Consiglio Nazionale delle Ricerche

Istituto di Elettronica e di Ingegneria dell'Informazione e delle Telecomunicazioni

1st Italian Workshop on Embedded Systems

IWES 2016

September 19-20, 2016 — Pisa, Italy

Reliable Wireless Communications for Industrial Networked Systems

G. Cena et Al.

CNR-IEIIT (Torino)

Wireless in industrial scenarios

- Wireless communication is typically deemed not reliable enough for the use in real-time control systems
  - *Interference* from nearby wireless nodes and networks, especially those based on CSMA/CA mechanisms and the like
  - *Disturbance* from industrial equipment generating electromagnetic noise, including (multipath) fading effects
- But it seem many people are interested in it…
  - A lot of papers have appeared in the past decade in the *scientific literature* that propose techniques aimed at improving reliability
  - Vendors as well are offering *commercial* wireless solutions specifically tailored for industry (e.g., WirelessHART)
  - Applications of wireless networks (like *cable replacement*) are often advantageous and sometimes necessary

Reliable wireless communications

- When dealing with wireless networked control systems, the term “*reliability*” has to be somehow redefined
  - Radio communication suffers from unpredictable *transmission latencies* and higher *packet losses*
  - Hardly they can be used in *hard real-time* contexts
  - Unsuitable when *determinism* is a strict requirement
- Developing a brand new wireless transmission technology *explicitly* for the industry is likely not the best option
  - Too *expensive* (no synergies with the ICT world)
  - *Coexistence* may be a problem (spectrum is already crammed)
  - Most vendors prefer to rely on existing, well settled, proven solutions: *IEEE 802.15* and *IEEE 802.11*

Available wireless technologies

- COTS wireless solutions for industry (e.g., *WirelessHART*) rely on specific mechanisms for enhancing reliability
  - IEEE 802.15 with *channel hopping and blacklisting*
  - but they are *not* particularly *fast* (250 Kb/s)
  - and are *not* directly *interoperable* with Ethernet at the data-link layer
- Our work mostly focused on Wi-Fi
  - Much *faster* (bit rates in excess of 600 Mb/s)
  - Extremely *popular* in home and office automation
  - *Integration* is possible with PCs, mobiles, and existing WLANs
- … but most of our approaches and results can be actually applied to *any* wireless transmission technology

Reliable communication

- In theory, networks used to interconnect devices in control systems must *always* deliver *all* packets *timely*
  - Unless some serious fault happens (remedy: fault-tolerance)
  - Wired solutions: CAN, Real-Time Ethernet, FlexRay, etc.
- What does *reliable* mean when applied to wireless networks?
  - The likelihood to *completely lose* a packet should be as low as possible
  - The likelihood to *miss* a packet *deadline* should be as low as possible
  - but the *ether* is open, erratic and much more prone to errors than wires
- No way: wireless networks are *unsuitable* for hard real-time
  - but, if *countermeasures* are taken, they offer interesting performance
  - *Soft real-time* systems (even demanding ones) are not out of reach

How to improve reliability?

- *Seamless redundancy*
  - Suitably counteracts *temporary* phenomena, like disturbance and interference from external networks and wireless stations
  - Offers (*on average*) tangible improvements
- *Tuning protocol parameters*
  - EDCA QoS can be exploited to *differentiate* traffic (RT vs. BE)
  - Operating parameters like *interframe spaces* and *contention windows* can be purposely employed to improve medium access
- *Centralized access schemes*
  - Prevent internal interference and delays typical of distributed access schemes when *all* wireless stations agree to obey the same access rules (S/W protocol *overlay* located above adapters)
  - Very interesting performance when coupled with other mechanisms

Seamless channel redundancy

- In its simplest embodiment, *Parallel Redundancy Protocol* (PRP) applied to conventional Wi-Fi equipment
  - Two *distinct* channels are required for each link (e.g., STA ↔ AP)
  - Frames copies are sent on *both* channels at the *same time*
  - Receivers *retain* the first copy and *discard* the latter
  - Packet *losses* decrease noticeably and transmission *latencies* become shorter and more predictable

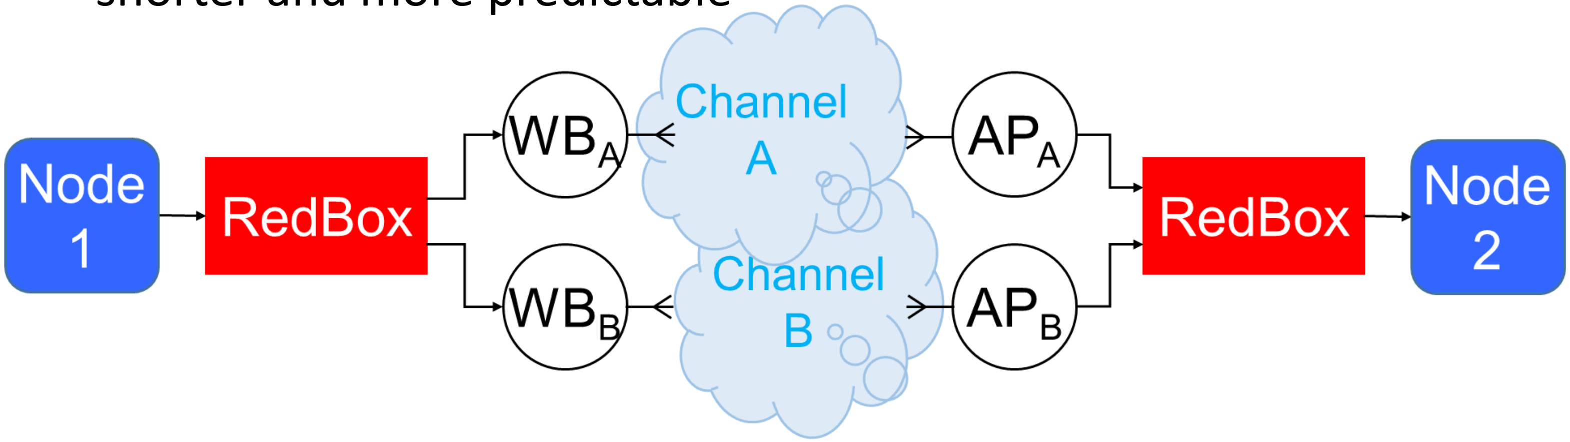

The next step

- Link-level seamless *Wi-Fi Redundancy* (Wi-Red)
  - Applies separately to the each *link* (unlike PRP, which is *end-to-end*)
  - Permits to exploit DL information (coming from *ACK* frames)
  - Reduces the *wasted* bandwidth and achieves higher *performance*
- Duplicate avoidance mechanisms
  - *Reactive Duplicate Avoidance* (RDA): Whenever an *ACK* is received on one channel, all pending transmissions of the same frame are *canceled* (in the transmission queue and possibly in the MAC)
  - *Proactive Duplicate Avoidance* (PDA): Save network bandwidth by temporarily *deferring* the transmission of the *second* copy of a frame (many solutions based on different heuristics)
- Performance analysis highlights *substantial* benefits

Wi-Red conceptual architecture

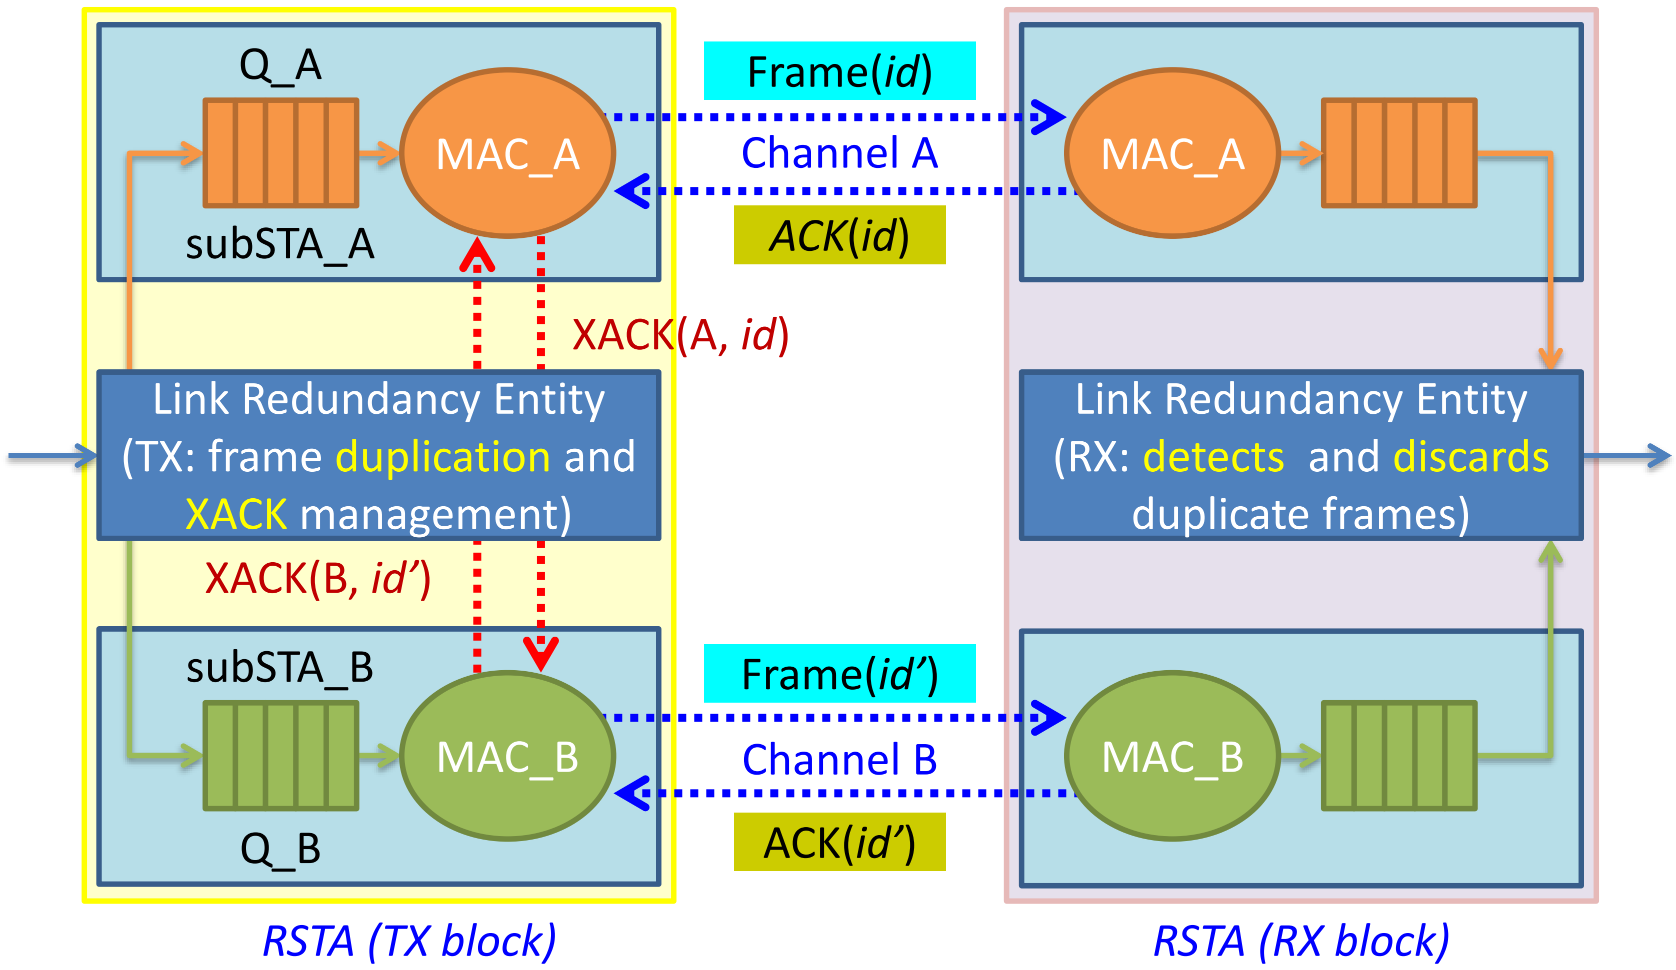

Experimental assessment

- *Simulation* shows that seamless redundancy is very effective in improving *communication quality* (and system reliability)
  - provided that behavior of channels is *statistically independent*
- Is this assumption true in the *real world*?
  - An *experimental* campaign was carried out on a *prototype* setup
  - A testbed was implemented on a PC running Linux with two Wi-Fi adapters, associated to two APs operating on *two distinct channels*
  - APs are connected to the PC through a switch
  - A *measurement task* sends a large number of packets (millions) using PRP rules and collects results about data delivery
  - Both *lab* and *industrial* environments were investigated
  - *Independence* of channels is typically verified quite well

Testbed

- Simple but extremely *effective* in order to determine transmission latencies and packet loss ratio
  - The PC acts as both *source* and *destination* of packets
  - Same *clock source* to measure timings

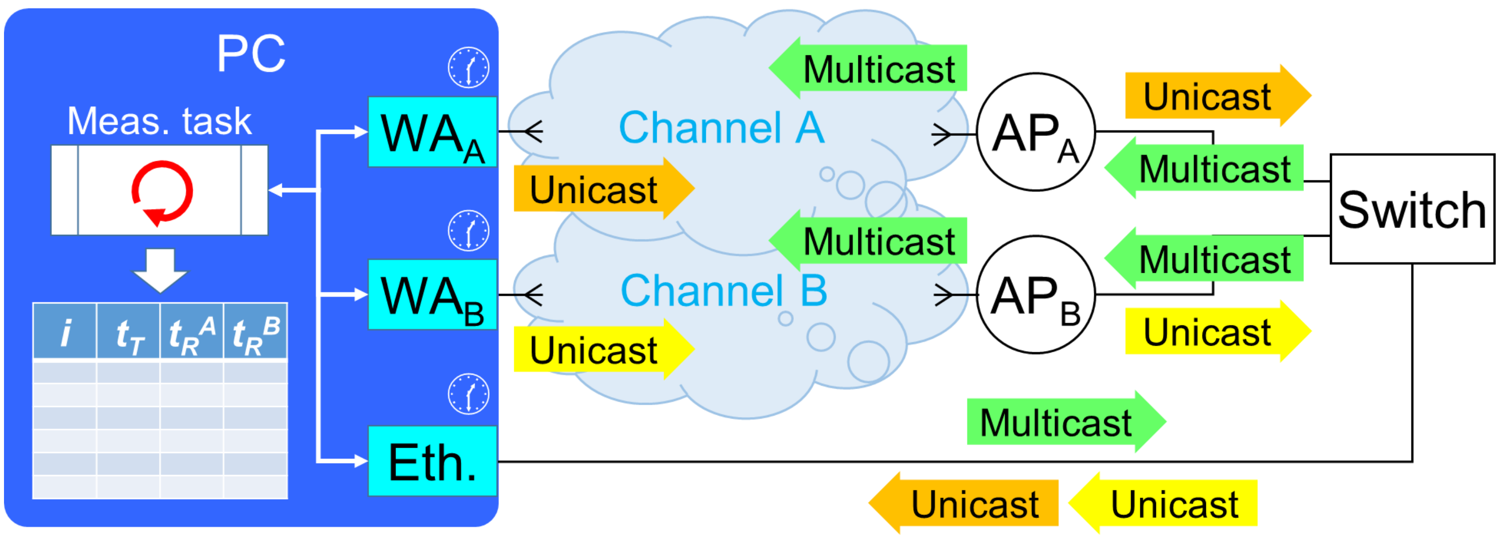

Guidelines

- In view of a *prototype implementation* preliminary guidelines were prepared for porting Wi-Red on real *embedded* devices:
  - *Adjacent Channel Interference* (ACI) between channels in a redundant link likely take place when antennas are located close to each other: *channels must be spaced as much as possible in the frequency range (e.g., by operating them in the 2.4 and 5 GHz bands)*
  - Activities causing *joint interference* on channels have to be prevented: *the network manager service must be rewritten from scratch so that network scans and reassociation for roaming are displaced in time*
  - Mechanisms related to the *delivery traffic indication message* (DTIM) in the access point cause unwanted latencies for multicast packets: *always enable Wi-Fi authentication and, possibly, encryption (overhead measured on modern adapters is really negligible)*

Sample results (guidelines followed)

Int.	Type	Ch.	$\bar{d}$	σ_d	$d_{p99.9}$	$d_{p99.99}$	$d_{p99.99}$	PLR [%]
No Interference	Unicast	1	1.34	1.54	15.90	25.70	55.17	0.0
		36	0.20	0.094	0.80	4.52	19.15	0.0
		1+36	0.20	0.085	0.77	3.90	16.77	0.0
	Multicast	1	2.13	2.50	26.56	43.35	117.98	10.58
		36	0.90	0.42	1.34	4.57	108.01	0.047
		1+36	0.90	0.42	1.28	4.41	108.01	0.0052
Interference	Unicast	1	1.43	1.91	20.86	34.48	67.93	0.0
		36	0.76	1.32	16.75	36.16	64.93	0.0
		1+36	0.49	0.42	4.05	7.04	14.77	0.0
	Multicast	1	2.07	2.47	25.82	40.38	77.95	9.41
		36	0.99	0.28	2.41	4.37	108.03	10.92
		1+36	1.06	0.89	12.80	26.68	108.03	1.03

Interference on channels 1 and 36 is very unbalanced – Latencies are expressed in ms

Centralized transmission scheduling

- The key for *deterministic* behavior in networks with shared transmission support is *coordinating* data exchanges
  - Intra-system *collisions* can be prevented completely
  - This does not apply to nodes belonging to *neighboring* networks…
- The simplest way to do so is using *centralized* approaches
  - Superframes in IEEE 802.15.4 (WSN)
  - PCF (and HCCA) in IEEE 802.11 (Wi-Fi)
  - Their behavior mostly resembles *cyclic executives*
- If messages are characterized by *firm* deadlines the best approach is using an *Earliest Deadline First* (EDF) scheduler
  - Transmission order is directly driven by *timing requirements* of data

Conceptual model for data exchanges

- *Time-critical data* exchanges are mapped on transmission services of the underlying network
  - A simple approach is to use a *three-way* packet exchange
  - If the source is the coordinator, request and data packets coincide

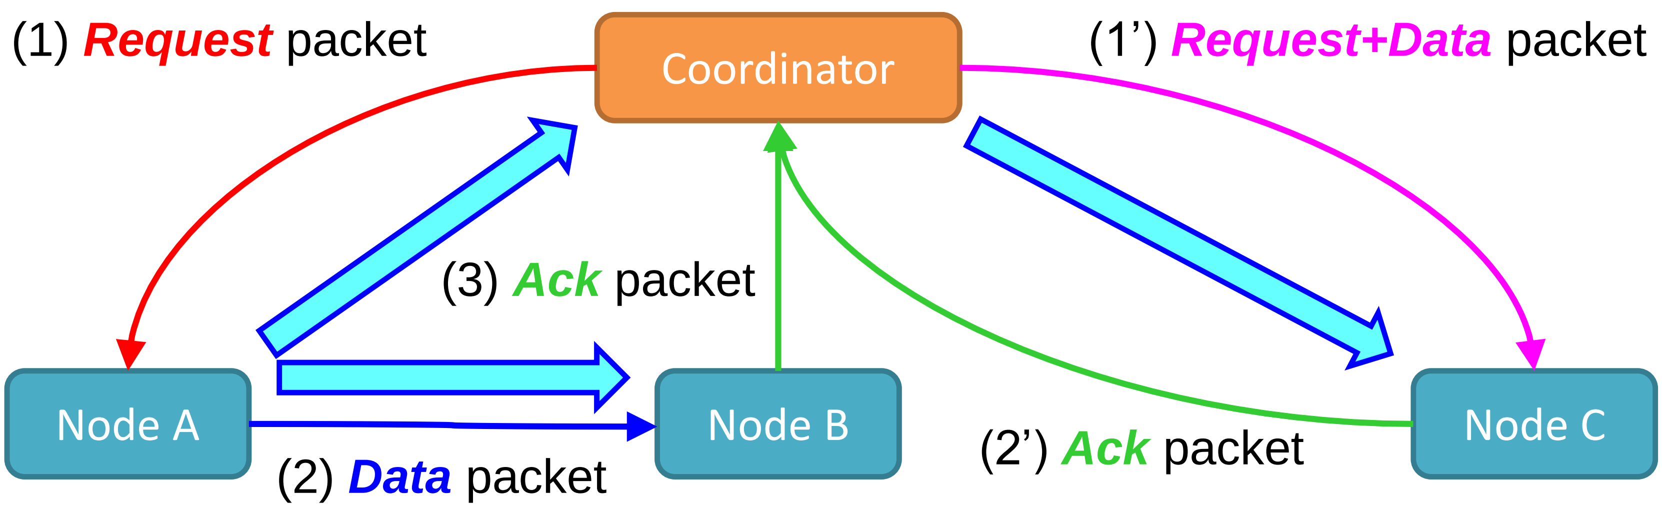

Reliability and timeliness

- *Retransmissions* are typically used to cope with frame losses
  - But they take (*variable*) time to be carried out
- In order to improve *reliability* not impairing *timeliness*
  - A certain amount of *planned retransmissions* can to be included directly in the stream set (by augmenting it suitably)
  - The coordinator has to *directly manage* both the first transmission attempts and retransmissions (as HCCA but according to EDF)
- Schedulability analysis is used to assess feasibility
  - Derivatives of Baker's test can be found in the literature
  - Both the relative deadline D_i and the number R_i of planned retries can be configured on a per-datum basis
  - We aim at obtaining a certain *deliver success probability* (DSP)

Further improvements

- The proposed scheme can be noticeably improved by
  - Disabling *random backoff*: collision avoidance is useless (and detrimental) because of the centralized access
  - Setting minimal *interframe gaps*: improves robustness against interference with external STAs and WLANs
  - Using *redundant channels*: decreases the packet error probability
  - Reusing the *unused bandwidth*: permits to accommodate additional retransmissions besides the planned ones without impairing feasibility
- A prototype setup has been implemented that combines all the above *mechanisms* and *tricks*
  - DSP experimentally evaluated under very *hostile* concurrent traffic
  - Results are good and quite close to *theoretical* expectations

Experimental evaluation

- Both simulation and experimental measurement
  - *Cumulative DSP* (evaluated on all streams) vs. *error probability (e)*
  - Bandwidth reuse provides *no guarantees* but is appealing *on average*

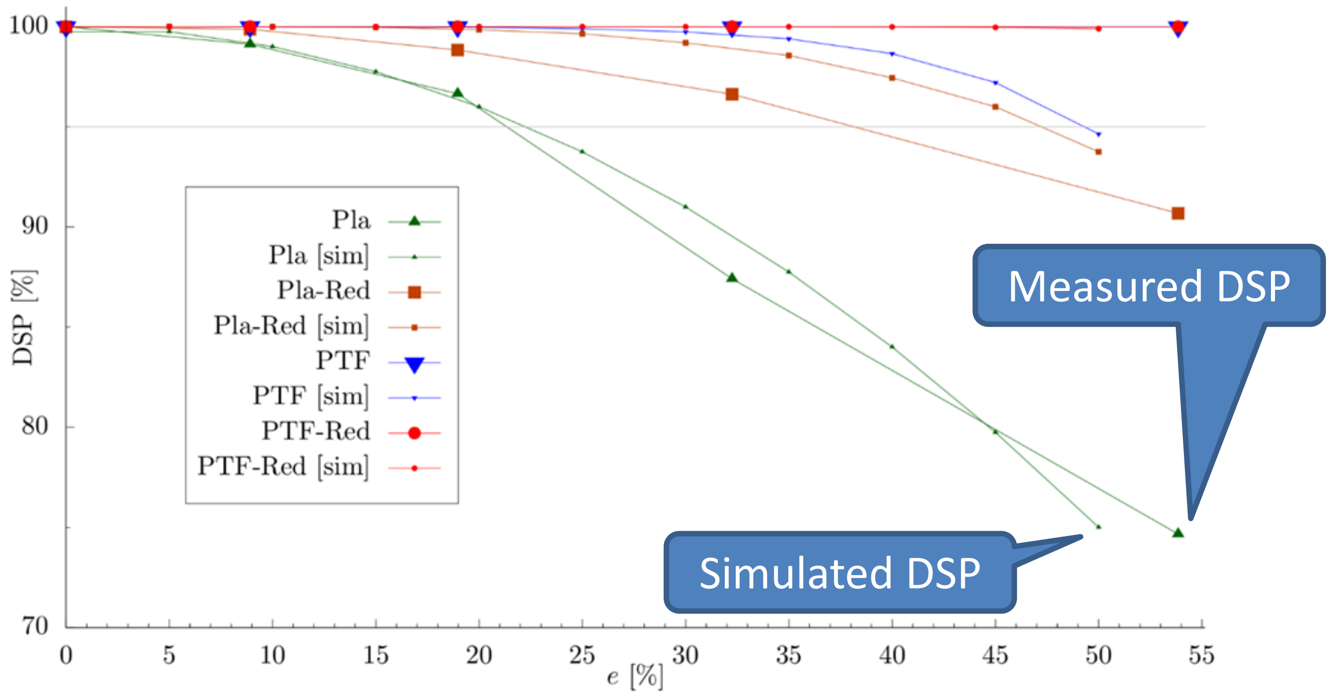

Distributed solutions

- Not as reliable as the centralized ones but interesting
  - Can be applied to *any* IEEE 802.11e (EDCA)-compliant H/W
  - Random backoff is *disabled*
  - Interframe spaces are *fixed* and depend on message priority
  - Kind of a *linear arbitration* (as opposed to binary)
  - *Feasibility* analysis "*à la CAN*" can be applied
  - Suitable for *event-driven* systems with sporadic transmissions
- Very important: *no modifications* are required to the MAC
  - Just minor changes to *drivers* are needed
  - Almost all existing commercial *adapters* can be used in theory
  - Planned for the inclusion in *prototypes* in our next research projects

Thanks for your attention

Any question?