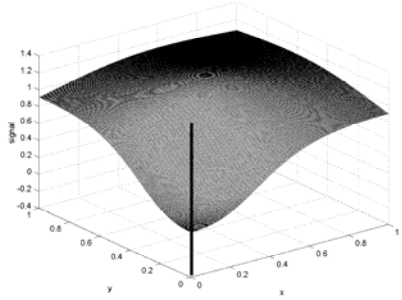

curve fitting

A co-design approach for efficient data processing (31)

- approximate interpolations (11)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

3 data points (nodes) selected

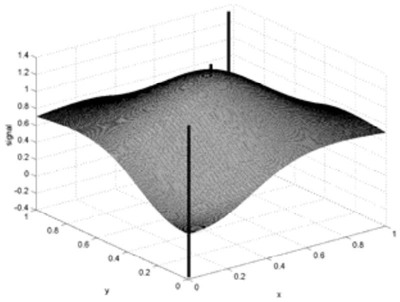

curve fitting

A co-design approach for efficient data processing (32)

- approximate interpolations (12)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

4 data points (nodes) selected

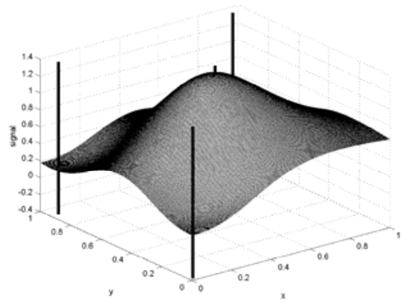

curve fitting

A co-design approach for efficient data processing (33)

- approximate interpolations (13)
  - consider a signal that varies with location (*x*,*y*)

how to carefully select nodes?

iteratively select the node with highest error

5 data points (nodes) selected

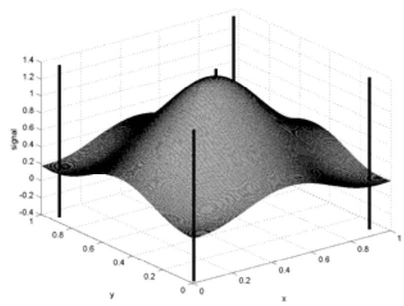

curve fitting

A co-design approach for efficient data processing (34)

- approximate interpolations (14)
  - consider a signal that varies with location (*x*,*y*)

the (PD)² protocol can be exploited to efficiently select the node with priority 1/error!

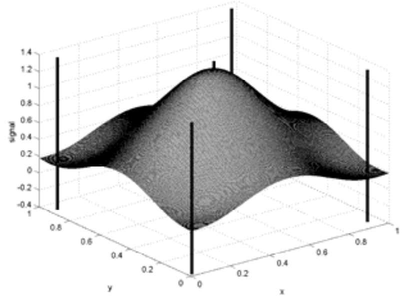

curve fitting

A co-design approach for efficient data processing (35)

- approximate interpolations (15)
  - consider a signal that varies with location (*x*,*y*)

tolerance to faulty values?

a single faulty value may jeopardize the whole approach

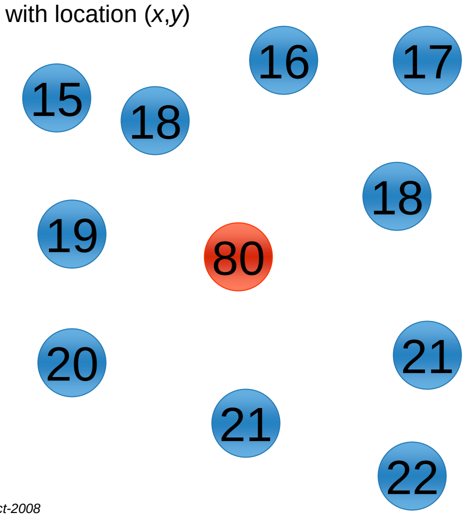

A co-design approach for efficient data processing (36)

- approximate interpolations (16)
  - consider a signal that varies with location (x,y)

tolerance to faulty values?

sensor readings are assumed to exhibit spatial locality

and

correct sensor readings are always in majority

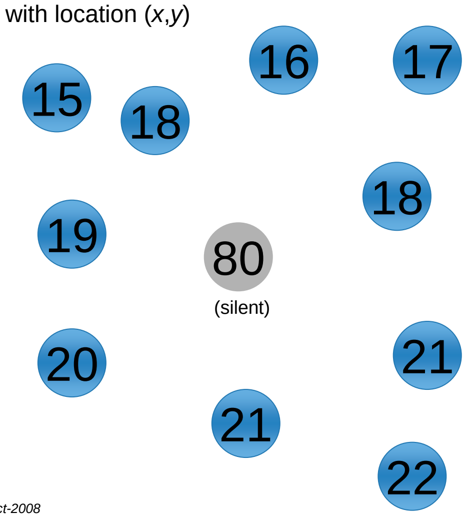

Outline

- What are Cyber Physical Systems?
- The challenge: Large-scale, dense sensor networks
- A co-design approach for efficient data processing
- **Dominance in wireless**

Dominance in wireless (1)

- the effectiveness of the $(PD)^2$ protocol in wired technology as already been proved
- obviously $(PD)^2$ would better leverage on wireless
  - is that possible?
    - potential problems (beyond the wild nature of radio communications)
      - wired-AND behavior
      - nodes need to listen while transmitting (observation: actually only when transmitting a recessive bit)

Dominance in wireless (2)

- we already mentioned Gordon Moore's law
- what about Arthur C. Clarke's laws (1917-2008)?

"When a distinguished but elderly scientist states that something is possible he is almost certainly right. When he states that something is impossible, he is very probably wrong"

"The only way of discovering the limits of the possible is to venture a little way past them into the impossible"

"Any sufficiently advanced technology is indistinguishable from magic"

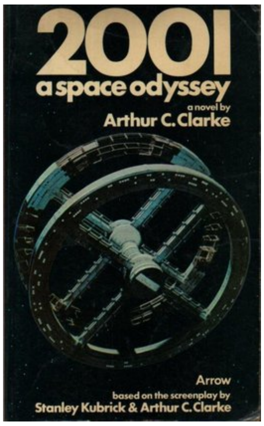

Dominance in wireless (3)

- or Gordon Bell's law (born 1934)?

"Roughly every decade a new, lower priced computer class forms based on a new programming platform, network, and interface resulting in new usage and the establishment of a new industry" (paper from 1972)

"**As of 2005**, the computer classes include:
- mainframes (**1960s**)
- minicomputers (**1970s**)
- PCs and workstations evolving into a network enabled by Local Area Networking (**1980s**)
- web browser client-server structures enabled by the Internet (**1990s**)
- small form-factor devices such as cell phones and other cell phone sized devices (c. 2000)
- wireless sensor networks, aka motes (c. **>2005**)
- home and body area networks (**> 2010**)"

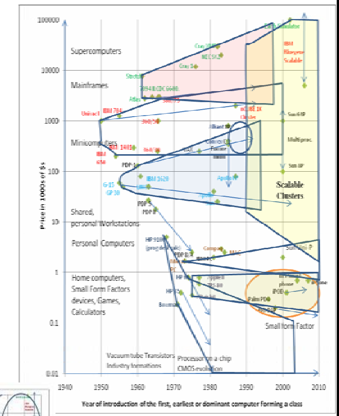

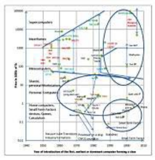

Dominance in wireless (4)

- or Robert Metcalfe's law (born 1946)?

"the value of a network grows as the square of the number of its users"

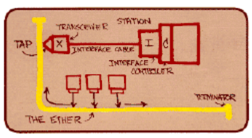

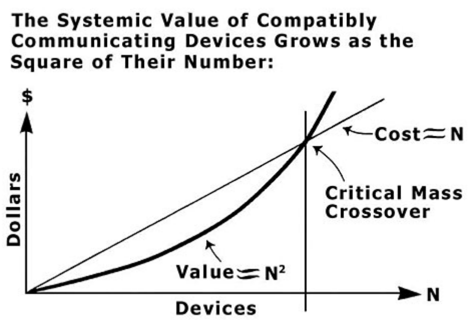

Dominance in wireless (5)

- but then there is…
- Edward Murphy's law (1918-1990)!

> "if there are two or more ways to do something, and one of those ways can result in a catastrophe, then someone will do it" (original version, c. 1952)

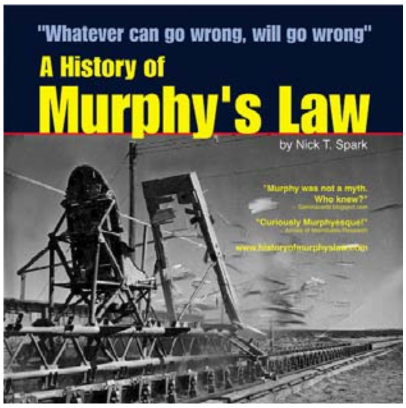

Dominance in wireless (6)

- WiDom (1)

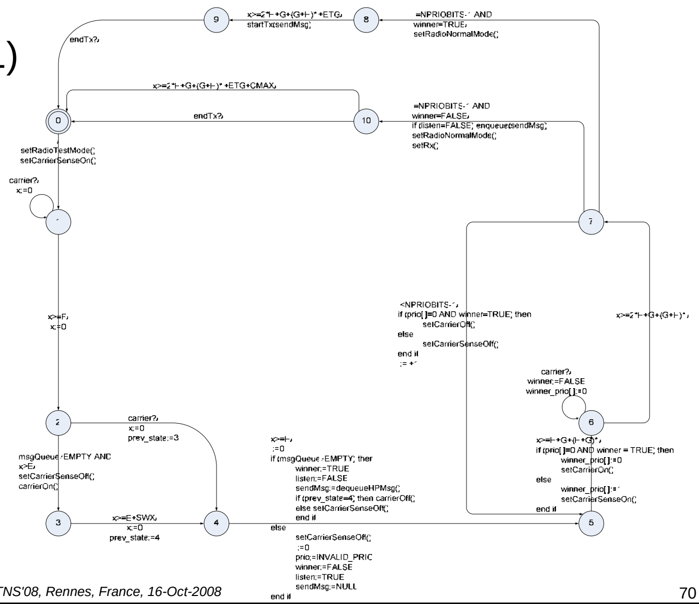

Dominance in wireless (7)

- WiDom (2)

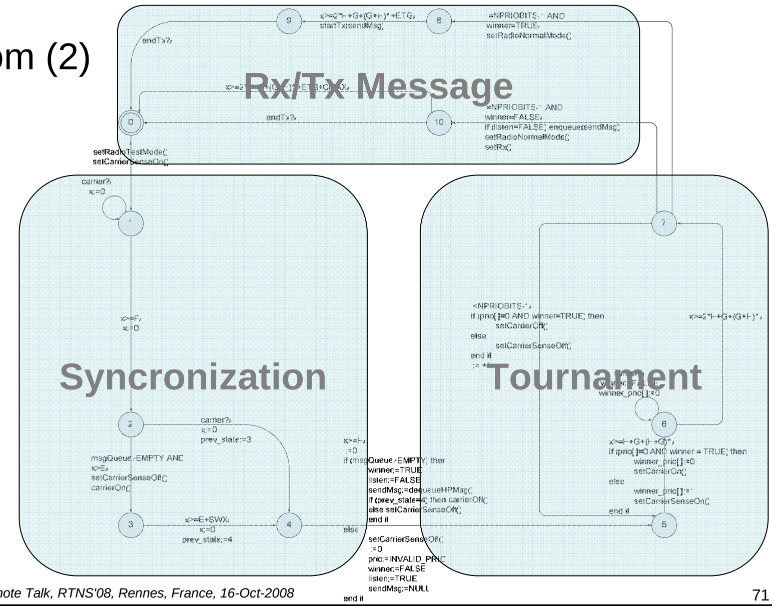

Dominance in wireless (8)

- WiDom (3)
  - synchronization

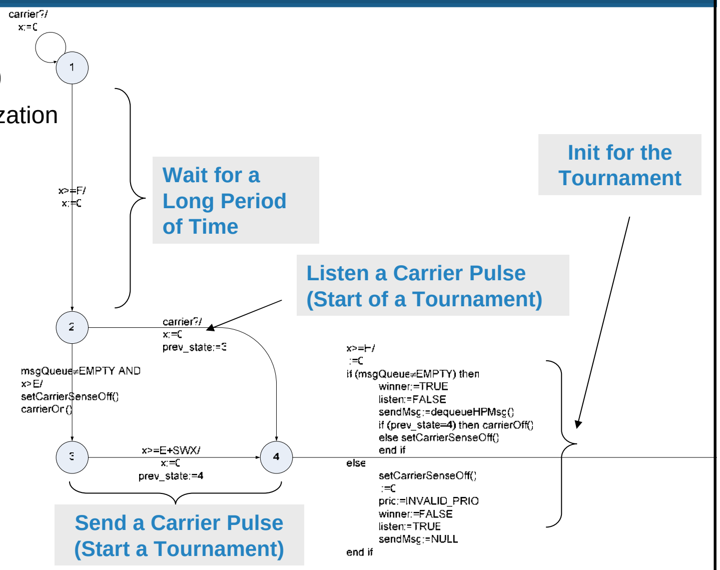

Dominance in wireless (9)

- WiDom (4)
  - tournement

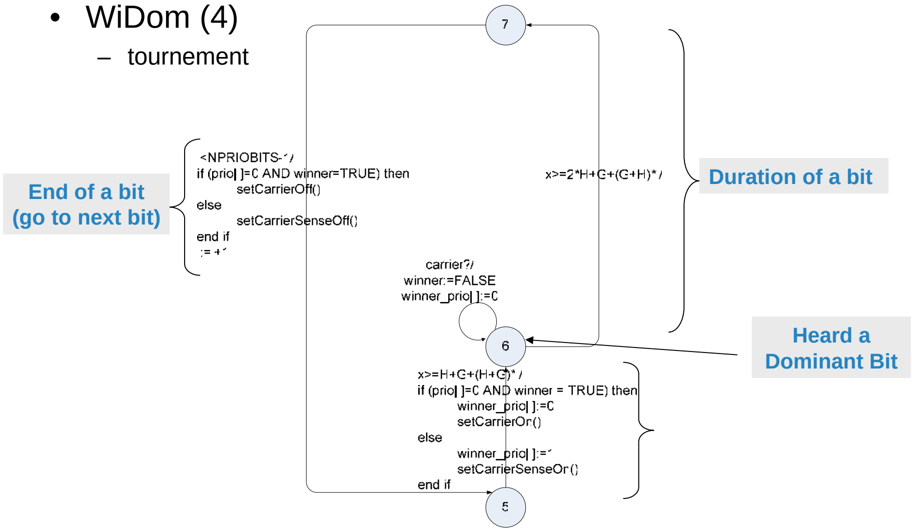

Dominance in wireless (10)

- WiDom (5)
  - Rx/Tx message

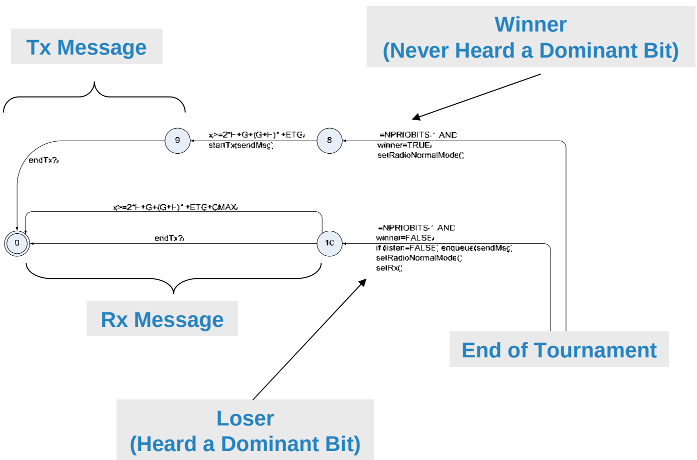

Dominance in wireless (11)

- WiDom: example of operation
  - Application, MAC protocol and Radio activity in two nodes

Priorities: N_1=010; N_2=011

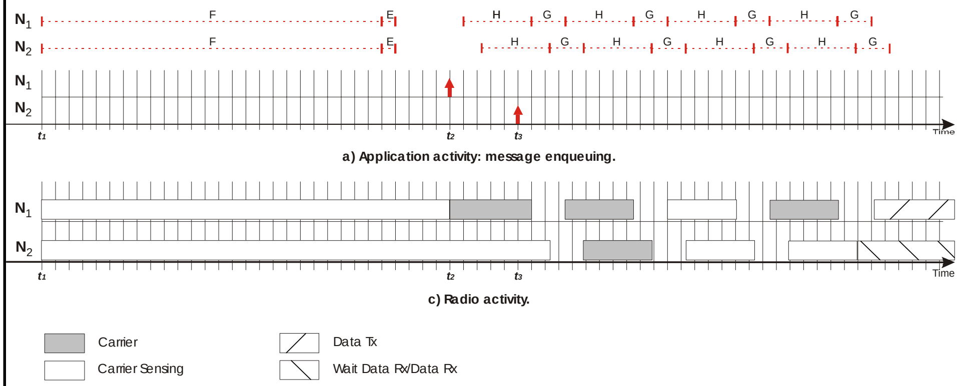

Dominance in wireless (12)

- WiDom: it crucial to minimize the overhead and energy
  - efforts in specialized hardware

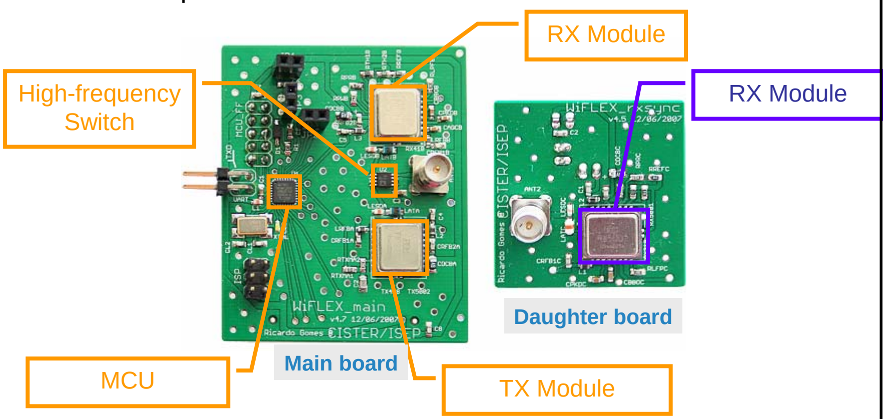

Dominance in wireless (13)

- WiDom: it crucial to minimize the overhead and energy
  - efforts in specialized hardware

on MicaZ

on CMU-FireFly

Dominance in wireless (14)

- some performance tests (1)
  - distance
    - 10 nodes divided in two groups
    - distance between the two groups varied from 1 to 20 meters in 1m steps
    - for each step, 150000 tournaments using random priorities were performed

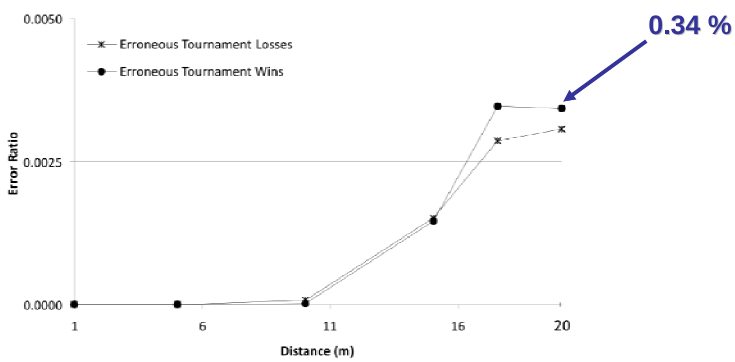

Dominance in wireless (15)

- some performance tests (2)
  - comparing the time to compute MIN with COTS (IEEE802.15.4) technology

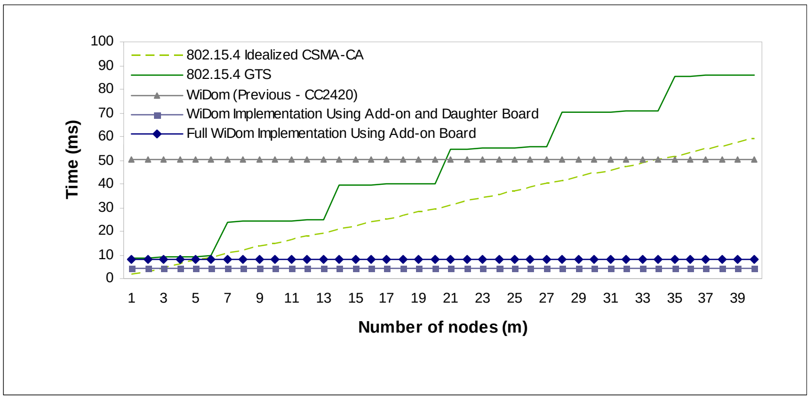

Dominance in wireless (16)

- extending to multiple broadcast domains (1)
  - (to allow the real large-scale)

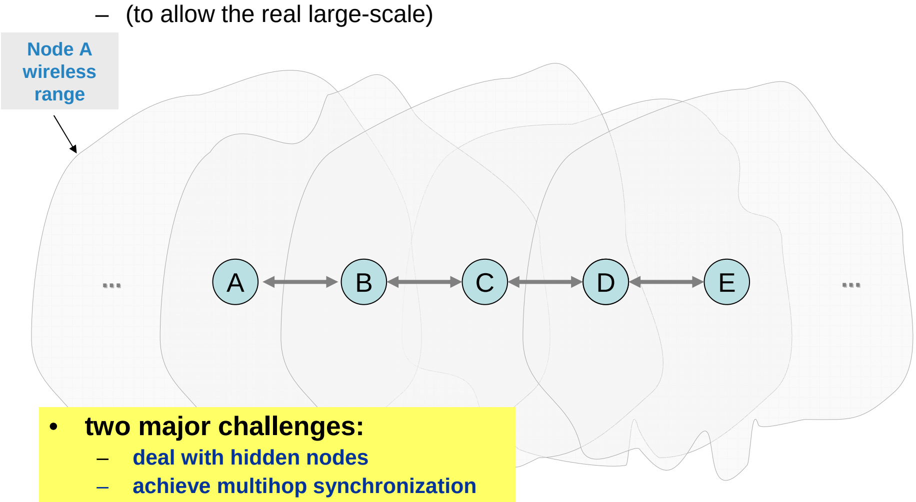

- **two major challenges:**
  - **deal with hidden nodes**
  - **achieve multihop synchronization**

Dominance in wireless (17)

- extending to multiple broadcast domains (2)
  - basic idea
    - re-propagate priority bits two-hops away

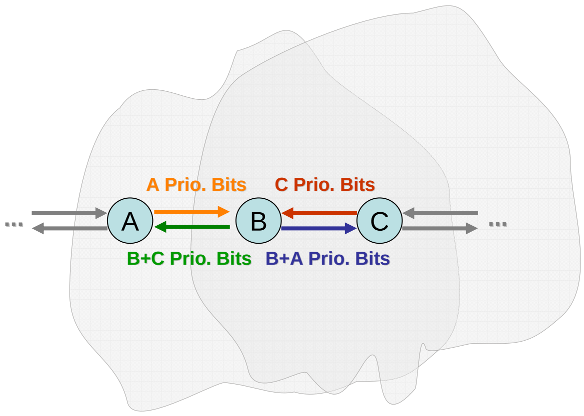

Dominance in wireless (18)

- extending to multiple broadcast domains (3)

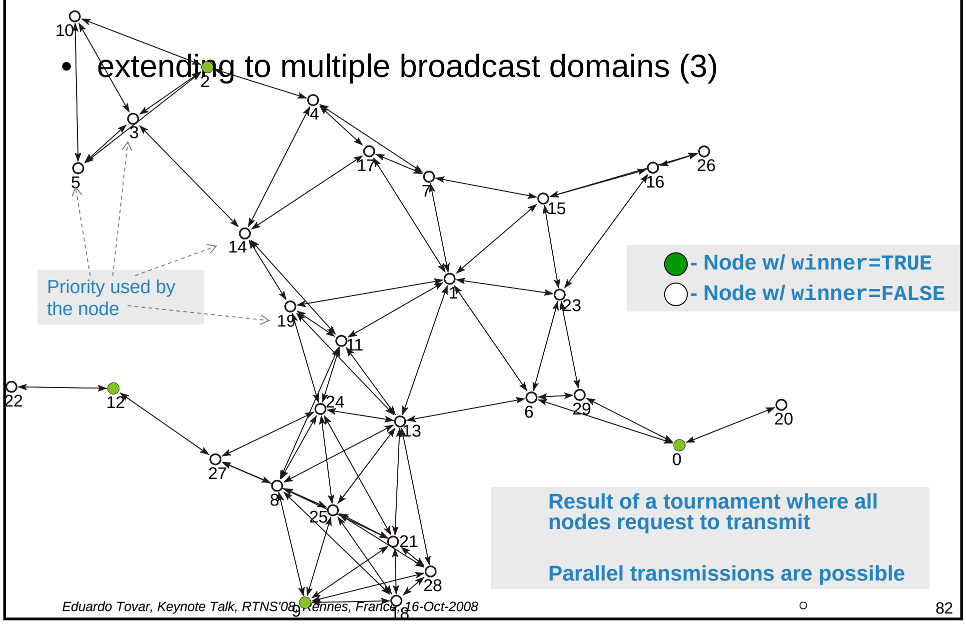

Dominance in wireless (19)

- extending to multiple broadcast domains (4)
  - aggregate quantities in multiple broadcast domains
    - algorithms to partition the network such that each partition is a broadcast domain

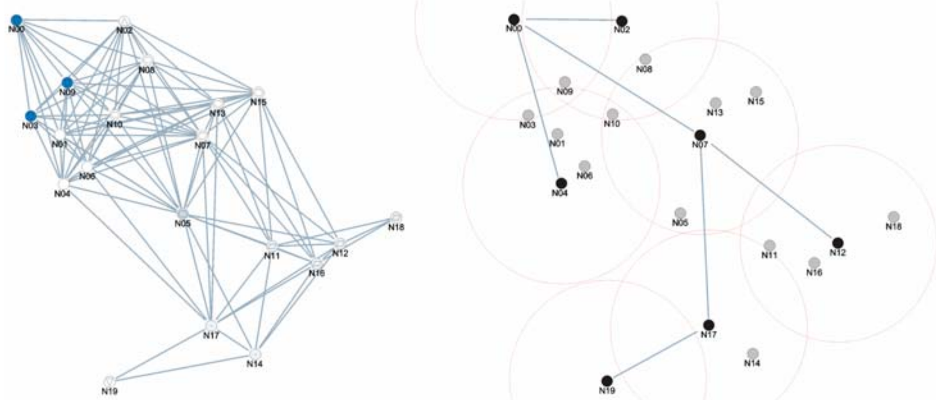

a) Network Example and Partitions Formed

b) Virtual Ranges of Partition Leaders

Conclusions

- the problem of performing scalable and efficient information processing in large-scale CPS must be solved
  - otherwise the usefulness of large scale, dense deployments is reduced significantly
- we believe that it is important to take a “clean-slate” approach (a co-design approach)
  - **co-design** distributed algorithms for sensor data processing and underlying networked distributed systems with corresponding resource management schemes such that the utilization of resources is low
  - in order to attain the best possible performance for systems in the long term

Where do I come from?

knew about these?






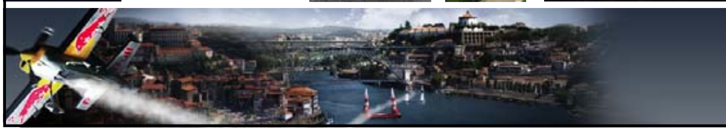

somewhere around here
(Porto, Portugal)