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

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

WIDE
Decentralized and Wireless Control of Large-Scale Systems

Starting date: 01 September 2008 Duration: 3 years

Deliverable number	**D2.5**
Title	**Protocols and distributed processing for state estimation and monitoring**
Work package	WP2 – Wide-area WSN for distributed control
Due date	M34
Actual submission date	01/10/2011
Lead contractor	KTH
Author(s)	Mikael Johansson `mikaelj@ee.kth.se` Olaf Landsiedel `oland@ee.kth.se` Zhenhua Zou `zzou@ee.kth.se` Euhanna Ghadimi `euhanna@ee.kth.se` Andrea Abrardo `abrardo@dii.unisi.it`
Nature	Report
Revision	v1.1 (April 16, 2012)

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 describes the efforts within WIDE to develop protocols for estimation, monitoring and control. Three protocol alternatives are considered, that all try to achieve high end-to-end reliability and low delays. One set of protocol can give guarantees for meeting hard deadlines in a centrally managed small to medium size mesh networks, while the others are scalable protocols based on distributed network management and with strong latency and energy performance, but without hard real-time guarantees.

Contents

1 Introduction

Real-time decision-making, be it condition-monitoring, state-estimation, control or alarm-handling, put constraints on the latency and reliability of end-to-end transmissions. When the wireless network is based on battery-operated low-power wireless nodes, then the energy consumption of the network nodes add an additional dimension to the problem complexity. Other issues, such as quantization and source/channel coding, is less relevant when operating on a packet-based radio standard such as 802.15.4, and it also has limited effect on the performance of estimators [Joh05]. This report details three protocols that have been developed within WIDE to balance reliability and latency of end-to-end communication, on the one hand, and energy-consumption in the network, on the other.

The first set of protocols consider unicast-forwarding of packets across a multi-hop network where individual links are unreliable and may drop packets. In this case, we can derive limits of performance for the latency-reliability trade-off, and characterize the energy consumption required to meet given deadline and reliability targets, hence characterize the full performance space. The theoretical analysis assumes that the network is centrally managed and runs a globally synchronized multi-channel TDMA protocol (as is the case in, e.g., *Wireless*HART or 802.15.4e).

The second protocol focuses on efficient contention-resolution for single-hop communication in dense networks with highly correlated traffic. One example of such a traffic scenario is when a number of sensors, all within range of each other, simultaneously measure a physical quantity and attempt to transmit to their data to a common gateway. Another example of such a situation is when nodes are heavily duty-cycled to save energy, so that whenever the gateway node wakes up, there is a high chance that the other nodes have packets to send to it. Contrary to our first set of protocols, this protocol is a decentralized medium access protocol. An interesting feature of this protocol, not yet explored in our work, is that it could readily support priorities between the contending nodes.

Our final protocol is a scalable solution to highly resilient low-delay forwarding. In contrast to our other protocols, this protocol uses opportunistic forwarding and limited flooding to ensure low-delay forwarding in dynamic networks. We have not included any ambitious literature study in this report; such efforts can be found in the individual papers upon which this report is based, as well as in recent survey articles about mission-critical sensor networking protocols [SRS11].

Some words regarding the content of this report are in order. While this deliverable was initially supposed to cover both in-network processing techniques and protocol developments, the Consortium decided to concentrate the attention on protocol development for estimation, monitoring, and control, considered a more useful contribution to the remaining workpackages. Earlier work in Task 2.3 for in-network processing based on distributed source-coding for efficient estimation and monitoring performed during Year 1 by UNISI (which is no more a project partner of the project since the end of Year 2) and was not initially included in this report. A short paragraph describing the work performed was included in the revised version of this deliverable upon request from the reviewers.

2 Deadline-constrained maximum reliability routing

2.1 Motivation

The current industrial standards for low-power wireless (*Wireless*HART, ISA 100) and emerging standard initiatives such as 802.15.4e are based on a centrally managed time-synchronized channel hopping MAC. The centralized management allows us to construct and implement transmission schedules with global optimality properties. In particular, within WIDE, we have developed solutions to maximize the probability that packets meet a given hard deadline. The solution constructs a jointly optimal routing and scheduling policy for a single unicast flow. We have been able to find optimal strategies when link losses are independent in time (e.g. follow a Bernoulli loss process) and when they are correlated and losses occur in bursts. In addition, we have been able to account for transmission energy consumption in these solutions, to characterize and explore the optimal trade-off between reliability and energy consumption for a given deadline.

It is important to understand that the optimal solution to the deadline-constrained maximum reliability routing problem does not, in general, separate into a routing problem and the problem of scheduling transmissions across these given routes. Figure 1 shows an example that highlights this issue. Packets are normally forwarded along the more reliable two-hop path. However, if a packet has experienced multiple losses and remains at the source when there is only a single slot left until the deadline, the optimal transmission decision should be to try the less reliable direct path. Hence, we notice that the optimal routing path does not only depend on the link qualities, but also the remaining time to deadline in a non-trivial manner.

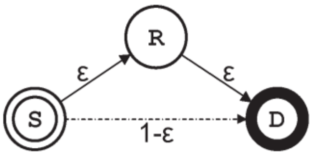

Figure 1 The longer path consists of two very reliable links. However, a packet with deadline 1 will be transmitted over the unreliable direct path.

2.2 Basic concepts

Low-power wireless links are notorious for high loss rates. Although erasure events on links are independent in some scenarios (e.g., static networks with long inter-packet transmission times), packet losses are typically correlated in time [WKHW02]. If a packet is routed over a bursty link which happens to be in a bad state, it will suffer a long delay waiting for the channel to recover. This has a negative impact on deadline-constrained reliability since there will be fewer opportunities for retransmissions on the remaining path to the destination. Furthermore, in low-power wireless applications, nodes are battery-powered and energy-efficiency is critical. It then becomes important to use retransmissions judiciously and avoid transmitting on links that are likely to be in bad state, even if this comes at the price of a slightly smaller end-to-end reliability.

To improve routing performance, several recent and emerging standards, such as *Wireless*HART, ISA100, IEEE 802.15.4e and IETF Roll support multi-path routing for enhanced reliability in wireless mesh networks with bursty links. However, they offer little insight on how to optimize routing and forwarding in these systems to ensure efficient and reliable real-time communications. In this line of work, we have developed optimal forwarding policies for reliable and energy-efficient real-time communication over lossy networks. We focus on a communication scenario where a single packet is to be transmitted on the network to the destination by a deadline. Periodic communication of samples from a sensor to a controller in a networked control system can then be realized by periodically repeating this optimal solution. Our objective is to maximize the delivery reliability while minimizing the number of transmissions to save energy consumption.

We assume that the packets are routed along a destination-oriented directed acyclic graph (DODAG) on a wireless multi-hop network. Communication is slotted, and a single time slot allows the transmission of a packet and the associated link level acknowledgment. Communication links are unreliable, with erasure events modeled through a Markov chain. For simplicity, we will illustrate our result for the case that the loss process on each link is described by a two-state Markov chain with a “good” (G) and a “bad” (B) state corresponding to successful transmission and packet loss, respectively; see Figure 2 for an illustration. This model reduces to the Bernoulli model of independent packet erasures when $p_B = p_G$ and $q_B = q_G$.

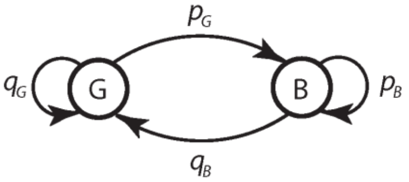

Figure 2 - Two-state Markov chain model for link packet erasures.

Let $\boldsymbol{\omega}_i(t) = [\omega_{ij}(t)]$ represent the state of all links (i, j) outgoing from node i at time t, with $\omega_{ij}(t)$=1 if the link is in good state and $\omega_{ij}(t)$=0 if the link is in bad state.

We consider a scenario where a single packet, generated by an arbitrary node i at time $t = 0$, should be transmitted over the DODAG to the sink node within a deadline of D time slots. The goal is to forward packets in a way that optimally balances deadline-constrained reliability and the expected transmission energy cost. To this end, each node i is associated with a utility

$$U_i(t) = R_i(t) - \delta C_i(t)$$

that describes the tradeoff between the reliability $R_i(t)$ for packet delivery within the next $D - t$ time slots and the associated energy consumption $C_i(t)$, given information up until time $t \leq D$. The variable δ describes the energy cost per transmission with full power but can also be used to balance the two conflicting objectives when we want to generate the Pareto-optimal trade-off surface between deadline-constrained reliability and expected energy consumption of the forwarding operation.

The optimal policy hinges on what knowledge we can assume about the state of the underlying Markov chain at the time when we make the forwarding decision. Let $\boldsymbol{I}_i(t)$ denote the link state information available at node i and time t. In our work, we have considered two information patterns: In information pattern (a), $\boldsymbol{I}_i(t) = \boldsymbol{\omega}_i(t)$ which denotes that the state of the outgoing links at time t prior transmitting is known; in information pattern (b), $\boldsymbol{I}_i(t) = \boldsymbol{\omega}_i(t-1)$ which denotes that the state at time $t - 1$ is known.

We then develop a dynamic programming framework to compute the optimal forwarding policy given the link state information and the time to deadline. Let $U_i^*(t|\boldsymbol{I}_i(t))$ be the maximum utility conditioned on the link state information $\boldsymbol{I_i}(t)$, and it is computed as:

$U_i^*\left(t\middle|\boldsymbol{I}_i(t)\right) = \max\{\max_j U_i^j\left(t\middle|\boldsymbol{I}_i(t)\right), U_i^i\left(t\middle|\boldsymbol{I}_i(t)\right)\}$ where

$$U_i^j\left(t\middle|\boldsymbol{I}_i(t)\right) = \sum_{\boldsymbol{\omega}_i(t)} \Pr\{\boldsymbol{\omega}_i(t)|\boldsymbol{I}_i(t)\}(\underbrace{\mathbf{1}_{\{\omega_{ij}(t)=1\}} U_j^*(t+1)}_{\text{Successful forwarding}}$$
$$+ \underbrace{\mathbf{1}_{\{\omega_{ij}(t)=0\}} \textstyle\sum_{\boldsymbol{I}_i(t+1)} \Pr\{\boldsymbol{I}_i(t+1)|\boldsymbol{\omega}_i(t)\}(U_i^*(t+1|\boldsymbol{I}_i(t+1)))}_{\text{Failed forwarding}} - \underbrace{\quad\delta\quad}_{\text{Tx Penalty}};$$

$$U_i^i\left(t\middle|\boldsymbol{I}_i(t)\right) = \underbrace{\textstyle\sum_{\boldsymbol{I}_i(t+1)} \Pr\{\boldsymbol{I}_i(t+1)|\boldsymbol{I}_i(t)\}\, U_i^*(t+1|\boldsymbol{I}_i(t+1))}_{\text{Staying at node i}}.$$

Thus, the optimal forwarding policy is to forward the packet to a node such that the utility is maximized:

$$j_i^*\big(t\big|\boldsymbol{I}_i(t)\big) = \begin{cases} i & \text{if } U_i^i\big(t\big|\boldsymbol{I}_i(t)\big) \geq U_i^j\big(t\big|\boldsymbol{I}_i(t)\big)\ \forall j; \\ \arg\max\limits_j\{U_i^j\big(t\big|\boldsymbol{I}_i(t)\big)\} & \text{otherwise} \end{cases}$$

The dynamic programming recursion proceeds backwards in time from $t = D$ to $t = 0$. The solution to this problem allows to characterize the set of achievable latency-reliability pairs and to trace out the Pareto frontier between achievable deadline-constrained reliability and transmission energy cost. It also allows to "roll out" a global transmission schedule that can be implemented in *Wireless*HART compliant devices.

2.3 Results

We demonstrate our algorithms with numerical examples on the topology in Figure 3, where a source (node 1) sends packets to a sink (node 6), and analyze the end-to-end reliability for different loss models and information patterns. We assume homogeneous links with unconditional loss probability 0.5 for both the Bernoulli and GE models.

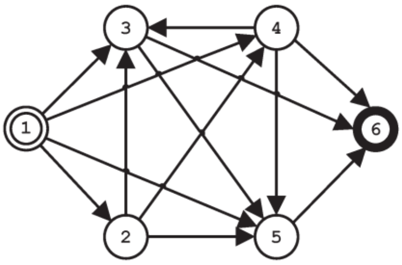

Figure 3 - Example DODAG topology

Figure 4 shows the achievable deadline-reliability curve under different channel models and information patterns. We can see how the available link-state information impacts the achievable reliability. Information pattern (a), where nodes know which outgoing links will be successful when they make forwarding decisions, always outperforms information pattern (b), where nodes cannot be certain that transmissions will be successful. The difference between the two decreases as the links get more bursty, since the link state at the previous time instant then becomes an increasingly accurate prediction of the link state at the current time.

We also compare the optimal forwarding policy with the minimum end-to-end delay routing scheme that chooses the minimum ETX [CABM03] single path in Figure 5. The minimum ETX single path can obtain the same performance under the Bernoulli loss model, but its performance degrades when links become more bursty.

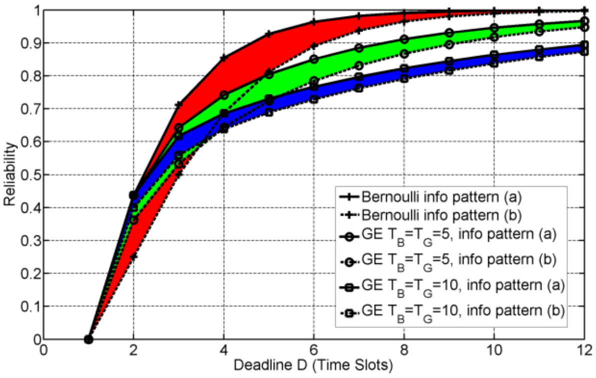

Figure 4 Deadline-Reliability for information patterns (a) and (b).

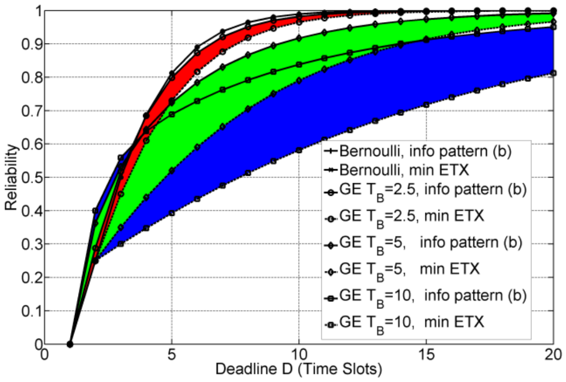

Figure 5 Comparison with min ETX path routing.

Next, we analyze the trade-off between achievable end-to-end reliability and transmission energy cost for different deadline constraints. All links have unconditional loss probability 0.5

while the burstiness parameter p_B is drawn uniformly from $[0.65, 0.85]$. The trade-off between reliability and transmission energy cost for different deadline constraints is shown in Figure 6. The maximum reliability is obtained for $\delta = 0$, i.e., when no considerations for the transmission energy cost are made. In this case, nodes will always try to transmit, even if all its links are likely to be in a bad state, provided that the remaining time to deadline is not less than the minimum hop count to the sink. As seen in Figure 7, the final few percent of reliability comes at a very high energy price. For $D = 12$, the final 3% of reliability doubles the transmission energy cost.

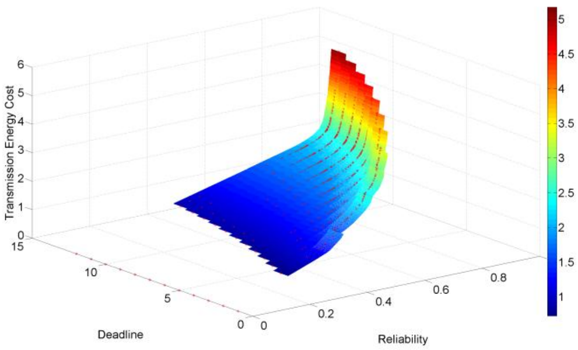

Figure 6 Deadline-Reliability-Energy surface for information pattern (b) with GE loss model parameters p_B.

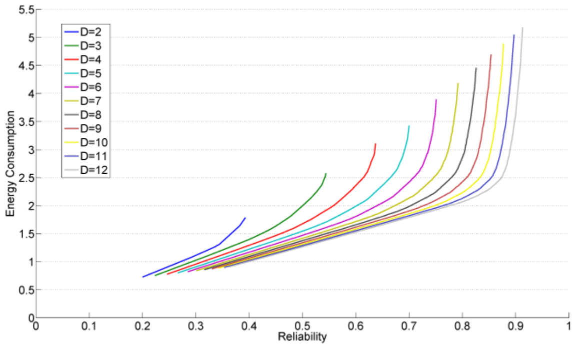

Figure 7 Reliability-Energy curves for information pattern (b) with GE loss model parameters p_B.

The work on deadline-constrained scheduling has been, in part, described in earlier deliverables but is being continuously refined and has also formed the foundation for the optimal co-design framework developed in WP4.

3 Efficient contention-resolution in networks with correlated traffic

3.1 Overview

Wireless sensor networks experience traffic bursts due to spatially temporally correlated events, bulk transfers and route and code updates [JP07]. In control applications, for example, sensing events are correlated when time-triggered control architectures (such as classical digital control solutions) are used, but also event-triggered controllers tend to fire multiple events at the same time (e.g. when a disturbance hits a process).

Traffic bursts aggravate the hidden terminal problem, as nodes that are hidden to each other may attempt to simultaneously send data to the same neighbor, causing data collisions and losses. Traditional contention resolution mechanisms are backoff-based and are suspectible to hidden terminals. The capture effect phenomenon on the other hand, allows a radio to correctly receive a data transmission even with simultaneous colliding transmissions. The capture effect requires that (1) the overlapping transmissions differ in signal strength, and (2) that the stronger transmission is initiated before the interfering weaker transmission(s).

We have conducted a profiling study to quantify the effects of hidden terminal and capture phenomena on the overall performance of the CSMA-based contention resolution mechanism. To this end, we considered a star topology with a single receiver. Dependent on packet reception rate (PRR) between each two nodes, their link might be either one or bi-directional. The number of (non-detectable) neighbor links, divided by all possible links defines our hidden terminal metric.

We extract a realistic range of hidden terminal metrics by performing a set of experiments on the publicly available TWIST sensor network testbed [HKW06]. Our experiments show that the hidden terminal metric varies significantly among different nodes in the network. The hidden terminal metrics in the TWIST ranges between 11.0% and 29.4%. Using our experiment traces, we furthermore model the capture effect effectiveness by the amount of links that their signal strength differs with more than 3 dB.

Figure 8, depicts goodput as a measure of performance versus increasing hidden terminal metric. Goodput is defined as the ratio of time that the receiver is receiving uncorrupted network-layer data from any of its neighbors. Plot shows that the performance degrades with increasing size of the network $N = 5,\ 10,\ 20$.

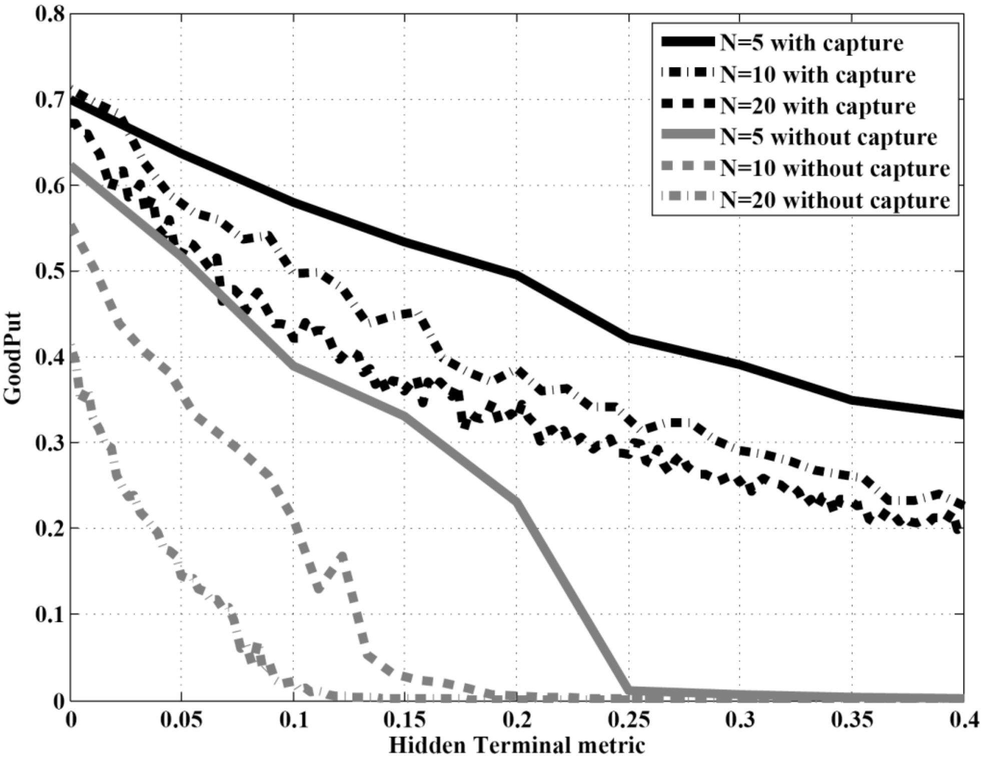

Figure 8 The performance of CSMA with random backoff degrades significantly with increased hidden terminals

The emerging class of receiver-initiated duty-cycled MAC protocols [DDC10], promises both reduced congestion and improved resilience against hidden terminals, in comparison to traditional sender initiated protocols [BYA06]. Within WIDE, we have analyzed, modified and redesigned a particular protocol called Strawman [ÖWT10]. Strawman's contention resolution mechanism – designed for receiver initiated duty-cycled protocols – mitigates the hidden terminal problem through an RTS/CTS-like handshake, but is lightweight and readily implementable on publically available hardware platforms.

Request-To-Send/Clear-To-Send mechanisms (RTS/CTS) have long been employed to mitigate the hidden terminal problem, but they suffer from high overhead in low-power sensor networks [PHC04], and are not often employed in practice. Strawman solves the hidden terminal problem efficiently by measuring which of multiple colliding random-length RTS transmissions is the longest. The contender that (randomly) picks the longest length is granted channel access and sends its data.

While Strawman has many promising properties, the initial design also has some drawbacks that limits its throughput and scalability. In our work, we have modelled the basic Strawman mechanism and derived an enhanced version with improved performance. We have also designed, analyzed and evaluated an optimal non-uniform request length distributions which outperforms the uniform distribution. Through extensive simulations, we have analyzed how hidden terminals and the capture-effect affects our contention-resolution mechanisms, and demonstrated improved performance of our mechanisms.

3.2 Protocol design

The basic Strawman protocol is illustrated in Figure 9. Each Strawman contention period consists of four consecutive messages: PROBE, REQUEST, DECISION, and DATA. The receiver broadcasts a Strawman PROBE message to notify neighbors that it is ready to receive data. All neighbors that have data for the receiver contend for the channel by sending an immediate REQUEST. Multiple REQUESTs may thus collide at the receiver. The length of each REQUEST message is chosen randomly by drawing samples from a uniform distribution [ÖWT10]. The receiver samples the channel for activity during the REQUESTs, and estimates the payload length of the longest REQUEST. The receiver then sends a DECISION message containing the length estimate. The contender whose REQUEST length matches the one specified in the DECISION is granted channel access, and sends its DATA message.

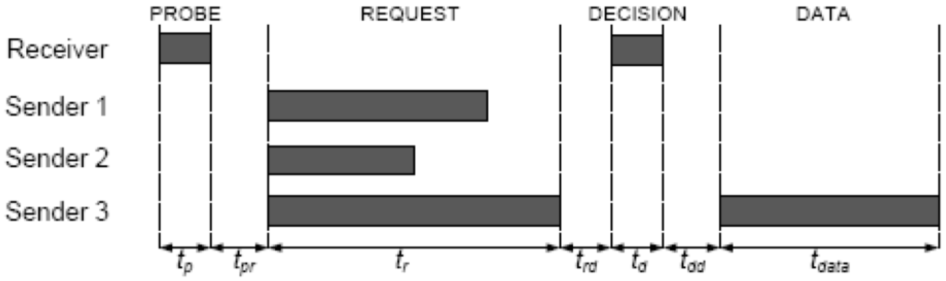

Figure 9 The Strawman contention resolution mechanism grants channel access to the contender with the longest request transmissions.

Parameter	**Description**
N	Number of contenders
K	Maximum number that nodes can pick in REQUEST phase
t_p	Time for transmitting the PROBE packet
t_{pr}	Delay between PROBE and REQUEST packets
t_r	Duration of REQUEST phase
t_{rd}	Delay between REQUEST and DECISION packets
t_d	Time for transmitting the DECISION packet
t_{dd}	Delay between DECISION and DATA packets
t_{data}	Time for transmitting DATA packet

The next contention round is initiated when the DATA has been received, or after a timeout. The PROBE message has dual purpose: it also acknowledges the last received DATA packet. Note that if two contenders pick the same random length, and hence are both granted channel access leading to a DATA collision, the timeout will trigger another contention period, and both data packets will be retransmitted due to the lack of acknowledgement. Note also that the Strawman contention resolution is used only if a receiver detects a data collision, and thus has zero overhead when events are not occurring simultaneously.

We model a snapshot of the network with a receiver node and a set of N contending transmitters. Upon detecting a collision, the receiver sends a Strawman PROBE packet as illustrated in Figure 9. Transmitters send a REQUEST packet with random length $\mathrm{x_n}$chosen with a given distribution in region $[1, \mathrm{k}]$, where the maximum length K is referred to as the Strawman resolution. Success probability of a single Strawman contention round $\mathrm{P_{N,K}}$ is the probability that at least one contender draws a number $\mathrm{x_n} = \mathrm{k}$ (with probability $\mathrm{P_k}$) while all other contenders draw smaller numbers, i.e.

$$
\begin{gathered}
P_{N,K} = Np_K(1 - p_K)^{N-1} + \; Np_{K-1}(1 - p_K - p_{K-1})^{N-1} \\
+ \cdots + Np_1(1 - p_K - p_{K-1} - \cdots - p_1)^{N-1} \\
N \sum_{k=1}^{K} p_k \left(1 - \sum_{r=k}^{K} p_r\right)^{N-1}
\end{gathered}
$$

Reliability is defined as the number of transmissions successfully received divided by the total number of contention rounds. Given above metrics we first tune the basic Strawman protocol (with uniform distribution) to get the best performance in terms of success probability and goodput. Second, we derive the optimal distribution which maximizes the probability of success per round can be derived using convex optimization techniques. The interested reader may refer [GSÖ11] for detailed analysis.

Although computing the optimal probability is tractable, it does not provide a closed-form expression for the optimal probability distribution. To develop a better understanding of the optimal distribution, we studied two different approximations that simplify the computation of the request-length distribution, still yield near-optimal performance. These approximations are derived by observing the shape of the optimal distribution p^*. First approximation follows a geometric distribution and is similar to SIFT [TJB04] - a technique which is used to approximate an optimal distribution in CSMA based contention resolution.

Although the SIFT approximation distributes probabilities coherently with the shape of the optimal probability distribution, it does not scale very well with the increasing number of contenders. The trapezoidal method is an alternative approximation that aims at imitating the shape of the optimal distribution more accurately. The approximation is inspired by the shape of the distribution for $\mathrm{k} \geq 2$. It turns out that $\mathrm{p_2^*}$ and $\mathrm{p_K^*}$ are two critical points in the mass of probability distribution function. We approximate these two probabilities for any given number of nodes then, $\mathrm{p_{k-1}^*}$, ... $\mathrm{p_3^*}$ are computed based on a linear interpolation function between $\mathrm{p_2^*}$ and $\mathrm{p_K^*}$. Eventually, $\mathrm{p_1^*}$ is given by putting the sum of probabilities equal by 1.

Figures 10 illustrates the success probability of a Strawman round with optimal distribution, trapezoidal and SIFT approximations, respectively, computed for fixed resolutions K = 8 and large number of contenders. Obviously, such a high range of contenders is not intended to reflect any practical scenario, but only to analyze numerically the protocol behavior. We observe that for almost all number of contenders the optimal distribution provides almost flat success probability curve which guarantees the low overhead of the protocol implementation. Small overhead of request phase offers the same good performance with huge increase of the network domain. Moreover, we can see in the figure that in all cases, the trapezoidal approximation always yields better performance rather than the SIFT approximation.

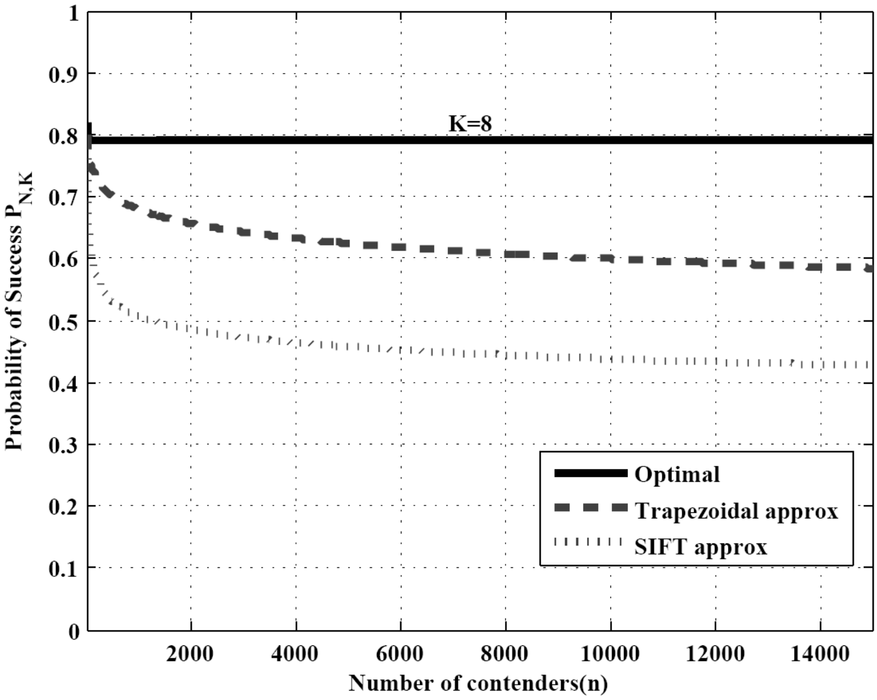

Figure 10 Scalability of Strawman request length optimization techniques.

3.3 Evaluation

We evaluate how our improved Strawman mechanism performs in terms of goodput, reliability, and compare it against the initial Strawman proposal.

Figure 11 validates the Strawman mechanism with trapezoidal approximation designed with resolution $K = 16$ and N ranging from 10 to 25 contenders. The figure shows an accurate match between the simulation and experimental values of the success probability and goodput of Strawman for all cases. Leveraging on this match, we will continue our evaluation through extensive simulations.

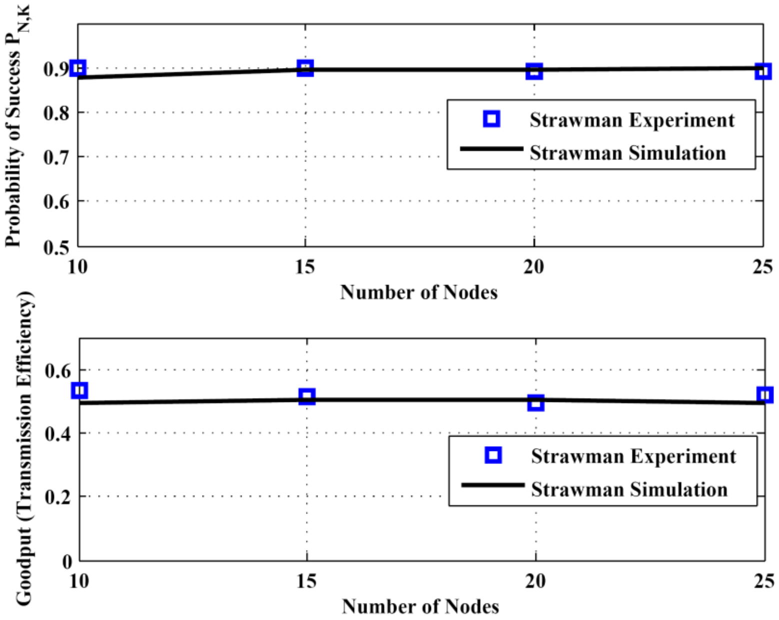

Figure 11 Success probability and goodput (the portion of network layer) validation with experimental result.

In the first experiment we simulate Strawman in a network without any hidden terminals, and measure reliability. A large number of (interfering) collisions lowers the performance. Figure 12 compares the reliability of Strawman with CSMA. Our first observation regarding the optimal and approximate request-length distributions of Strawman is that the performance of Strawman remains largely unaffected by increasing number of contenders. This behavior is due to the scalability property of these distributions. Figure 12 shows In the case of no hidden terminal both CSMA and Strawman have similar performance for uniform and optimal distribution. Moreover uniform distribution does not scale well with an increasing number of contenders.

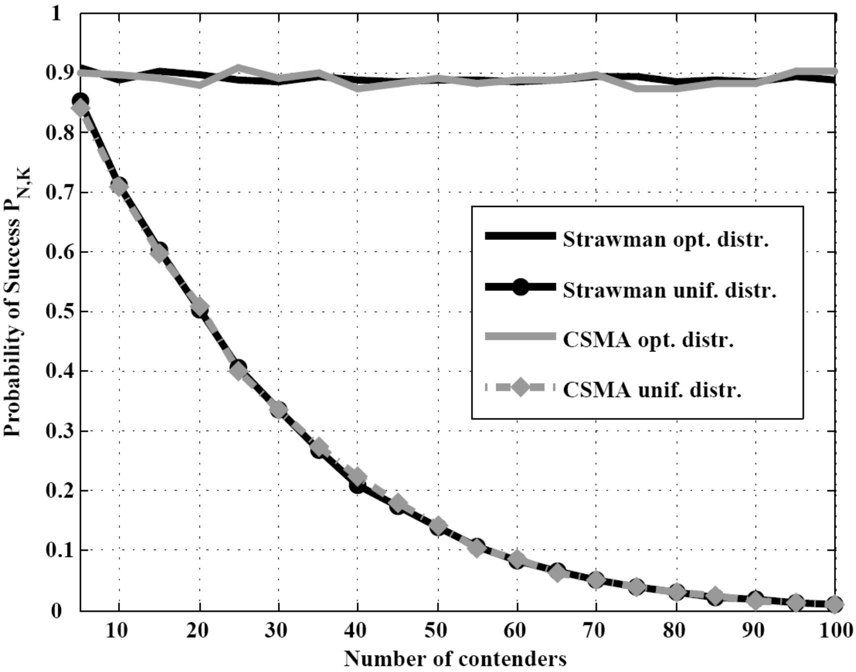

Figure 12 The reliability—the ratio successful transmissions—of Strawman and CSMA in an idealistic network setting without hidden terminals.

We now include hidden terminals as profiled previous Section. We do not, however, yet include the (positive) effects of the capture effect phenomenon. Neither Strawman nor CSMA is designed to exploit the capture effect, and so it is interesting to study how they behave without capture effect. Moreover, capture effect efficiency differs with network types and radio hardware. Figure 13 shows how the reliability of the mechanisms is affected by hidden terminals. Whereas Strawman with optimal distribution is unaffected by the addition of hidden terminals, CSMA suffers significantly. Moreover, Strawman with uniform distribution outperforms CSMA case. This is due to the inherent RTS/CTS like handshake mechanism in Strawman which mitigates hidden terminal problem.

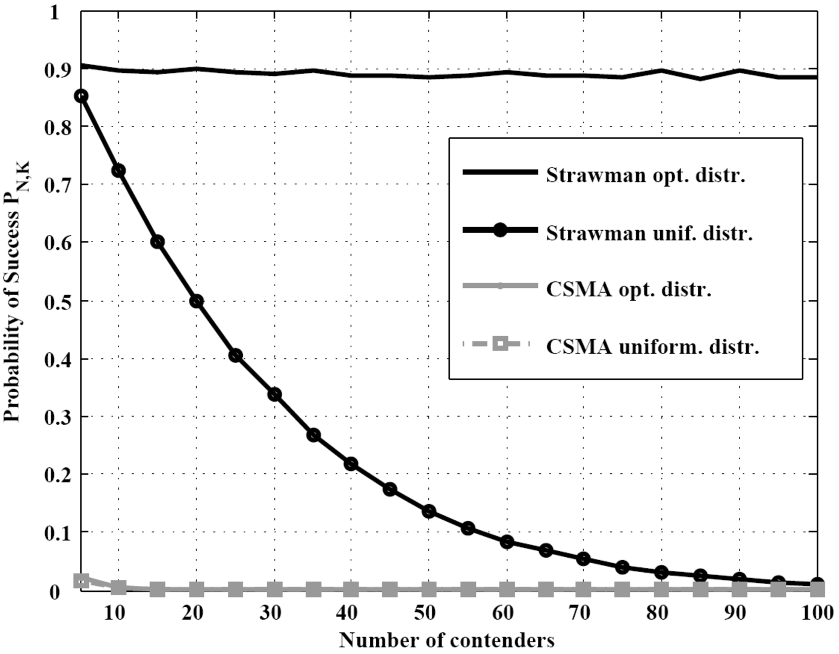

Figure 13 reliabilityof Strawman and CSMA with hidden terminal but no capture effect.

We finally enable the capture effect phenomenon, thus fully mimicking the testbed in previous Section. For CSMA, we can observe an increase in reliability (Figure 14) . Strawman is not simulated with capture effect in these experiments, and thus has the same performance. As these experiments show, Strawman outperforms CSMA even when CSMA benefits from capture effect. We observe an interesting phenomenon in these experiments: CSMA with uniform distribution appears to perform better with more contenders. At first glance, this is counter-intuitive: CSMA with uniform distribution was shown to scale badly even without hidden terminals (Figure 12). After careful studying of experiment logs we attribute this behavior to a complex interaction between capture effect and the random backoff-distribution. The uniform distribution renders more collisions at early stages of a transmission, whereas the optimal distribution achieves a higher probability of a single transmission. If a neighbor in the uniform-based network manages to successfully initiate a transmission to the receiver (due to capture effect), the probability is high that several other neighbors are also transmitting (with a lower signal strength), thus blocking the rest of the network from interfering transmissions. This initial result motivates us to further study the relation between capture effect and hidden terminals, and demonstrates that protocols must be evaluated in realistic but controlled environments.

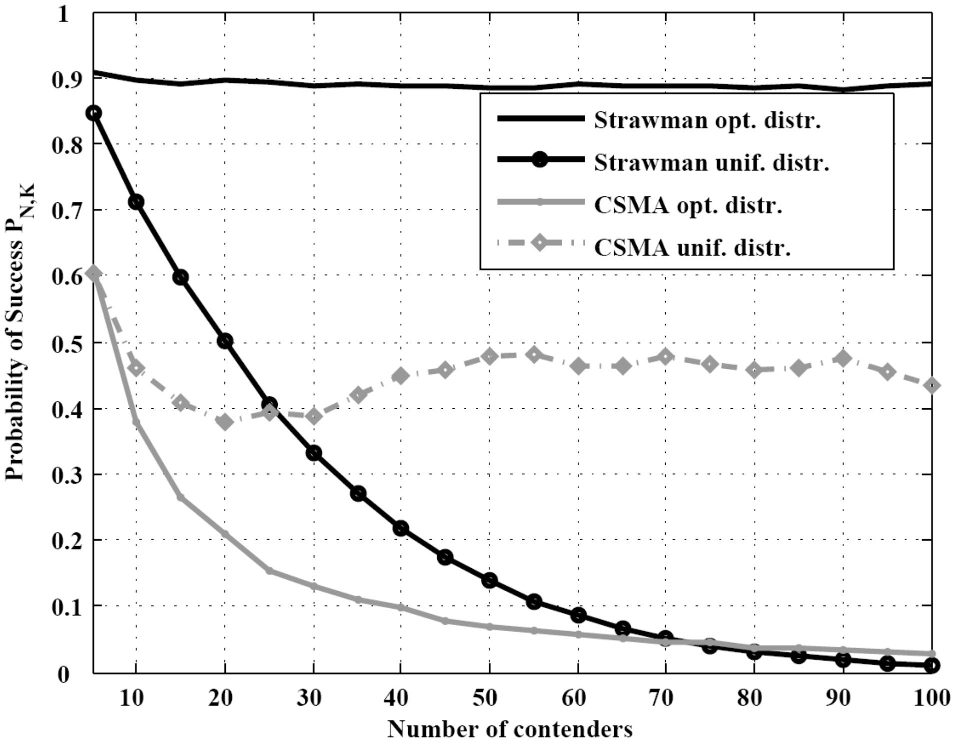

Figure 14 reliabilityof Strawman and CSMA with hidden terminal and capture effect.

This evaluation has compared Strawman with CSMA in three scenarios with increasing realism. We observe that hidden terminals, with or without capture effect, may greatly degrade performance of contention resolution mechanisms. Strawman is, however, shown to yield the same high performance both with and without hidden terminals.

4 Opportunistic forwarding for sensor networks

4.1 Motivation

In Wireless Sensor Networks (WSNs), forwarding of packets to their intended destination commonly relies on a two-step process: First, the routing protocol determines the next hop utilizing a routing metric and link estimations. Second, the MAC protocol waits for the intended destination to wakeup and to successfully receive the packet. In this work, we depart from this unicast design paradigm of sending a packet to a single forwarder. Instead, we transmit packets opportunistically in duty-cycled sensor networks: the first awoken neighbor that successfully receives the packet and offers progress towards the destination forwards it. This opportunistic forwarding leads to three key advantages: It improves energy efficiency, reduces end-to-end delay, and increases resilience to wireless link dynamics.

Originally, opportunistic routing was developed to improve throughput in multi-hop, mesh networks over lossy wireless links. It benefits from the fact that in mesh networks radios are always on and hence can overhear messages with practically no additional cost. In contrast, wireless sensor networks are commonly duty cycled to ensure long node and network lifetime. Moreover, WSN applications demand for high-energy efficiency, high reliability, and low delays. This limits the direct application of opportunistic routing to WSNs. The main distinction of this work over existing work on opportunistic routing is, that it applies the concept of opportunistic routing to WSNs and adapts it to the specific demands of sensor networks and applications.

Low-power links in WSNs are highly dynamic. To ensure stable routing, routing protocols in WSNs commonly rely on link estimation to identify links of consistently high quality for forwarding. Restricting forwarding to this subset of stable links trades routing progress for stability as dynamic links are often far ranging. Our main departure from this work is that our opportunistic forwarding explicitly utilizes all neighbors, i.e., both stable and unstable links, for packet forwarding.

4.2 Basic Concept

Our work, denoted ORW, targets duty cycled protocol stacks, such as asynchronous, low power listening [BYA06]: In these protocols a sender transmits a stream of packets until the intended receiver selected by the routing protocol wakes up and acknowledges it, see Figure 15a and Figure 15b. To integrate opportunistic routing into duty cycled environments, we depart from the traditional unicast forwarding in one key aspect: The first node that (a) wakes up, (b) receives the packet, and (c) provides sufficient routing progress, acknowledges and forwards the packet, see Figure 15c. For example, in Figure 15a node A can reach node C either directly via an unreliable link or via B. If A → C is too unreliable, i.e., PRR less than 50%, traditional routing will ignore A → C and rely on A → B → C. ORW extends this, by also including the unreliable link A → C into the routing process: If A → C is temporary available and C wakes up before B, ORW will utilize it for forwarding. This reduces the average number of hops, delay, and energy consumption.

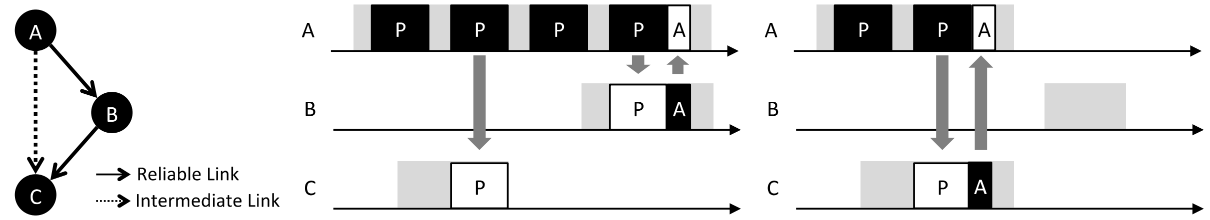

(a) Sample topology: Node A reaches C via B on reliable links or directly on an unreliable link.

(b) Traditional routing: Although C might overhear some transmission from A, packets are addressed to B to ensure stable routing.

(c) Opportunistic Routing in ORW: The first node that wakes up, receives a packet, and provides sufficient routing progress acknowledges and forwards it.

Figure 15: Basic idea of ORW: ORW compared to traditional routing on an asynchronous, low power listening MAC. ORW aims to reduce energy consumption and delay while providing high routing progress.

This design enables an efficient adaptation of opportunistic routing to the specific demands of wireless sensor networks: (1) In contrast to opportunistic routing in mesh networks, forwarder selection in ORW focuses on energy efficiency and delay instead of network through- put: It minimizes the number of probes until a packet is received by a potential forwarder. (2) Using packet streams it integrates well into duty cycled environments and ensures that many potential forwarders can overhear a packet in a single duty cycle. Thereby, ORW exploits spatial and temporal link diversity to improve resilience to wireless link dynamics. (3) The fact, that only a small number of nodes receive a probe at a specific point in time simplifies the design of a low-overhead coordination scheme to select a single forwarder. This limits the cost of the consensus protocol and ensures low control traffic.

4.3 Results

We compare the performance of our work, ORW, and the state of the art (Collection Tree Protocol, CTP [GFJ09]) in three evaluation scenarios. Figure 16 shows that ORW significantly improves energy efficiency and delay, while achieving similar, but slightly hops counts in all scenarios.

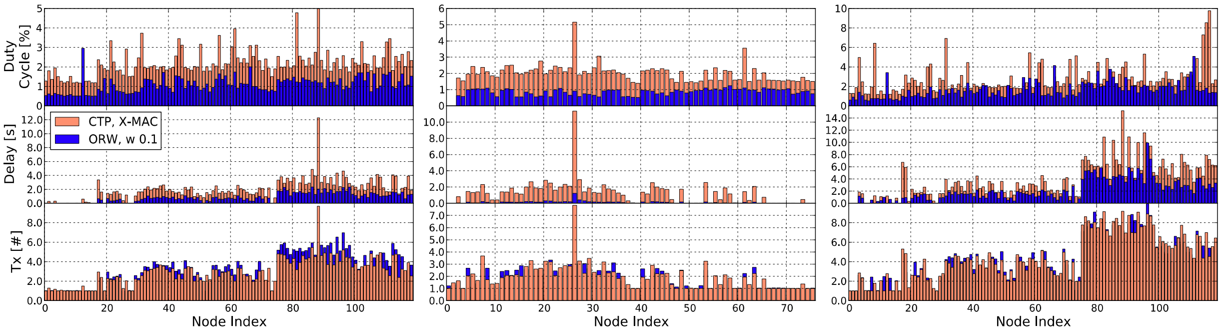

(a) Indriya, 0 dBm. (b) Twist, 0 dBm. (c) Indriya, -10 dBm.

Figure 16: Per Node Comparison of our work (ORW) and the state of the art (CTP): ORW improves duty cycles and delays while achieving slightly higher hop counts than CTP.

On average ORW doubles the energy efficiency, individual nodes show improvements up to a factor of 8.5. The results show that ORW strongly benefits from network density: It reduces the duty cycle the most on the dense Twist deployment with 0dBm. Similar, it shows the worst improvements, nonetheless a factor of 1.5, on Twist with −25dBm. Additionally, it improves delay by factor from 1.4 to 10.8 depending on network density and achieves hop counts that are similar, but slightly higher when compared to CTP.

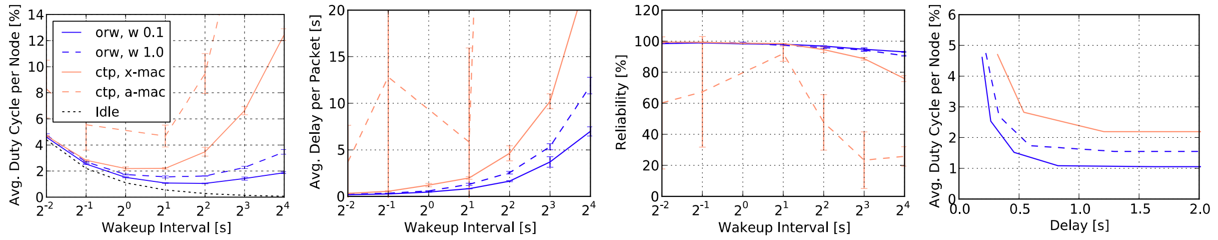

(a) **Duty Cycle**, Indriya 0 dBm (b) **Delay**, Indriya 0 dBm (c) **Reliability**, Indriya 0 dBm (d) **Duty Cycle over Delay**, Indriya 0 dBm

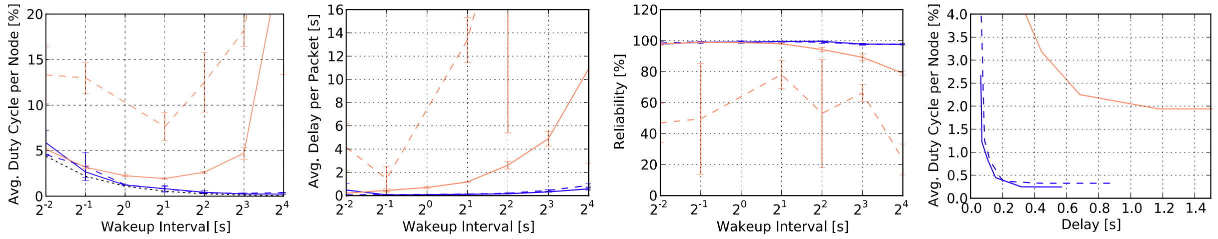

(e) **Duty Cycle**, Twist 0 dBm (f) **Delay**, Twist 0 dBm (g) **Reliability**, Twist 0 dBm (h) **Duty Cycle over Delay**, Twist 0 dBm

Figure 17: Performance of ORW and CTP for wakeup intervals between 0.25 and 16 seconds.

The previous experiments used a fixed wakeup interval of 2 seconds, i.e., a node wakes us every two seconds to receive data from neighboring nodes. We used this interval to ensure a fair comparison, as it leads to the optimal duty cycle in CTP and X-MAC at an IPI of 4 min (see Figure 17a and Figure 17e). Next, we discuss the impact of the wakeup interval on duty cycle, delay, and reliability. Figure 17a and Figure 17e show that ORW benefits much more than CTP from reduces wakeup intervals. It allows ORW to reduce its duty cycle further while increasing delay. The figures also depict the idle baseline, i.e., the energy that is consumed by just the wakeups of X-MAC without any data transfer. This base line defines the lower bound for the duty cycle. For Indriya (see Figure 17a) stays close to this line than CTP and in the dense Twist testbed (see Figure 17e) ORW is marginally above the base line throughout all experiments. These results show that ORW efficiently exploits network density.

Delay increases with increasing wake-up intervals (see Figure 17b and Figure 17f). However, the increase for ORW is significantly lower for ORW than for CTP with X-MAC. As a result, ORW can operate at much lower duty cycles. ORW exploits network density to reduce delay.

Hence, the delay increase in Twist is much lower than the one in Indriya. Figure 17d and Figure 17h show the resulting duty cycle for a given average delay. This underlines that ORW can operate at much lower duty cycles for a given average delay.

Both, CTP and ORW show high reliability (see Figure 17c and Figure 17g). However, at high wakeup intervals reliability of CTP decreases. This is due to queue overflows on individual nodes. In contrast, ORW avoids this by using multiple forwarders. Overall, the results for other inter-packet intervals (IPI) and channels show similar performance gains of ORW over CTP. Just the optimal duty cycle for CTP with X-MAC varies depending on traffic load. As reference, the figure also depicts results for CTP on A-MAC: Its performance is much lower in all key metrics.

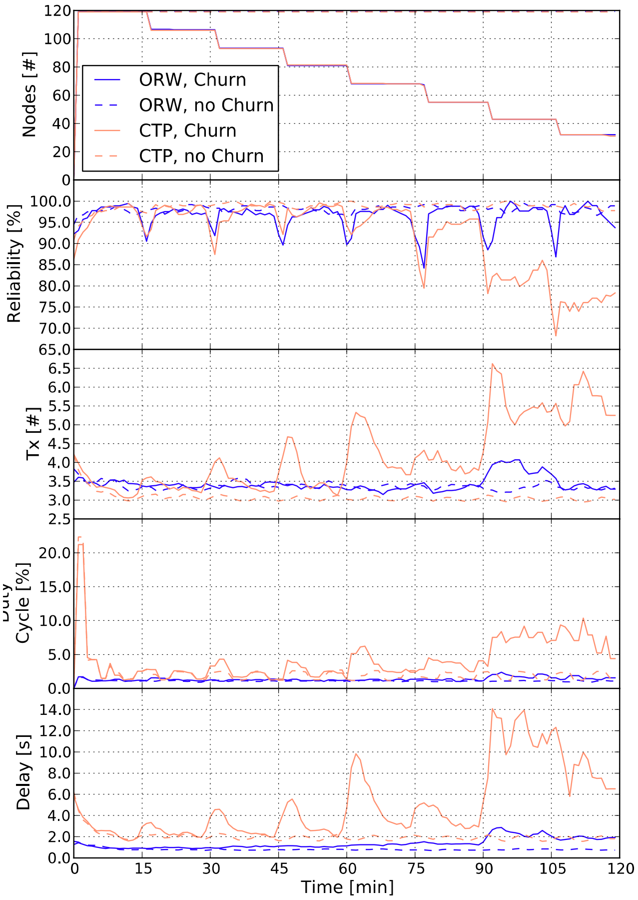

Figure 18: While both ORW and CTP achieve similar reliability under churn, CTP pays a much higher price in terms of energy, delay and transmissions.

Next, we evaluate the impact of dynamics on both ORW and CTP. For this we disable about 10 nodes every 15 minutes (Indriya testbed). Throughout the course of two hours it reduces the number of nodes from 120 to about 30.

Figure 18 depicts the impact of churn on the key metrics of reliability, transmissions, duty cycles, and delay. It shows that both protocols show spikes of reduced reliability under churn. For increased churn these spikes grow strongly for CTP. Hence, ORW maintains connectivity much longer in the resulting sparse network. Additionally, it shows the benefits of anycast routing in ORW over unicast routing in CTP: Churn has only minimal impact on the duty cycle, delay, or transmission of ORW, while these show sharp increases after churn in CTP.

5 Initial work by UNISI on distributed coding

The initial work on distributed information processing performed by UNISI considered distributed radio communication systems where two or more nodes need to transmit to a common remote destination [ABF09,AFMP09]. This model applies to many scenarios, such as, for example, cellular networks, wireless local area networks with one access point (AP), ad-hoc wireless networks, wireless sensor networks, etc. In these scenarios, collaboration between the nodes might bring significant advantages, leading to the so-called *collaborative diversity*. In a cooperative system, each user is assigned one or more partners. The partners overhear each other, process the received signals, and retransmit proper messages to the destination in order to provide extra information, with respect to the signal sent by a single source, to the AP. Even in the presence of noisy inter-partner channels, the virtual transmit-antenna array formed by cooperating nodes provides additional diversity and may improve the system performance in terms of error rate and throughput.

In the literature, many schemes have been proposed to exploit collaborative diversity. These schemes differ especially for the used relaying technique, i.e., on the basis of the information which is retransmitted by cooperating nodes to guarantee the highest ratio between diversity degree and resource consumption. The simplest schemes are those where the nodes re-transmit all the received information in an orthogonal way (typically with time division multiplexing): the used codes are not very efficient but the highest diversity is guaranteed. In other schemes, only a concise version of the information received by a cooperating node is transmitted, e.g., a parity bit. Finally, there are schemes where the nodes simultaneously access the shared radio medium—typically modeled as Gaussian multiple access channel (GMAC)—with Alamouti-like space time coding. In this case, the first direct transmission corresponds to the first row of the Alamouti code matrix (which, nevertheless, corresponds to transmissions at different moments, since the nodes cannot transmit and receive at the same time), while the simultaneous transmissions are associated with the second row. In the latter case, the nodes have only to transmit, so that the transmissions can be simultaneous. A scheme of this type allows a much higher efficiency of the previous schemes, since the multiple access interference is completely eliminated owing to the orthogonality of the Alamouti matrix—obviously, perfect synchronization between the nodes is required. In classical cooperation scenarios, therefore, the idea is that of making the nodes cooperate among themselves to implement a distributed channel coding scheme where different nodes retransmit, in some sense, the same information.

In many application scenarios, however, the information which resides in different nodes is *intrinsically correlated*. In other words, even without implementing any cooperation among the nodes, the same or, more generally, "similar" information is transmitted by the nodes. A significant application example where this situation typically appears is given by *wireless sensor networks*. The design of efficient transmission of correlated signals, observed at different nodes, to one or more collectors is one of the main design challenges in these networks. In the case of one collector node, this system model is often referred to as reach-back channel. In its simplest, this problem can be summarized as follows: two independent nodes have to transmit correlated sensed data to a collector node by using the minimum possible energy, i.e., by exploiting in some way implicit correlation among the data. In the case of orthogonal additive white Gaussian noise (AWGN) channels, a possible solution to exploit the correlation is based on joint source channel coding (JSCC) schemes, where no cooperation among nodes is required and the correlated sources are not source encoded but only channel encoded. The absence of direct cooperation between the source nodes is attractive in scenarios where the communication links between the source nodes may be noisy. This approach has attracted the attention of several researchers in the recent past, also because of its implementation simplicity. Note that, in the JSCC approach, the sources are encoded independently of each other (i.e., for a given source neither the realization from the other source nor the correlation model are available at the encoder) and transmitted through the channel. Correlation between the sources must be instead assumed to be known at the (common) receiver.

DII's activity in this filed was firstly aimed at presenting a simple perspective on cooperative coding strategies for distributed radio system. For ease of derivation, we have introduced a simple reference scenario (two source nodes and a common AP) which allow evaluating analytically the performance of various cooperative transmission schemes in the presence or absence of explicit cooperation. In all cases, we consider the presence of block faded channels and power control—under the constraint of maximum transmit power. In particular, we have analyzed various schemes, where the source correlation is exploited at the source nodes and/or at the AP. Considering a JSCC scheme in the case of orthogonal multiple access, we have introduced the concept of *correlation-induced* diversity gain, to be compared with a *cooperation-induced* diversity gain. Our results show that, in many cases, the presence of correlation between sources limits the necessity of explicit cooperation between them. The second goal of the research activity was that of deriving some "practical" example of low-density parity check (LDPC) coded JSCC scheme that can be used in the scenario of interest. The performance of such codes has been tested by means of computer aided simulations. Finally, we have developed an analytical framework to design optimal LDPC and JSCC coders, capable of exploiting data correlations in an efficient way, thus approaching the theoretical Shannon bounds**.**

6 Summary

This deliverable has summarized the efforts within WIDE to develop protocols for monitoring, estimation and control. As evidenced by the research in WP4, latency, loss and energy consumption in the network are the key parameters needed for a high-performing wireless control loop with long battery life time. Hence, our work has focused on developing reliable and low-latency protocols with good energy performance for a number of network scenarios that we believe are particularly relevant for networked control.

On one extreme, we have developed theory and methodology for maximizing the probability of meeting a hard per-packet deadline in *Wireless*HART networks. Our algorithms assume full knowledge about network topology and channel statistics and, in the version presented in this report, need to be executed centrally to compute the global transmission schedule.

This work has been complemented by the development of protocols that operate without central management, still yield high reliability, low latency and good energy performance. One effort has focused on improving contention-resolution mechanisms for correlated traffic in a single-hop scenario. We have analyzed the StrawMAN mechanism and proposed and implemented changes that gives strong performance improvements over the basic design. A particular effort has been devoted to understanding the impact of hidden terminals and the capture effect in this scenario. Another effort has been devoted to the development of opportunistic forwarding protocols for wireless sensor networks. While opportunistic forwarding protocols have been proposed for other wireless technologies, such as WiFi networks before, we use a new delay-centric metric and believe that our solution is the first one targeting and meeting the stringent energy requirements of sensor networks. The protocols has been fully implemented and evaluated in large-scale testbed deployments, demonstrating strong performance improvements over existing state-of-the art protocols such as CTP.

7 References

[Joh05] M. Johansson, "On the source-channel trade-off in networked control", IFAC World Congress, Seoul, Korea, 2005.

[SRS11] P. Suriyachai, U. Roedig and A. Scott, "A survey of MAC protocols for mission-critical applications of wireless sensor networks", IEEE Communications Surveys and Tutorials, pp 1-25, Issue 99, March 2011.

[JP07] A. Jindal and K. Psounis, "Modeling spatially correlated data in sensor networks", ACM Transactions of Sensor Networks, 2007.

[DDC10] P. Dutta, S. Dawson-Haggerty, Y. Chen, M. Liang, and A. Terzis, "Design and Evaluation of a Versatile and Efficient Receiver-Initiated Link Layer for Low-Power Wireless", In ACM SenSys, 2010.

[BYA06] M. Buettner, G. V. Yee, E. Anderson, and R. Han. "X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks", In ACM SenSys, pages 307–320, Boulder, Colorado, USA, 2006.

[ÖWT10] F. Österlind, N. Wirström, N. Tsiftes, N. Finne, T. Voigt, and A. Dunkels, "StrawMAN: Making sudden traffic surges graceful in low-power wireless networks", In ACM HotEMNETS, Killarney, Ireland, June 2010.

[PHC04] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks", In ACM SenSys, pages 95–107, 2004.

[TJB04] Y. C. Tay, K. Jamieson, and H. Balakrishnan. "Collision-Minimizing CSMA and its Applications to Wireless Sensor Networks". IEEE Journal on Selected Areas in Communications, 2004.

[HKW06] V. Handziski, A. Kpke, A. Willig, and A. Wolisz. "TWIST: A Scalable and Reconfigurable Testbed for Wireless Indoor Experiments With Sensor Networks". In ACM REALMAN, 2006.

[GSÖ11] E. Ghadimi, P. Soldati, F. Österlind, H. Zhang and M. Johansson, "Hidden Terminal-Aware Contention Resolution with an Optimal Distribution", IEEE MASS, Valencia, Spain, October 2011.

[GFJ09] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection Tree Protocol", in Proc. of SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems, 2009.

[ABF09] A. Abrardo and G. Ferrari, "A perspective on distributed radio systems with cooperative coding", URSI Radio Science Bulletin (Commission-C topics), 2009 (invited)

[AFMP09] A. Abrardo, G. Ferrari, M. Martalo and F. Perna, "Joint channel decoding with feedback power control in sensor networks with correlated sources", Proc. Second International Symposium on Wireless Communication Systems, ISWCS, Siena, Italy, 2009.