Collaborative Project
Small-medium-scale focused research project (STREP)

Grant Agreement n. 224168 FP7-ICT-2007-2

WIDE

Decentralized Wireless Control of Large-Scale Systems

Starting date: 01 September 2008 Duration: 3 years

Deliverable number	**D2.3**
Title	**Towards ultra-reliable short-range wireless sensor networks**
Work package	WP2 – Wide-area WSN for distributed control
Due date	M22
Actual submission date	01/12/2008
Lead contractor for this deliverable	KTH
Author(s)	Mikael Johansson `<mikaeljl@ee.kth.se>` Pablo Soldati `<soldati@ee.kth.se>` Johannes Bleuel `<Johannes.Bleuel@e-senza.de>`
With the help of	Zhenhua Zou `<zhenhua.zou@ee.kth.se>` Haibo Zhang `<haibo@ee.kth.se>`
Revision	v1.0 (January 7, 2011)

Dissemination Level

PU	Public
PP	Restricted to other programme participants (including the Commission Services)
RE	Restricted to a group specified by the consortium (including the Commission Services)
→ **CO**	Confidential, only for members of the consortium (including the Commission Services)

Executive summary

This report focuses on mechanisms for ultra-reliable short-range wireless networks. Some important sources that causes unreliability of current wireless technologies are described, and mechansims and protocols for reliable wireless low-power networking are described. This includes the SenzaNET framework for reliable and industrial control and performance bound and optimal routing and scheduling policies to maximize system reliability in the *Wireless*HART standard.

Contents

1 Introduction **3**

2 Sources of unreliability **3**

2.1 Attenuation . . . 3

2.2 Fading . . . 4

2.3 Interference with other RF sources . . . 4

3 Mechanisms for ultra-reliable short-range WSNs **5**

3.1 SenzaNET: design and goals . . . 5

3.1.1 Ultra-low power . . . 6

3.1.2 Bounded communication latency . . . 6

3.1.3 True mesh networking . . . 6

3.1.4 Robust and secure transmission . . . 7

3.1.5 Comparison of SenzaNET with WirelessHART . . . 7

3.2 Improving the single link . . . 7

3.2.1 Node placement and antenna configuration . . . 8

3.2.2 Channel selection . . . 8

3.3 Scheduling and routing to maximize end-to-end reliability . . . 9

3.3.1 A Model for transmission scheduling in *Wireless*HART . . . 9

3.3.2 An abstract model for transmission scheduling . . . 11

3.3.3 Markov models for packet erasures on links . . . 11

3.3.4 Pre-scheduling of transmission opportunities . . . 12

3.3.5 Dynamic transmission scheduling . . . 13

3.4 Validation on real data . . . 16

4 Application to the Barcelona water distribution network **17**

1 Introduction

The aim of the work leading up to this deliverable has been, as stated in the Description of Work, to

> "Report isolating reliability bottlenecks of industrial wireless standards, proposals for alleviating these and, when applicable, arguments for why specific technologies are insufficient for the requirements of the water distribution network."

We have approached this problem in several ways. Since many years, it has been clear that existing low-power solutions such as Zigbee has severe reliability problems, partially because it is a single-channel solution and in partially due to restricted use of diversity in time and space. The *Wireless*HART standard was first released the year before this research started. *W*irelessHART allows to fully exploit diversity in time, space and frequency, but has the drawback that the standard relies on and provides functionality for centralized scheduling of network resources but does not provide any support or guidelines for constructing the actual schedule. Hence, it is non-trivial for a user of the technology to reap the full reliability benefits of the technology and bad scheduling could easily lead to networks that are less robust than the Zigbee technology that it was intended to replace. Hence, the final parts of this report describes the work that has been carried out within WIDE on fundamental limitations and optimal policies for transmission scheduling and routing in *Wireless*HART. This work has been mainly theoretical and been reported in numerous conferences and archival journals. Another drawback with *Wireless*HART is precisely the centralized management of network resources. Such a system can be fragile (since there is a single point of failure, *i.e.* the network manager) and typically results in very long network convergence time and large footprint implementation on nodes. As a complement, ESenza has been developing the SenzaNET technology. A lightweight distributed sensor network framework for industrial control. The protocol stack and its features are also described in this report. Admittedly, there is a number of alternative standards that are beginning to emerge. This includes ISA100, Bluetooth low-energy, and the IPv6/6loPAN/IETF ROLL/IEEE 802.15.4E-suite. However, since these standards have not yet had a large impact on the target application area of WIDE, we have not considered them in the work leading up to D2.3.

2 Sources of unreliability

Most technological approaches to wireless sensor networks today use the 2.4 GHz ISM-band for radio transmission, since this band is license-free globally. For the same reason, the focus of standardization efforts is on 2.4 GHz technologies, too. For the 2.4 GHz band specifically, reliability of wireless links is mostly determined by three factors:

1. Attenuation of the radio signal along its propagation path
2. Multi-path propagation effects causing self-interference of the radio wave (Fading)
3. Interference with other sources of electromagnetic waves

This chapter outlines the potential impact of these effects on the reliability of wireless communication links whereas the next chapter describes possible countermeasures and how they can be applied to achieve sufficient reliability; see also the extensive discussion in Deliverable D2.2.

2.1 Attenuation

Attenuation causes a gradual reduction of intensity of the transmitted radiowave, which is mostly due to the spreading of signal energy with distance. Atmospheric gases cause additional absorption, which is negligible in most cases. Although water generally absorbs 2.4 GHz waves, water vapor, fog or even rain do not significantly contribute to attenuation, due to the size of water drops being much smaller than the wavelength (12,5 cm). In contrast, water-containing obstacles with a size comparable to the wavelength or larger will attenuate heavily. The attenuation introduced by obstacles strongly depends on the materials used and on the pathlength of the radiowave penetrating the obstacle. The reliability of a wireless communication link against effects of attenuation is determined by the senders transmit power as well as receivers sensitivity, the link budget. Establishing a sufficient link budget while maintaining low power consumption of the communication hardware is one of the basic challenges in the field of wireless sensor networking. Typical values for common WSN hardware can be assumed as follows:

Transmit Power:	10 dBm
Receiver Sensitivity:	-98 dBm
Resulting link budget:	108 dBm

Typical Attenuation values are

Free space attenuation:	80 dB for a 100 m link (increases by 6dB when distance doubled)
Plasterboard office wall:	3 dB
Office window:	3 dB
Brick wall:	6 dB

2.2 Fading

Fading occurs when a an electromagnetic wave propagates from sender to receiver over multiple paths and interferes with itself at the receiver. In the simplest case of two wave propagating, complete suppression of the signal occurs when the difference in distance traveled is half the wavelength of the wave, so that the waves maximum and its minimum overlap. In the real world, the situation is much more complex: Radio waves will propagate in several directions and they will be reflected and attenuated by the objects around. The extent of reflection and attenuation depends on size, shape and material of the objects making up the environment. In industrial environments, the presence of metallic objects is of specific importance for the propagation of radio waves, since they are causing reflections. Reflection at metallic objects is almost without attenuation, meaning that the energy of the wave is preserved. While such reflection could be beneficial for the range over which a radio link can be operated, multi-path fading effects tend to be stronger in metallic environments which makes the link more unreliable. To avoid communication losses from fading, wireless systems use multiple RF channels at slightly different frequencies / wavelengths. When a specific channel suffers heavily from fading, an adjacent channel typically does not, since the difference in wavelength changes the conditions for self-interference. Alternatively, diversity of propagation path by using multiple antennas can resolve fading challenges.

2.3 Interference with other RF sources

Due to the heavy usage of the 2.4 GHz band, many sources of interference exist. Wireless-LAN (IEEE802.11b/C and g), Bluetooth and analog video transmitters are the most common sources.

The IEEE802.15.4 radio standard used as communication medium for many WSN implementations and also throughout the WIDE project is characterized by high robustness against interference from other radio sources. This is achieved by spreading the radio signal over the full bandwidth of 5 MHz in each channel, which is much larger than the actual bandwidth of the payload signal. At the receiver, an incoming radio signal is first correlated with the IEEE802.15.4-specific spreading code. Signals

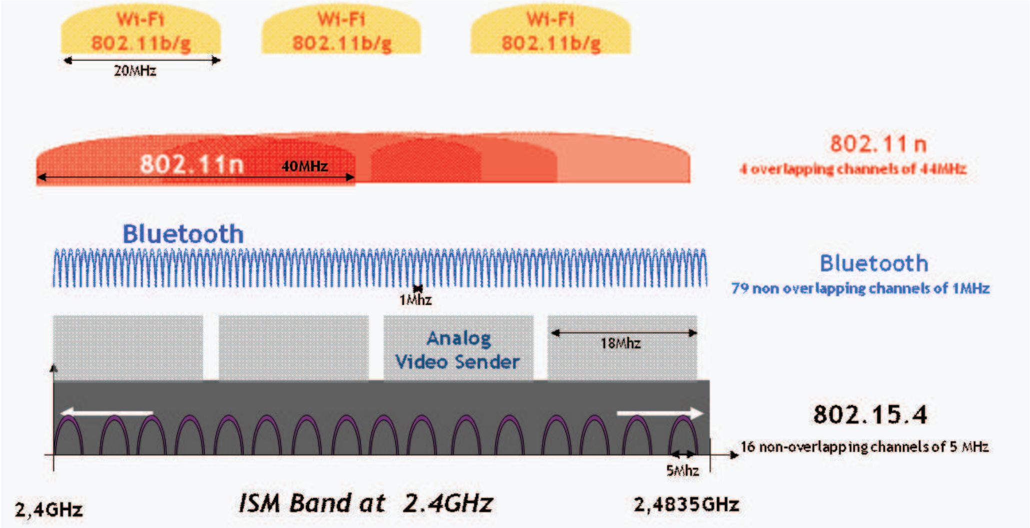

Figure 1: Co-existing wireless technologies in the 2.4GHz ISM band.

from other sources, e.g. WLAN, use different spreading codes, so that they are suppressed at the receiver directly. Being designed for low data rates, enough additional bandwidth is available to also introduce redundancy and thus allowing for higher noise levels while maintaining low packet error rates.

3 Mechanisms for ultra-reliable short-range WSNs

Designing ultra-reliable short-range wireless sensor networks implies compensating for failures which can occur on individual radio links and proper mechanisms for establishing and managing the networking of nodes and the end-to-end transmission of data.

It also has to be kept in mind, that “reliability” must be defined according to the needs of the application. In case of WIDE, where control applications are targeted, it is not enough to have high reliability but data packets must also be transmitted with a certain maximum latency. If a transmission takes longer, the data packet is considered lost and will be dropped.

This chapter describes two complementary activities within the WIDE project. On the one hand, the SenzaNET technology has been developed under WIDE as a practical plug-and-play framework for wireless communication in industrial environments. On the other hand, theoretical work has been done to investigate the performance limitations of *Wireless*HART networks and to develop optimal routing and scheduling policies.

3.1 SenzaNET: design and goals

The SenzaNET technology was developed under WIDE as a plug-and-play framework for the wireless communication of universal devices in automation environments. This framework allows the integration of sensors, actuators and instrumentation in a low power Wireless Personal Area Network (WPAN). SenzaNET is built on top of the IEEE802.15.4 radio standard. Compared to WirelessHART, the end-to-end latency and the effective information rate in a dynamic environment is improved, while the high reliability levels required by industrial applications are maintained. This is achieved by a decentralized networking protocol, which does not require communication to a centralized Network-Manager. Instead, the wireless network is organized into branches. For each branch, a selected

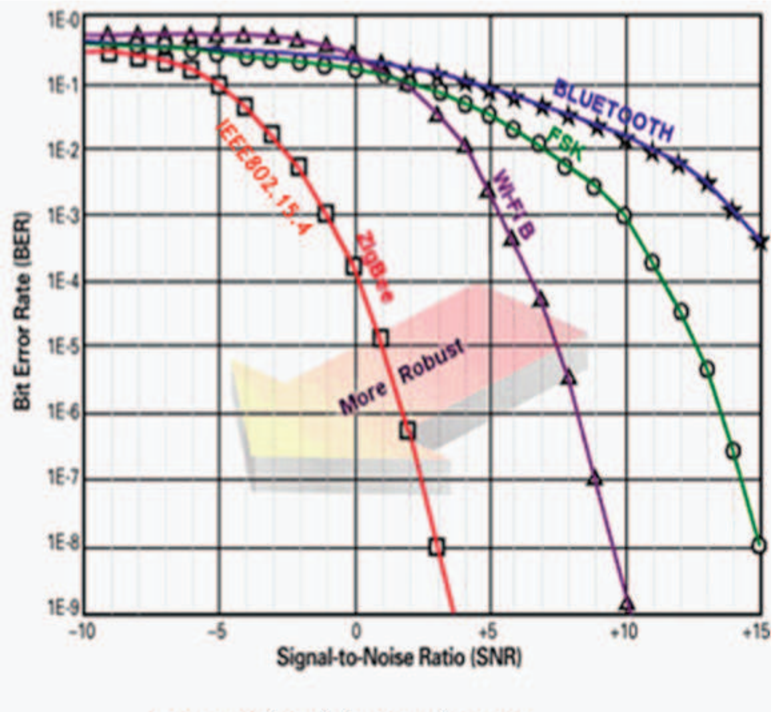

Source: Embedded-computing.com

Figure 2: Bit error rates for co-existing wireless technologies.

wireless node handles all tasks of assigning communication slots, so that overall network traffic is reduced and re-arrangement of routing paths and communication slots can be much faster.

3.1.1 Ultra-low power

Extended operation on battery power or the use of energy scavenging is a key requirement in the deployment of wireless device networks. SenzaNET utilizes an accurate time synchronization algorithm, allowing all SenzaNET nodes to remain in standby mode when not required to perform a measurement or wireless transaction. In standby mode, nodes can operate on extremely low amounts of energy, since most hardware components are powered off. As a result, overall power consumption is dramatically reduced and largely correlated with the desired sample rate, rather than unnecessarily drained by idle states.

3.1.2 Bounded communication latency

A secondary benefit of the time synchronization approach is the ability to provide balanced medium access and predictable transmission slots. This keeps latency within tolerable limits, enabling real-time monitoring of assets and guaranteed delivery of time-critical information such as alarms and control commands.

3.1.3 True mesh networking

The self-organizing and self-healing properties of SenzaNET provide maximum fault tolerance and deployment flexibility. SenzaNET nodes establish connections and transmission paths by themselves, and are capable of multi-hop routing for formation of arbitrary topologies and bridging of extended distances. In contrast to ZigBee, SenzaNET routing nodes do not depend on mains power and can run on batteries. This makes SenzaNET particularly suitable for environments where powering individual nodes designated to act as routers would be difficult or prohibitively expensive.

3.1.4 Robust and secure transmission

For maximum reliability, SenzaNET employs automatic retries, acknowledgements, and a channel hopping scheme. Network security is provided through encryption of all data transmissions, and each individual data packet is integrity protected. In addition, join requests by new nodes can be authenticated via access control list so that only known and legitimate nodes are granted access, based on their unique MAC address.

3.1.5 Comparison of SenzaNET with WirelessHART

When comparing SenzaNET and the WirelessHART standard, the obvious difference is the decentralized design of SenzaNET: In WirelessHART, all networking intelligence is centralized and a so-called NetworkManager determines all routes in the network and also re-organizes the network based on updated information of the RF signal-strength between each node-pair in the network. In SenzaNET, all devices are able to operate independently from a central NetworkManager using local intelligence. The devices can determine locally when the connection to their parent is assumed bad, when to disconnect from it and when and how to reconnect to a new parent node. Centrally, the server application SenzaWMS is available for user interaction with the wireless sensor network, to authenticate a device in the network, to store data, view status of each node and to send new configuration parameters. The main design motivation of SenzaNET is to provide a network which is simple to use, simple to configure and allows easy integration with a variety of existing automation systems. In contrast to comprehensively standardized protocols like WirelessHART, SenzaNET does not require extensive infrastructure for device management, but offers a set of common interfaces which are supported by a almost all automation systems. For field devices, digital and analog interfacing is offered, whereas on the side of PLC controllers, mostly serial interfacing (RS232), TCP/IP and Profibus are being used. Interfacing to SCADA-systems is achieved through Modbus-TCP.

Devices in WirelessHART are supplied with a large set of pre-defined commands for configuration and parametrization and no networking intelligence. Devices in WirelessHART are supposed to listen to commands from the NetworkManager when it somes to setup communication links. This makes the NetworkManager very complex and its operation very tricky. All decisions on network topology and re-formation have to be taken centrally, increasing networking latency, thus limiting network flexibility substantially. Especially in case of relocation of devices, network latency becomes unacceptably high due to the complex network communication required between NetworkManager and all devices involved in re-routing.

To summarize, SenzaNET compares with the WirelessHART and ZigBee standards as follows:

3.2 Improving the single link

It is natural to first attempt to improve the reliability of the single link. This can be done in a variety of ways, including clever node placement and antenna configuration, mitigation of external interference, power control, channel selection, etc; This section focuses on design rules for node placement, antenna configuration and channel selection. The final part of the report considers system-wide optimization of reliability.

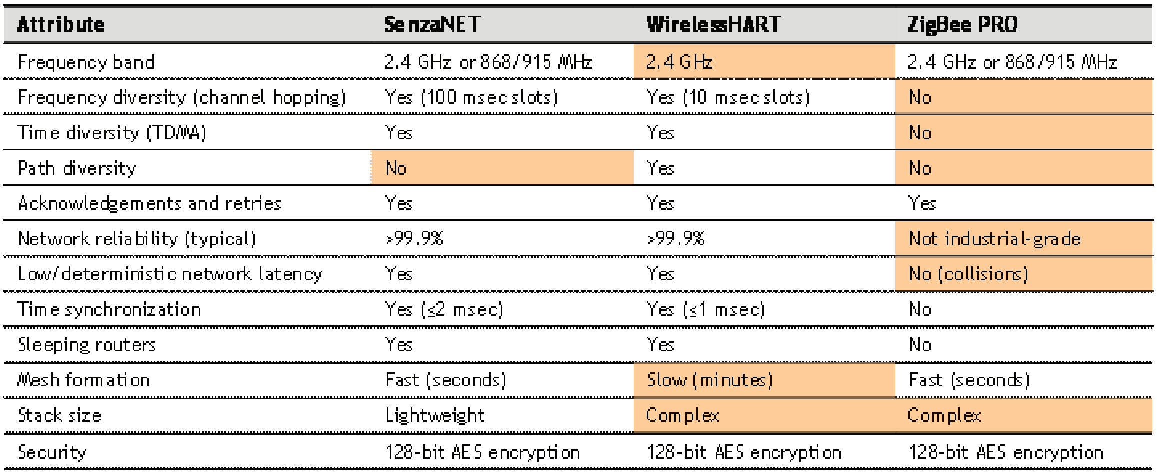

Attribute	SenzaNET	WirelessHART	ZigBee PRO
Frequency band	2.4 GHz or 868/915 MHz	2.4 GHz	2.4 GHz or 868/915 MHz
Frequency diversity (channel hopping)	Yes (100 msec slots)	Yes (10 msec slots)	No
Time diversity (TDMA)	Yes	Yes	No
Path diversity	No	Yes	No
Acknowledgements and retries	Yes	Yes	Yes
Network reliability (typical)	>99.9%	>99.9%	Not industrial-grade
Low/deterministic network latency	Yes	Yes	No (collisions)
Time synchronization	Yes (≤2 msec)	Yes (≤1 msec)	No
Sleeping routers	Yes	Yes	No
Mesh formation	Fast (seconds)	Slow (minutes)	Fast (seconds)
Stack size	Lightweight	Complex	Complex
Security	128-bit AES encryption	128-bit AES encryption	128-bit AES encryption

Figure 3: Compaisong between *Wireless*HART, SenzaNET and Zigbee.

3.2.1 Node placement and antenna configuration

Based on field testing results, engineering guidelines were identified for designing wireless sensor networks for field applications:

- Sensor nodes should not be placed next to WLAN and Bluetooth radio sources, a minimum distance of 2 m should be respected. This to avoid overmodulation of the receiver
- Minimum 3 neighboring nodes should be visible to each device, so that enough redundant wireless links can be established
- Devices (antenna) should be mounted ¿0.5m from any vertical surface and ¿1.5m off the ground
- When neighboring nodes are vertically separated from each other, the direct connecting line should have a pitch of less than 45 . This holds for omnidirectional antennas, since the power radiated varies with the vertical angle. Tilting the Antenna accordingly would be an option in cases where a pitch less than 45 is not possible.

Omnidirectional dipole antennas generally showed the best performance during field tests, directional and high-antennas are suitable under Line-of-sight conditions only. Small-scale PCB-antennas (inverted-F) are benefitial due to their small size and very cost-efficient design, but range might be lower by a factor of 2-3 compared to dipole antennas.

3.2.2 Channel selection

Selecting suitable RF channels requires a trade-off: The more channels are being used, the more options for avoiding fading effects and interference with other RF sources are available. Using too many channels, however, makes it more difficult for network nodes to join the network or to re-connect, since they have to search more channels to find the network. We found a channel hopping sequence using 3 channels to represent a good compromise. Channels 11, 19 and 26 of the 2.4 GHz band are recommended, which means the two ends of the available spectrum plus a channel in the middle are used. This way, maximum tolerance against fading can be achieved, other sources of interference can largely be avoided because they will in most cases affect only one of the three channels chosen and an energy-efficient yet reasonable fast network search is still possible.

3.3 Scheduling and routing to maximize end-to-end reliability

Even with clever node placement, advanced antennas and careful network planning (in terms of channel selection and mitigation of external interference), it is very likely that individual links will still be unreliable. There are many reasons for this: node placement is often constrained by the physical layout of the process to be monitored and it might be impossible to avoid significant shadowing; there is often no single channel that provides reliable communication across a complete site; the presence of mobile elements such as cranes and fork lifts introduces a significant risk for blockage of otherwise strong links; etc. Improved reliability can then be engineered on the system level, exploring channel diversity in time, space and frequency. Within the WIDE project, we have developed techniques for optimal transmission scheduling in systems that combine multi-channel TDMA with multi-path routing. Representative technologies include the *Wireless*HART and ISA100 standards, but the results also have bearing for upcoming proposals such as IETF ROLL and new IEEE 802.15.4 extensions.

3.3.1 A Model for transmission scheduling in *Wireless*HART

*Wireless*HART is an extension of the wired HART protocol for process and control applications. A *Wireless*HART network is composed of *field devices* connected to the process equipment, *gateways* that enable communication between host applications and field devices within the network, and a *network manager* that is responsible for network health monitoring, network configuration, maintaining routing tables and scheduling communication between devices. The network may also include adapters for connecting to existing HART-compatible devices and handhelds to configure, maintain and control plant assets. The network is a full mesh network, in the sense that all field devices are able to source, sink and forward packets on behalf of other devices in the network.

The routing layer in *Wireless*HART supports multi-path multi-hop routing. A distinct feature of the current version of the standard is that all traffic has to go through the gateway. In other words, traffic from one field device to another will always be forwarded from the first device to the gateway, and then from the gateway out in the network to the second device. The medium access control layer is based on multi-channel TDMA and performs channel hopping at each slot boundary. *Wireless*HART has several mechanisms for promoting network-wide clock synchronization within 1ms accuracy and uses time slots of 10ms length. One time slot allows for channel switching and the transmission of a single packet and the associated acknowledgement. Transmission opporunities can be dedicated or shared. We focus on dedicated time-slots where only a single transmitter-receiver pair can communicate at any given channel. *Wireless*HART operates over low-power radios compliant with the IEEE 802.15.4-2006 standard, which supports 16 channels in the 2.4GHz ISM band. Channel blacklisting is employed to avoid channels with consistently high interference levels (e.g. due to coexistence with 802.11). In practice, some channels might also be blacklisted to protect wireless services that share the ISM band with *Wireless*HART. Hence, the number of channels is limited and efficient channel utilization is instrumental. To logically structure the global transmission schedule, *Wireless*HART supports multiple superframes. As illustrated in Figure 4, each such superframe is typically used for scheduling one networking operation such as the collection of measurements from a subset of sensors, or the dissemination of commands to a set of actuators. These superframes are then merged into one, and unless scheduling conflicts between different superframes are resolved during the global scheduling phase, the standard prescribes how nodes should behave to resolve such situations. A thorough description of the *Wireless*HART standard can be found in the book [1].

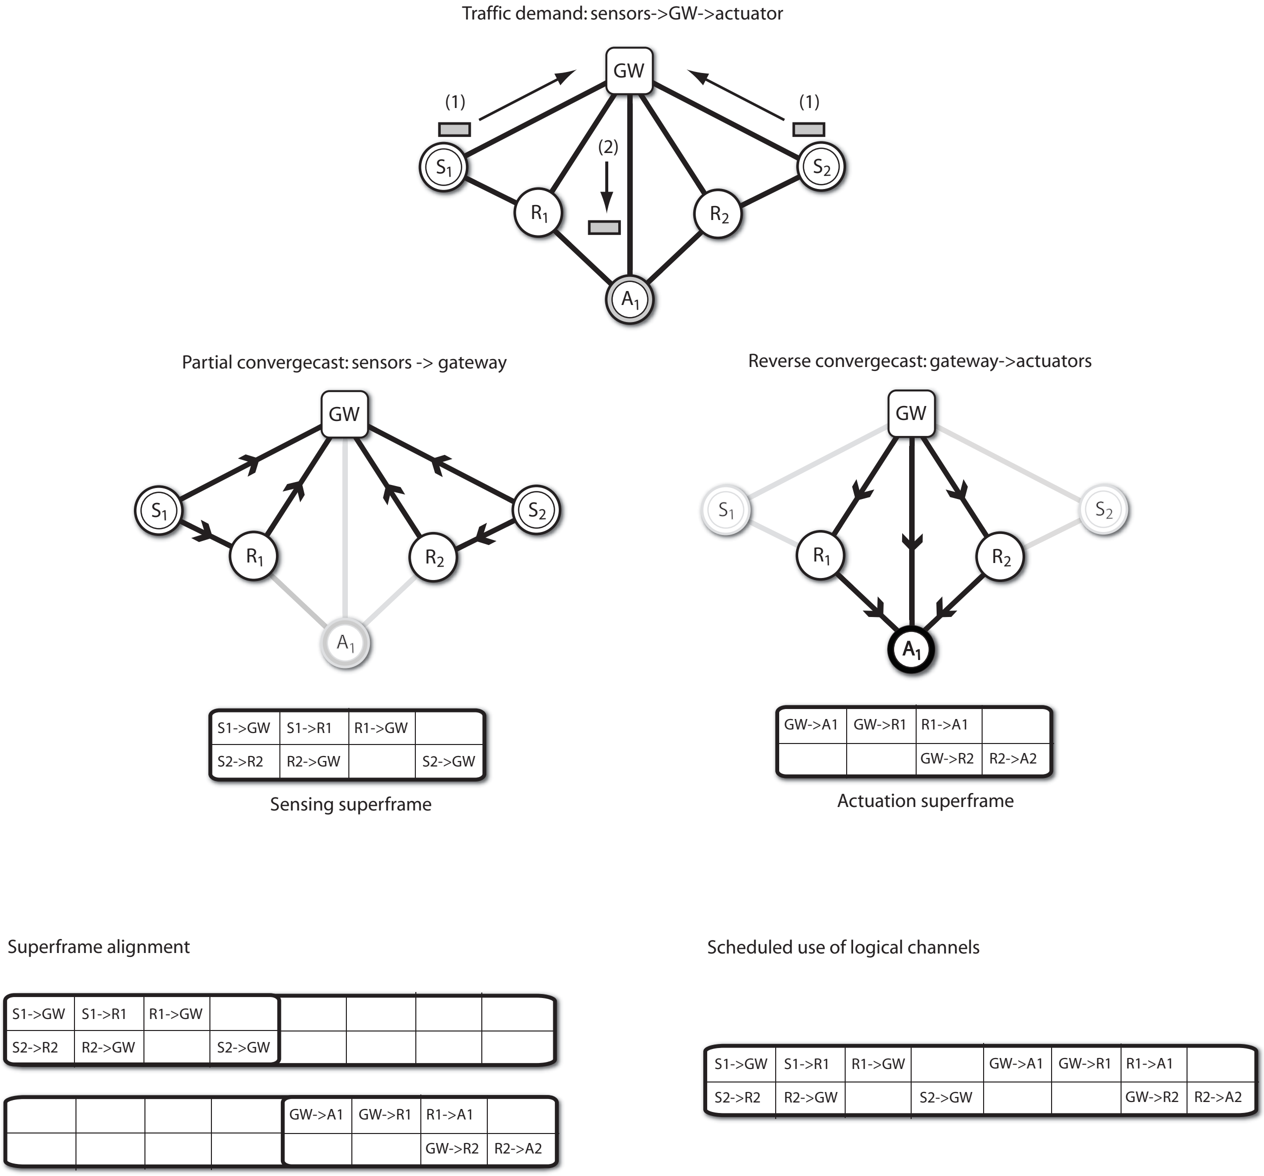
Traffic demand: sensors->GW->actuator

Partial convergecast: sensors -> gateway

Reverse convergecast: gateway->actuators

Sensing superframe

Actuation superframe

Superframe alignment

Scheduled use of logical channels

Figure 4: Illustration of our design methodology. In the first step, the traffic demand is broken down into basic networking operations, such as unicast, broadcast, convergecast, etc. In this case, the network needs first to collect data from the two sensors using a convergecast operation, compute the control action (not conisdered here) and then disseminate the command to the actuators using a reverse convergecast. In the second step, the basic operations are then scheduled on different superframes. Finally, in the third step, the superframes are ordered logically in time and adjusted so that the composite schedule does not have any conflicts and uses a minimum of logical channels.These logical channels are mapped onto physical channels using the channel hopping mechanisms in the *Wireless*HART standard.

3.3.2 An abstract model for transmission scheduling

We will now develop a mathematical model for analysis and design of transmission schedules in *Wireless*HART based on the short description above. The model considers a single superframe in isolation – the problem of merging multiple superframes into one conflict-free superframe is not considereded in this report.

We assume that the devices are static and represent the network topology with a graph $G = (V, E)$ where vertices $v \in V = \{v_0, v_1, \ldots, v_N\}$ represent devices and edges in E denote device pairs that can communicate with each other.

We focus on communication from devices to the gateway and assume that all such traffic is routed along a spanning tree $T = (V, E')$ with $E' \subseteq E$ rooted at GW. We refer to T as the routing topology of the network to stress that it could potentially be very different from the actual physical placement of the devices. For every device v_i we let f_i denote its parent and ch_i be the set of its children in T.

Time is synchronized and slotted with standardized length that allows the transmission of one data packet and its associated acknowledgement. Transmission opportunities in a given time slot on a given channel are dedicated and can only be assigned to one device pair. Devices are equipped with a half-duplex radio trasceiver, which means that devices cannot transmit and receive in the same time slot. There maximum number of concurrent transmissions is equal to the number of channels C available for communication.

3.3.3 Markov models for packet erasures on links

There is a large body of literature of measurements and models for wireless channels in different scenarios, including indoor and outdoor environments with fixed or mobile transmitters. However, the majority of models consider the time-varying behavior of the wireless channel and not the resulting packet-level performance, and few models and public data traces are available from industrial deployments in the 2.4GHz ISM band relevant to our work (see, *e.g.* D2.1). A notable exception is the work by Willig *et al.*, who performed measurements and modeling work on industrial 802.11 communications [2].

In the development of higher layer mechanisms, it is customary to use simple stochastic models of erasure events on links. The simplest such model is the Bernoulli model, where packet losses on a given link are independent over time and occur with a fixed probability in each time slot. A problem with this model is that it fails to capture link burstiness, *i.e.* that link losses tend to occur in relatively long sequences, possibly followed by rather long periods where practically all transmissions are successful. The simplest stochastic model that captures link burstiness is the two-state Markov model due to Gilbert and Elliot [3, 4], shown in Figure 5. This model has two states, G (for good) and B (for bad), corresponding to successful transmission and packet loss, respectively. The Markov chain evolves in discrete time and state changes coincide with the time slot boundaries. Hence, the probability of successful packet transmission at time slot t given that the Markov chain was in good state during time slot $t-1$ equals q_G, and the conditional probability of successful transmission at time t given that the Markov chain was in bad state during time slot $t-1$ is q_B. The average packet

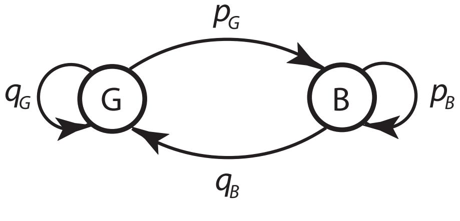

Figure 5: Two-state Markov chain model for packet erasures over a link.

loss probability is

$$\pi_B = \frac{1 - q_G}{1 - q_G + q_B}, \tag{1}$$

and the unconditional probability of successful transmission is $\pi_G = 1 - \pi_B$. Let T_G and T_B be the average sojourn time (in number of time slots) for the good and bad state, respectively *i.e.*

$$T_G = \frac{1}{1 - q_G}, \qquad T_B = \frac{1}{q_B}. \tag{2}$$

The two-state Markov chain model reduces to the Bernoulli model of independent packet erasures when $p_B = p_G = p$ and $q_B = q_G = q = (1 - p)$. In our theoretical work, we will assume that the Markov chains for different link are independent of each other.

Despite its simplicity, the GE model has proven to be reasonably accurate in capturing real packet loss behavior and the parameters q_G and q_B can be readily estimated from loss traces, see *e.g.* [2].

We estimate average error burst length T_B and average loss probability π_B from the data trace collected at KTH. The data trace includes 6×10^4 transmission results of a single link. q_G and q_B can be computed by plugging in Eq. (1) and Eq. (2). Figure 6 shows the error burst length probability predicted by two loss models, whose parameters are estimated from the real data trace. GE model can better capture the error burstiness of the real data trace since it also has a slower decay rate.

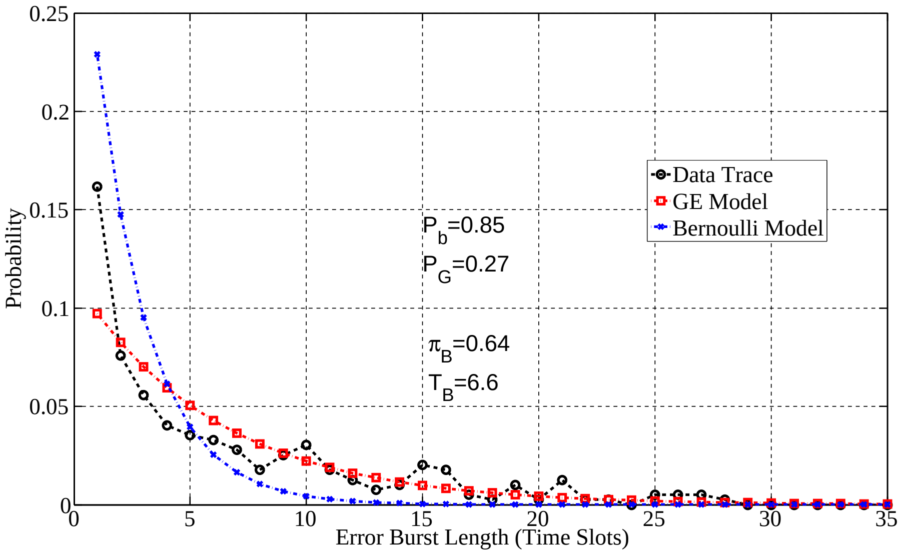

Figure 6: Comparison of two channel models and the real data trace.

3.3.4 Pre-scheduling of transmission opportunities

One natural and widely used approach for dealing with packet erasures is to use retransmissions: if a transmitter does not receive an acknowledgement from the intended receiver within a given time interval, it assumes that the transmission failed and schedules a new tranamission attempt. This basic “automatic repeat request” mechanism exists in many variations in practically all wireless systems. However, when we perform centralized scheduling of transmission attempts we have to allocate such retransmission opportunities on beforehand and activate them only when required.

To make things simple, consider the (uni-cast) transmission of a single packet along a sequence of $N + 1$ nodes. The time-optimal way to schedule the transmission is simply to activate the links sequentially from source to destination. If all transmissions are successful, the packet reaches the destination in N time slots. Assume now that a packet transmission on link l fails with probability p_l. Then, the unicast transmission scheduled in this way fails with probability

$$1 - \prod_l (1 - p_l).$$

In this model, reliability can only be enhanced by allowing retransmissions, and hence letting the end-to-end transmission take longer time. As a simple example, if we allow two retransmission attempts to each link, the unicast transmission will take at least $2N-1$ retransmission attempts (since the final link is only scheduled at time slots $2N-1$ and N) and will fail with probability $1-\prod_l(1-p_l^2)$. However, this solution is neither optimal in terms of maximizing reliability under a deadline constraint that the transmission must be completed in $2N$ time slots, nor does it extend to any deadline which is not a multiple of N. Such extensions will be given next.

We consider a unicast retransmission flow in which a single copy of the packet is routed from the source to the GW via a line routing topology. Associated with each packet is a strict latency bound T (time slots) for packet delivery from source to destination. We are interested in the following pre-scheduling scheme: T time slots are statically allocated to nodes. Each node is assigned to x_n dedicated time slots, with $\sum_n x_n = T$. Our objective is to develop optimal scheduling policies aiming at maximizing the probability that the packet is delivered within the deadline T. We refer to this problem as the *deadline-constrained unicast scheduling*, which can be formulated as follows:

$$\begin{array}{lll} \underset{\mathbf{x}}{\text{maximize}} & \sum_n \log q(\mathbf{x}) = \sum_n \log(1 - p_n^{x_n}) & \\ \text{subject to} & \sum_n x_n = T & \\ & x_n \in \{1, 2, \ldots, T\} & \forall\, n, \end{array} \tag{3}$$

The above formulated problem can be solved by the following greedy algorithm.

Algorithm 1 Greedy algorithm.

```
Initialize $\mathbf{x}^{(0)} = [0 \ldots 0]$.
For $t = 1 \ldots T$:
    1. Compute $\Delta = \log(q(\mathbf{x}+1)) - \log(q(\mathbf{x}))$
    2. Let $n^\star = \arg \max_{n=1,\ldots,N} \Delta$
    3. Set $x_{n^\star} = x_{n^\star} + 1$
```

Theorem 1 *The greedy algorithm 1 yields the optimal time slot allocation* $\mathbf{x}$ *that maximizes the reliability of pre-scheduling of a unicast flow along a single-path.*

3.3.5 Dynamic transmission scheduling

In pre-scheduling scheme, all time slots are dedicated slots which are statically allocated to dedicated nodes. For any node allocated with more than one time slots, if the transmission attempted in the first slot is successful, the remaining allocated slots will be wasted as they can not be shared with other nodes. Although pre-scheduling is energy-efficient as the nodes can turn off radios during non-scheduled time slots, it is not optimal in terms of maximizing end-to-end reliability for deadline-constrained traffic. To address this problem, we investigate dynamic transmission scheduling scheme in which time slots are virtually shared among non-conflicting nodes, and packets are dynamically routed on the basis of the time left to meet the deadline without prior knowledge of which links have failed on the network.

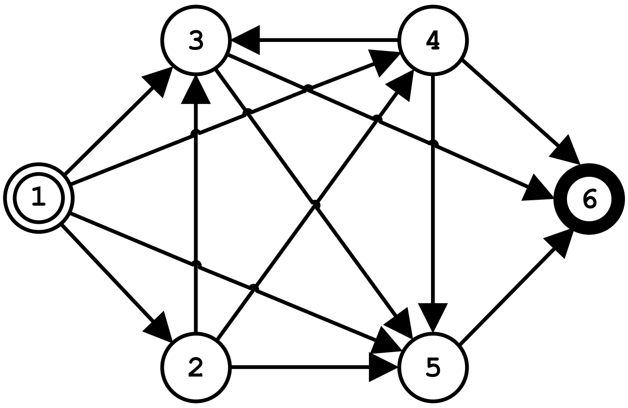

Figure 7: Example DODAG topology

We consider a scenario where a single packet, generated by an arbitrary node at time $t = 0$, should be transmitted over a destination oriented directed acyclic graph (DODAG) to the sink within a deadline of D time slots. Our aim is to find the joint routing and scheduling policy that maximizes the probability that the packet reaches its destination within the deadline. Packet losses are not only considered independent, but also are correlated in time. The erasure events on links follow an underlying two-state Markov chain model presented in Section 3.3.3. The optimal policy also hinges on what knowledge we can assume about the state of the underlying Markov chain at the time when we make the scheduling decision. We investigate the deadline-constrained maximum reliability routing problem under three information patterns: (a) nodes can know the transmission results ahead of time; (b) nodes have full state knowledge at time slot $t - 1$; (c) nodes do not know the state of their links until they are used.

A suite of such problems are solved using dynamic programming. We first derive the optimal one-step next-hop decision under different loss model and information pattern. We define $R_i(d)$ as the the probability that node i will be able to deliver the packet to the sink node within maximum d time slots. Clearly, the local optimal decision given the maximum delay constraint d can be computed using the estimated link state and the success delivery reliability of maximum delay constraint $d - 1$ of itself and its immediate neighbors. Specifically, the optimal next hop under information pattern (b) and Bernoulli loss model is

$$j_d^\star = \arg\max_{j \in \mathcal{N}_i} \big(q(i,j) R_j(d-1) + p(i,j) R_i(d-1) \big), \tag{4}$$

in which j is the immediate neighbors of node i, $q(i,j)$ and $p(i,j)$ are the reliability and loss probability of the links respectively. For other loss models and information patterns, we should apply different optimal local routing decisions. In GE loss model, the states (Good or Bad) of all outing links have to be taken into account. In information pattern (a), a simpler routing decision is possible since the node can preview the transmission result.

In all losses models and information patterns, the initial values are the same where the reliability is one at destination of any maximum delay constraint, and the success delivery probability with maximum delay constraint 0 is zero at all other nodes.

Then, denoted by dynamic programming principle, letting each node forward packets according to their optimal local forwarding decisions at each maximum delay constraint maximizes the probability that the packet generated by the source node arrives at the sink node within the deadline D.

We analyze the end-to-end reliability for the different loss models and information patterns by considering the topology in Figure 7, where a source (node 1) sends packets to a sink (node 6). Figure 8 considers homogeneous links with the same (fixed) average behavior in good and bad state, *i.e.* $\pi_G = \pi_B = 0.5$. The average burst length in good and bad state, *i.e.* T_G and T_B respectively, in

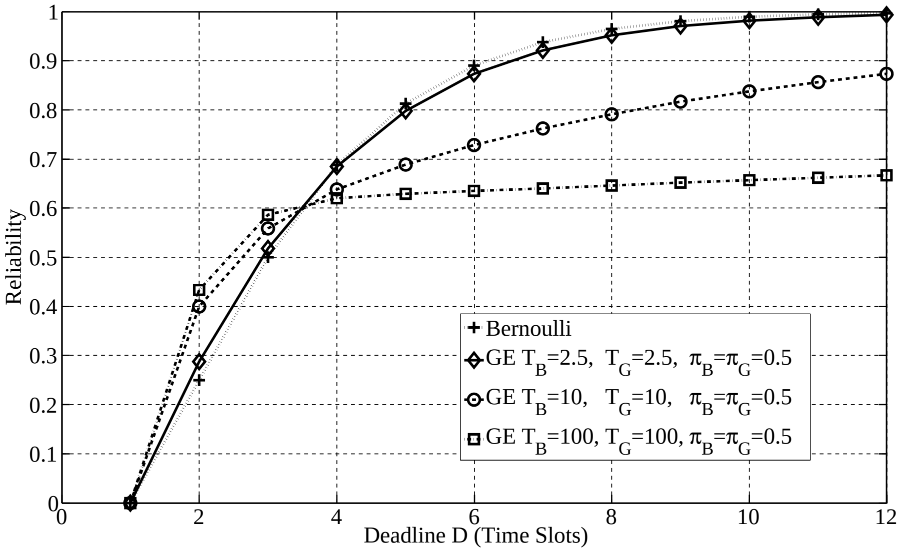

Figure 8: Deadline-Reliability curves for information pattern (b) with different T_B and T_G in GE model.

equation (2) are changed symmetrically to maintain $\pi_G = \pi_B$. We first notice that the larger the burst length, the better is the reliability for very short deadlines. This follows from Eq. (1) and Eq. (2) by noticing that keeping π_B fixed and increasing T_B yields a larger T_G. When T_G is large, it is very likely that the links that are initially in good state will remain good for a long time, and that routing the packet across these links will be successful at the first attempt. However, this benefit disappears with longer deadlines. The more important observation is that with longer T_B it becomes harder to obtain very high reliability even for long deadlines. In this case, if a packet gets blocked at a node where all outgoing links are in bad state, it is likely to suffer a long delay for channel recovery before it can be forwarded.

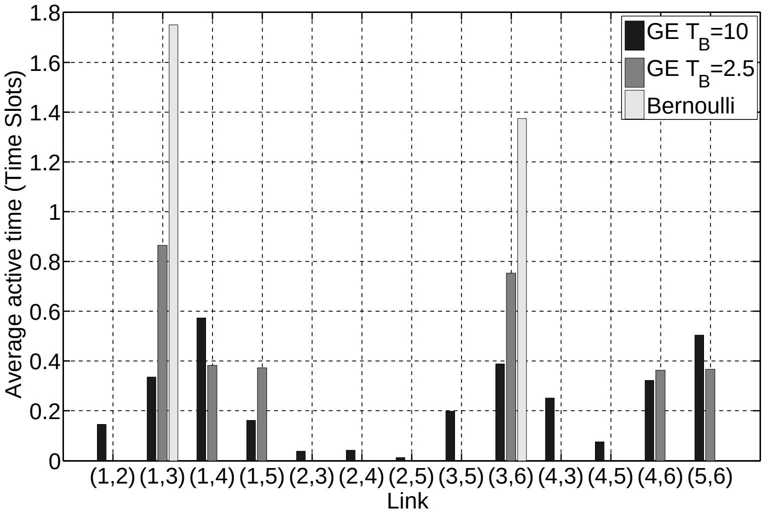

Figure 9: Average active time under information pattern (b) with $D = 4$.

Figure 9 shows the average number of times that each link has been used before the packet is dropped or received at the destination under information pattern (b) over 10^6 realizations of Monte Carlo simulation. In this specific example with homogeneous links, the single path $(1 \to 3 \to 6)$ is optimal when the packet losses are independent and described by the Bernoulli model. However, as links become increasingly bursty, it becomes increasingly beneficial to use multi-path routing to

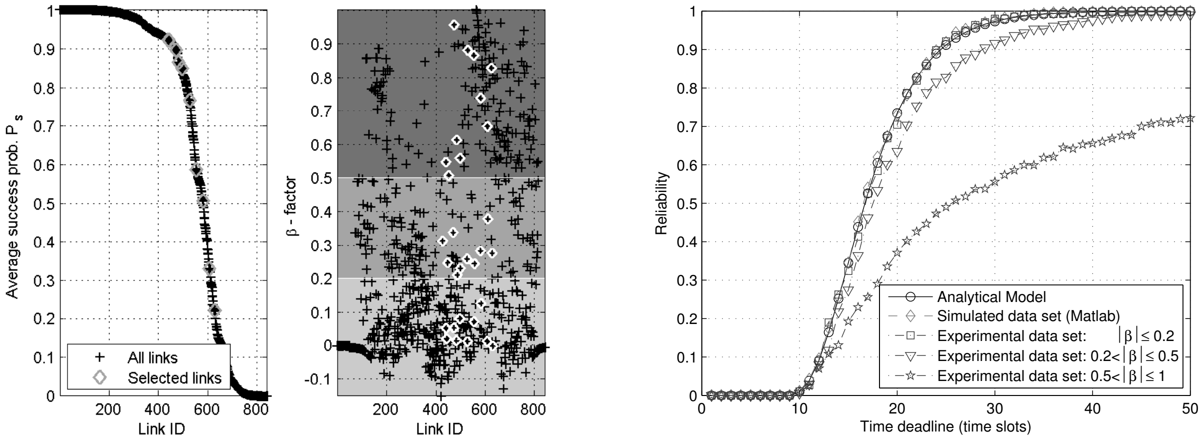

(a) Data link sets analysis in terms of PRR and β-factor.

(b) Model and scheduler validation for a 10-node line network.

Figure 10: Analytical model and scheduler validation against synthetic and experimental data in a 10-node line network. We identify 10 triplets of traces with the same PRR, but different packet loss correlation. The synthetic set of data is generated in Matlab using a Bernoulli process with the same PRR as the selected experimental data traces.

avoid links that are in, or likely to be in, bad state. For bursty links described by the GE model with $T_B = 2.5$, the packets follow three paths $1 \to 3 \to 6$, $1 \to 4 \to 6$ and $1 \to 5 \to 6$. For longer burst length $T_B = 10$, all possible paths are used.

3.4 Validation on real data

In addition to the simulation-based validataion on synthetic data presented above, we have also validated our models and scheduling schemes against real data. The evaluations use both publically available traces and data that we have collected at KTH and in Barcelona. For the ease of description, we consider a 10-node line topology and deadlines $D \in [1, 50]$ and present the evaluation of the scheduler developed under Bernoulli loss assumption. The application to the Barcelona small-scale demo is addressed in the next section.

In compliance with the time slot length of WirelessHART, experimental data traces are collected with nodes sending packets with 10ms inter-packet interval. We characterize the packet loss correlation with the methodology in [11], which quantifies the link burstiness through a scalar, the β-factor, defined as the Kantorovich-Wasserstein (KW) distance of the empirical data compared to the KW distance of an independent link with the same packet reception rate (PRR). A perfectly bursty link has $\beta = 1$, while $\beta = 0$ corresponds to independent losses.

Figures 10(a) show the PRR and β for all traces. In particular, we define three ranges of β corresponding to increasing packet loss correlation, and we identify triplets of links with the same PRR belonging to the three different regions. Figure 10(b) shows that the predicted model reliability assuming Bernoulli losses matches exactly the deadline-reliability curves obtained with synthetic data generated in Matlab. Furthermore, although the analytical model assumes independent packet losses, it accurately predicts the achievable deadline-reliability curves when packet losses are correlated with burstiness factor $\mid \beta \mid \leq 0.5$. In particular, the analytical reliability perfectly matches the scheduler reliability for link traces with $\mid \beta \mid \leq 0.2$, while a small gap occurs for links with $0.2 < \mid \beta \mid \leq 0.5$. Not surprisingly, for larger β, i.e. very high loss correlation, the model overestimates the scheduler reliability since it assume independent packet losses. In this case, the Gilbert-Elliot model is a better fit.

4 Application to the Barcelona water distribution network

The control architecture of the Barcelona distribution network comprises long-range communication on a slow time-scale for collecting sensor data and disseminating set-points and ultra-reliable short-range communication on a fast time scale for the local set-point tracking.

As discussed already in D2.1, the small-scale demonstrator is rather trivial since no relays are needed (both range and reliability is sufficient) and since sensor-controller and controller-actuator communication are naturally separated in time (and hence do not interfer). For such a simple scenario, we cannot rely on spatial diversity but can only use time (retransmissions) and frequency (channel hopping) to improve reliability. Moreover, there are no significant challenges in the scheduling of multiple real-time streams, but they can be studied in isolation (using the abstraction described above).

Hence, our proposal for this scenario is to schedule back-to-back retransmissions with a channel hopping sequence that jumps sufficiently far in the 2.4 GHz band to decorrelate the channel. Clearly "sufficiently far" depends on what type of outage events we expect, but if 802.11 interference is the main source then 12-19-26 should be a useful channel hopping pattern. We have not yet collected measurements that allow us to judge if this channel hopping pattern decorrelates losses at the small-scale demo site but we have collected individual traces for each of these channels at the targeted deployment locations. As discussed in D2.1, the Gilbert-Elliot model parameters for each of these channel traces gave very similar parameters (as long as the interpacket generation times stayed the same). Hence, scheduling R back-to-back retransmissions with channel hopping, assuming that loss events are uncorrelated in frequency, would give a end-to-end loss rate of

$$\pi_B^R$$

while if we assume that the state evolution of all channels follow the same Markov chain (i.e. the states are totally correlated), then the corresponding quantity is

$$\pi_B p_B^{R-1}$$

Considering, for example, the 15 ms traces described in D2.1, we have $\pi_B = 0.0233$ and $p_B = 0.1795$. Although the losses are below 0.1% already for two retransmissions in both scenarios, the loss rates decay a factor $p_B/\pi_B \approx 7.7$ faster when the channel hopping sequence manages to decorrelate the channel. The loss rates as function of the number of retransmissions for the two scenarios are shown in Figure 11.

References

[1] D. Chen, M. Nixon, and A. Mok, Wireless*HART*^{TM}*:Real-Time Mesh Network for Industrial Automation*. Springer, 2010.

[2] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz, "Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer," *IEEE Transactions on Industrial Electronics*, vol. 49, no. 6, pp. 1265–1282, Dec. 2002.

[3] E.N.Gilbert, "Capacity of burst-noise channels," *Bell System Technial Journal*, vol. 39, pp. 1253–1265, 1960.

[4] E.O.Elliott, "Estimates of error rates for codes on bursty-noise channels," *Bell System Technical Journal*, vol. 42, pp. 1977–1997, 1963.

[5] P. Soldati, H. Zhang, and M. Johansson, "Deadline-constrained transmission scheduling and data evacuation in WirelessHART networks," in *Proceedings of the 10th European Control Conference (ECC)*, 2009.

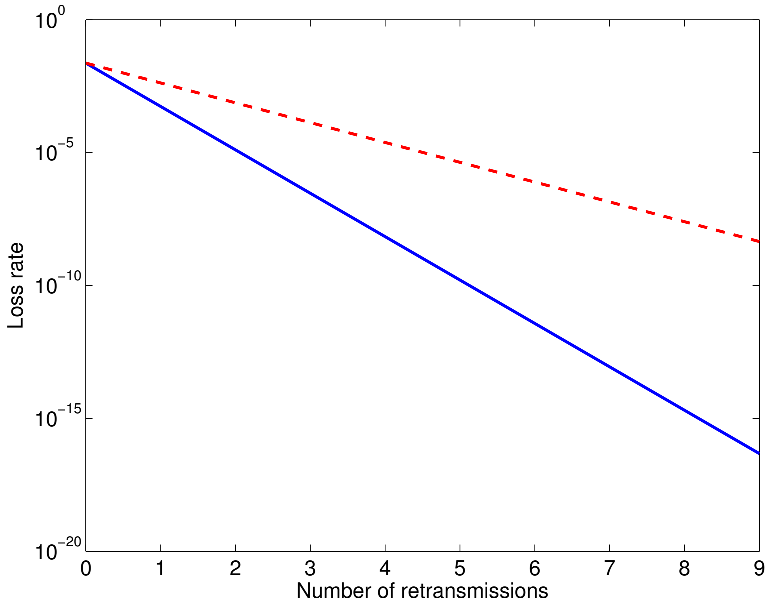

Figure 11: Loss rates for correlated channel (full) and decorrelated channel (dashed) as function of the number of back-to-back retransmissions. Channel data is taken from traces collected using 15 ms time slots.

[6] H. Zhang, P. Soldati, and M. Johansson, “Optimal link scheduling and channel assignment for convergecast in linear WirelessHART networks,” in *Proceeding of the 7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2009.

[7] J. Pesonen, H. Zhang, P. Soldati, and M. Johansson, “Methodology and tools for controller-networking codesign in wirelesshart,” in *Proceeding of the 14th IEEE International Conference on Emerging Techonologies and Factory Automation (ETFA)*, 2009.

[8] H. Zhang, F. Österlind, P. Soldati, T. Voigt, and M. Johansson, “Rapid convergecast on commodity hardware: Performance limits and optimal policies,” in *Proceeding of the 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 2010.

[9] P. Soldati, H. Zhang, Z. Zou, and M. Johansson, “Optimal routing and scheduling of deadline-constrained traffic over lossy networks,” in *Proceeding of IEEE Global Telecommunications Conference (GLOBECOM)*, 2010.

[10] Z. Zou, P. Soldati, H. Zhang, and M. Johansson, “Delay-constrained maximum reliability routing over lossy links,” in *Proceeding of the 49th IEEE Conference on Decision and Control (CDC)*, 2010.

[11] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis, “The beta-factor: measuring wireless link burstiness,” in *SenSys*, 2008, pp. 29–42.