

Not All Wireless Sensor Networks Are Created Equal: A Comparative Study On Tunnels

LUCA MOTTOLA

Swedish Institute of Computer Science

GIAN PIETRO PICCO

University of Trento, Italy

MATTEO CERIOTTI

Bruno Kessler Foundation, Trento, Italy

ȘTEFAN GUNĂ

University of Trento, Italy

and

AMY L. MURPHY

Bruno Kessler Foundation, Trento, Italy

Wireless sensor networks (WSNs) are envisioned for a number of application scenarios. Nevertheless, the few in-the-field experiences typically focus on the features of a specific system, and rarely report about the characteristics of the target environment, especially w.r.t. the behavior and performance of low-power wireless communication. The TRITon project, funded by our local administration, aims to improve safety and reduce maintenance costs of road tunnels, using a WSN-based control infrastructure. The access to real tunnels within TRITon gives us the opportunity to experimentally assess the peculiarities of this environment, hitherto not investigated in the WSN field. We report about three deployments: *i*) an operational road tunnel, enabling us to assess the impact of vehicular traffic; *ii*) a non-operational tunnel, providing insights into analogous scenarios (e.g., underground mines) without vehicles; *iii*) a vineyard, serving as a baseline representative of the existing literature. Our setup, replicated in each deployment, uses mainstream WSN hardware, and popular MAC and routing protocols. We analyze and compare the deployments w.r.t. reliability, stability, and asymmetry of links, the accuracy of link quality estimators, and the impact of these aspects on MAC and routing layers. Our analysis shows that a number of criteria commonly used in the design of WSN protocols do not hold in tunnels. Therefore, our results are useful for designing networking solutions operating efficiently in similar environments.

Categories and Subject Descriptors: C.2.1 [Computer-communication Networks]: Network architecture and design—*Network topology*

General Terms: Wireless, Low-power, Network, Topology, Tunnel

Additional Key Words and Phrases: Wireless sensor networks, Network topologies, Low-power wireless communications, Topology characterization, Link quality, Tunnel environment

Corresponding author: Gian Pietro Picco, University of Trento, Department of Information Engineering and Computer Science, Via Sommarive 14, I-38050 Povo (TN), Italy. E-mail: gianpietro.picco@unitn.it. Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2001 ACM 1529-3785/2001/0700-0001 \$5.00

1. INTRODUCTION

Wireless sensor networks (WSNs) are used in a growing number of applications, where they enable dense measurements with low-cost, untethered, and flexible deployments. Nevertheless, the literature on real-world WSN deployments typically focuses on reporting about the features of a specific system, and only marginally considers its target environment. A characterization of the latter, especially for what concerns the behavior and performance of low-power wireless communication, complements the report of a successful deployment, providing results that are more general, and therefore: *i*) can be used as guidelines by developers working in the same application area; *ii*) enables comparisons across different application areas. In this paper, the design and development of a WSN-based system for monitoring and controlling *road tunnels* is the motivation for a characterization of this unusual environment, and its comparison against others that are more popular in the literature.

Motivation and research goals. TRITon (Trentino Research & Innovation for Tunnel Monitoring, triton.disi.unitn.it) is a 3-year project funded by the local administration of Trento (Italy) with the goals of reducing the management costs of road tunnels and improving their safety. One of the applications currently under development exploits WSNs to provide adaptive control of the light intensity. Densely-deployed light sensors, placed along the sides of the tunnel, measure the light intensity and report it to a controller, which performs tuning of the illumination. In contrast to traditional solutions, which typically use pre-set intensity values or at best are based on a single external sensor, this solution adapts to fine-grained light variations, both in space and time. Therefore, it is capable of *dynamically* maintaining the legislated light levels, which are instead typically statically determined according to the worst-case scenario. This enables energy savings at the tunnel entrances, where sunlight greatly affects the perceived light intensity, but is also useful inside the tunnel, where the external light does not reach. In this case, the sensor readings allow for maintenance of the light levels required by law even when lamps burn out or are obscured by dirt.

WSNs are an asset in the above scenario, as the nodes can be placed at arbitrary points along the tunnel, not only where power and networking cables can reach. This drastically reduces installation and maintenance costs, and makes WSNs particularly appealing to augment the functionality of pre-existing tunnels, where changes to the existing infrastructure should be minimized. The downside to such flexibility is the reliance on an autonomous energy source. Currently this is provided by batteries, which need to be changed periodically, e.g., when tunnel maintenance occurs. As an alternative, we are also exploring the use of techniques to harvest energy from the environment, e.g., relying on solar light at the tunnel extremities or vibrations caused by vehicles in transit. In TRITon, the WSN supporting adaptive light control is expected to become *permanent* in four existing tunnels, whose lengths range from 400 to 1,400 m. The tunnels are situated on a major freeway nearby Trento and have an average traffic of $\sim 14,000$ vehicles per day.

Other projects have envisioned the use of WSNs in road tunnels [Costa et al. 2007] or in similar environments, such as metropolitan subways [Cheekiralla 2005], underground mines [Li and Liu 2009], and service pipes [Sabata and Brossia 2005]. None of these projects, however, has yet reported a quantitative characterization of the deployment environment. Nevertheless, tunnels are a peculiar setting, bringing elements from both indoor and outdoor deployments. As in the former, the deployment area is delimited by walls and

ceiling, making multi-path effects relevant. The walls are not, however, as densely packed as in a building. Instead, they determine a regular, tube-like shape that has already been proved, although mostly analytically, to act as an “oversized waveguide” [Molina-Garcia-Pardo et al. 2009]. On the other hand, tunnels are outdoor, therefore susceptible to environmental conditions affecting connectivity. Finally, the effect of vehicles transiting the tunnel brings an additional unknown. While tunnels and similar scenarios are rapidly gaining attention as targets for WSN deployments, the lack of quantitative, real-world evidence about these environments leaves WSN developers in the dark, without specific guidelines to drive their deployments.

The access to real tunnels in TRITon gives us a unique opportunity to fill this gap. Our focus, however, is not on a theoretic or analytical characterization of the tunnel environment. The availability of real tunnels within TRITon, and our own practical need to design an efficient WSN system, pushed our goals towards a pragmatic and *system-oriented* perspective, whose ultimate objective is to understand the *cause-effect* relations between the behavior at the physical layer and the performance of commonly employed sensor networking techniques. Our objectives were therefore to:

- (1) gather and analyze *quantitative evidence* accumulated during extensive experimental campaigns using mainstream WSN technology, and
- (2) distill *practical* findings to support the design and deployment of WSN applications in tunnels.

Methodology. We analyze over 320,000 data points collected during over 1,450 experiments of duration ranging from 2 seconds to 60 minutes, for a total in excess of 1,680 hours of experiments run between Dec. 20, 2008 and Jan. 16, 2009. These were carried out across three deployments, described in Section 2: *i*) an operational road tunnel, enabling us to assess the impact of vehicular traffic; *ii*) a non-operational tunnel, providing insights into analogous scenarios (e.g., underground mines) without vehicles; *iii*) a vineyard, serving as a baseline representative of deployments commonly found in the literature. To make the results comparable at least in the aspects we can control, we replicated the node placement, ran the same suite of test applications, and used the same hardware and software in all deployments. As we describe in Section 3, we selected WSN technology that appears to be the most popular, in an effort to provide results that are of wide applicability for the research community.

In accordance to our goals, our experiments span multiple network layers, as described in Section 4. At the physical layer, we gather data about packet delivery rate and link quality estimators. At the MAC layer, we assess the packet delivery rate in the presence of packet collisions and the performance of radio duty-cycling. Finally, at the routing layer we gather data about end-to-end delivery, the shape and dynamic characteristics of a tree-based routing topology, and the performance of popular techniques used by the routing protocol being tested, e.g., the ETX [Couto et al. 2005] metric and link-level acknowledgments. We also collect data about temperature and humidity, to assess their impact on the network performance.

The entire data set collected during the experiments is publicly available [D3S Group].

Contribution. Our study shows that the tunnel environment is indeed significantly different from those targeted in the traditional WSN literature, to the point that some commonly accepted techniques, trade-offs, observations, and rules-of-thumb lose their validity.




Fig. 1. Deployment sites.

For instance, it is common wisdom that WSN links tend to be brittle and short-range when using mainstream hardware. Instead, we observed that in tunnels links are generally stable and long-range, an effect of the tunnel behaving similarly to a waveguide. In this situation, CSMA MAC protocols tend to perform inefficiently, due to a high probability of packet collision. Similarly, previous testbed experiments [Srinivasan and Levis 2006] noted that the Link Quality Indicator (LQI) provided by IEEE 802.15.4 radios does not accurately reflect the corresponding link reliability, unless the link is near-perfect or hundreds of LQI values are averaged. As a consequence, ETX and its derivatives are generally preferred for building routing trees. Our experiments show that, in a tunnel, this is only partially true: even a few LQI values yield routing trees as good as ETX-based ones, but require significantly less network traffic.

These and the other findings discussed in Section 5 show that tunnels and similar environments undermine common assumptions and, worse, make popular protocols and techniques inefficient. The findings reported in this paper, supported by extensive quantitative evidence, help to build background knowledge for developing and deploying WSNs in these scenarios. Our survey of related work in Section 6 shows that we are the first to undertake such an effort.

2. DEPLOYMENT SCENARIOS

The first deployment is an operational road tunnel (named TRAFFIC hereafter), where the presence of vehicular traffic allows us to analyze the impact of objects moving inside the tunnel. The results we derive in this context are therefore potentially applicable to analogous scenarios such as railway or subway tunnels [Cheekiralla 2005]. The second scenario is a non-operational tunnel (TUNNEL). In this case, the absence of traffic allows us to assess the impact of the tunnel shape in isolation, making our results of interest in scenarios such as underground mines [Li and Liu 2009] and service pipes [Sabata and Brossia 2005]. Our third scenario is a vineyard (VINEYARD), representative of traditional outdoor scenarios [Beckwith et al. 2004]. VINEYARD offers a comparison point for the tunnel deployments, although the experimental results we derived for it are useful to shed additional light on this often-cited scenario. Figure 1 shows our three deployment sites.

The network we used in all experiments includes 20 battery-powered WSN nodes, whose hardware and software configuration is described in Section 3. The configuration of our deployments is essentially the same across all scenarios and mirrors the one in TRAFFIC, shown in Figure 2. Nodes are placed along two parallel lines on the opposite walls and skewed so that a node on one wall is never directly opposite to another node. In TRITon, this is motivated by the need to minimize transient occlusions of the light sensors caused by voluminous vehicles such as trucks.




Fig. 2. TRAFFIC deployment. The distance between nodes is the same in all scenarios.

The network *density* is also driven by the requirements of the TRITon project. The nodes are more densely deployed at the tunnel entrance, and become more sparse inside the tunnel. Indeed, light variations are expected to be larger and more frequent at the tunnel entrance. To capture these variations, it is necessary to attain higher sensing granularity in that area. A coarser sensing resolution is instead sufficient inside the tunnel, where light variations are likely to be smaller and less frequent. On the other hand, the *distances* between adjacent nodes on the same wall are dictated by the desire to replicate the very same configuration in VINEYARD. Indeed, unlike the tunnel scenarios where we could freely place nodes along the wall, in VINEYARD our choices were limited by the pre-existing location of the vine-supporting poles.

Each node is packaged in a plastic box, shown in Figure 3, and placed so that the board, containing the integrated antenna detailed in Figure 3(a), is parallel to the tunnel wall. This allows the light sensors, which are mounted on an extension board parallel to the node itself for ease of assembling, to face directly the road.

In all deployments we used a gateway, a Gumstix (www.gumstix.com) embedded PC featuring both Ethernet and Wi-Fi interfaces, to enable remote control of the experiments and collection of the results. The gateway is connected through USB to a WSN node where we collect the results of our experiments.

In the following, we provide additional details about the chosen scenarios and the corresponding WSN deployments.

2.1 Operational Tunnel (TRAFFIC)

The TRAFFIC deployment is in the “San Vigilio” tunnel, shown in Figure 1(a), situated about 4 Km from Trento on the road SS 45bis “Gardesana”. This is a main road with a relatively high traffic load, as it provides access to the city of Trento from the West, e.g., from the nearby lake Garda. The tunnel is one-way, with two lanes. The average number




Fig. 3. WSN nodes in the field.

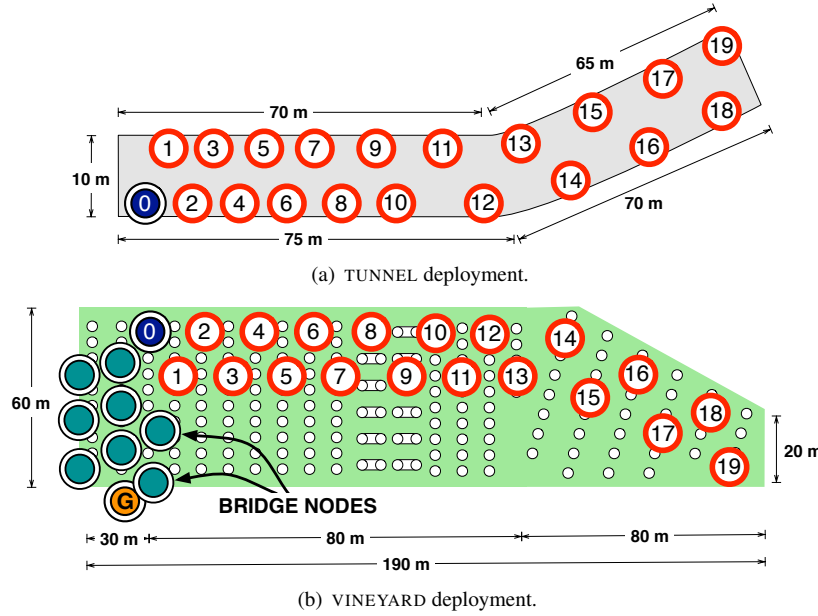


Fig. 4. Node placement in TUNNEL and VINEYARD.

of vehicles per hour varies between a few tens during the night to about a thousand during peak hours, e.g., in the morning. The traffic includes heavyweight vehicles such as trucks, which account for 2% of night traffic and up to 10% of morning traffic. The tunnel is carved into a hill. It is 400 m long, 8 m high, and the distance between walls is 10 m. In the 150 m where we placed our nodes, the tunnel is straight.

Deployment details. The nodes are installed approximately 1.70 m from the ground, the height at which light levels must be monitored in the TRITon project for the control loop to be compliant with the current regulations. Nodes are packaged in a $14 \times 17 \times 9.5$ cm³ box that, according to regulations, must be fire-proof and water-proof (IP65 or higher). Each node was equipped with four size D Duracell batteries. The gateway is connected to the internal network already present at the entrance of the tunnel, which relays data through the Internet to a remote control center.

2.2 Non-operational Tunnel (TUNNEL)

Within the TRITon project, we also had the opportunity to test our system in a non-operational tunnel called “Doss Trento”, shown in Figure 1(b), situated in the northern urban area of Trento. As it is no longer in operation, the TUNNEL deployment was much more accessible than the TRAFFIC one, where every change to the placement of nodes or to the software running on them required a number of authorizations from the local administration to enforce partial or total interruption of traffic and the execution of the required modifications by the company in charge of the tunnel maintenance. Instead, in TUNNEL we could access the site as needed, and perform the required tasks ourselves.

Unlike TRAFFIC, TUNNEL bends about 70 m from the entry, as shown in Figure 4(a) (not to scale). Although the tunnel is 290 m long, we were granted access only to the first 150 m. Apart from these details, the other values are similar to those in TRAFFIC. The tunnel is old and its structure is not completely impermeable: on rainy days, water was

seen to penetrate in the area around the gateway.

Deployment details. Compared to TRAFFIC, the only difference in TUNNEL, as seen by comparing Figure 2 and 4(a), is the mirrored placement of nodes, i.e., the walls on which the nodes with even and odd identifiers are placed are switched. This is motivated by the availability of the power outlets, necessary for the gateway and for a radio bridge installed at the entry. The latter is necessary to ensure connectivity to our institute because TUNNEL, unlike TRAFFIC, was not connected to any network infrastructure.

2.3 Vineyard (VINEYARD)

This scenario was made available by the Edmund Mach Foundation, in San Michele all'Adige, nearby Trento. The area available to us, shown in Figure 1(c), is located on a slope, with a length of 190 m and a width ranging from 20 to 60 m. The area can be divided in two sectors: a first one with a rectangular shape (approx. $110 \times 60 \text{ m}^2$) and a second that progressively narrows over a length of 80 m. Figure 4(b) provides a graphical representation (not to scale). In the second sector the position of the poles is fairly regular: 6 m between rows and 3 m between the poles inside each row. Instead, in the first sector the distance between rows varies between 1.3 and 3.9 m, and even the type of poles is different, with a mixture of vertical and v-shaped poles. As mentioned earlier in this section, these constraints determined the node distances we chose for all deployments.

Deployment details. The vineyard is far from the closest building where Internet connection and electricity are available. Therefore, it was not possible to co-locate the gateway with the nodes in the field. As shown in Figure 4(b), we placed the gateway in reach of Internet access and ensured connectivity between the root and the gateway through 8 additional WSN “bridge” nodes. To avoid affecting the experiments, these operate only outside of the experiment times. During each experiment, the root node buffers on external storage the information that would normally be forwarded to the gateway. The actual transmission to the gateway through the bridge nodes takes places when the experiment is over.

Each WSN node was packaged in a $13 \times 13 \times 7.5 \text{ cm}^3$ box, shown in Figure 3(b). The boxes were placed at the top of the vine-supporting poles, i.e., at about 1.5 m from the ground. Unlike the other two deployments, we equipped each node with only one pair of Size D Duracell batteries, as the poles could not support a heavier weight. It is worth noting that since we performed most of our tests in winter, the vineyard had no leaves.

3. WIRELESS SENSOR NETWORK TECHNOLOGY

The research goals we stated in Section 1 demand studying different layers in the stack, namely the physical (PHY), MAC (MAC), and routing (RTN) ones. The behavior of the physical layer is determined by the communication hardware and the characteristics of the deployment, and is therefore essential to understand the peculiarities of the tunnel scenario. In contrast, the mechanisms and protocols at the MAC and routing layers, affected by the physical layer, bear a direct impact on the behavior of WSN applications. As such, the study of these two layers is fundamental in understanding the performance that can be obtained in our target scenarios.

Looking at the landscape of *deployed* WSN systems, it appears that most of them leverage a limited set of simple networking techniques [Raman and Chebroly 2008]. Here, we employ the same solutions to assess the effectiveness of these popular mechanisms. A similar rationale motivated our choices of hardware and system support. For the former, we

use nodes functionally equivalent to the the widespread TMote Sky [Polastre et al. 2005]. These are equipped with a SHT11 temperature/humidity sensor, which in our experiments is used to study how the environmental conditions affect our results. We avoided the use of external antennas, as these would significantly increase the overall deployment costs and also make our results a function of the specific antenna employed [Raman et al. 2006]. In retrospect, this turned out to be a wise decision given that, as we show in Section 5, the on-board antenna already enables a range up to 150 m. This is sufficient for our application, where the distances among sensors are constrained by the position relative to the lamps, and is at most 50 m. For system support we use TinyOS [Hill et al. 2000], as this allows us to re-use implementations of popular networking protocols.

In the following, we provide a concise overview of the specific hw/sw mechanisms we tested at the different layers.

Physical layer (PHY). We use the IEEE 802.15.4-compliant ChipCon 2420 [Chipcon Tech.] as the radio chip and the on-board inverted-F micro-strip omnidirectional antenna. The 802.15.4 standard prescribes a direct sequence spread spectrum OQPSK modulation across 16 channels in the ISM 2.4 GHz band, yielding a nominal bandwidth of 250 Kbit/s. For each packet received, the CC2420 computes two indicators: the *received signal strength indicator* (RSSI) and the *link quality indicator* (LQI). RSSI measures the received signal strength, while LQI quantifies the correlation between a received symbol and the symbol this is mapped to after the radio completes decoding [Chipcon Tech.]. Intuitively, the higher this value, the “cleaner” the channel. Optionally, the CC2420 provides hardware-generated acknowledgments for reliable delivery of unicast packets.

MAC layer (MAC). The default MAC layer used in TinyOS is a Carrier-Sense Multiple-Access (CSMA) protocol based on clear channel assessments (CCA) and random back-offs for channel arbitration [TinyOS d]. On top of it, an asynchronous low-power listening (LPL) scheme [TinyOS a] is available to save energy by duty-cycling the radio. When LPL is employed, every node turns on the radio periodically (based on a *sleep interval SI*) to perform a receive check, and turns it off immediately after if no transmission is detected. Senders retransmit the outgoing packet for a duration of twice the *SI* of receivers, therefore ensuring that the latter eventually detect the transmission. Unicast transmissions are acknowledged to allow the sender to stop its transmission when the packet is received. The value of *SI* is set by the application, to balance the trade-off between energy consumption and network throughput.

Routing layer (RTN). The Collection Tree Protocol (CTP) [Gnawali et al. 2009; TinyOS c] is a representative of commonly employed tree-based routing schemes. At a high-level, the behavior of most of these protocols is characterized by two core elements: *i*) a routing metric used by every device to select a parent node in the tree; and *ii*) a set of reliability mechanisms to improve the fraction of messages correctly delivered to the root of the tree.

The version of CTP we used leverages the *expected transmissions count* (ETX) [Couto et al. 2005] as a basis for its routing metric. This metric is computed based on beacons broadcast by every node. In CTP, nodes select their parent in the tree by minimizing the *end-to-end* ETX, i.e., the number of retransmissions required to deliver a message to the root, possibly across multiple hops. Reliability is achieved mainly by using randomly-scattered retransmissions and link-layer acknowledgements. In CTP, the latter is implemented entirely in software, instead of using the CC2420 hardware-level acknowledgments, to retain better control of retransmission timeouts.




Fig. 5. Scheduling transmissions in PHY and MAC experiments.

4. EXPERIMENTS: GOALS AND SETUP

In each deployment we selected the radio channel to avoid interference caused by nearby devices transmitting in the same band, e.g., WiFi or Bluetooth. We also monitored the battery voltage before and after every experiment to verify that the battery never went below a threshold affecting the performance of the radio transceiver [Chipcon Tech.]. In all experiments, the packet size was 105 bytes inclusive of MAC header and 94 bytes of payload, representative of common WSN applications [Mainwaring et al. 2002]. Following the manufacturer's suggestions, we set the CC2420 output power to -1 dBm, one level less than maximum power. This is to alleviate distortion effects due to miscalibrated hardware or imperfections in the manufacturing process of the radio chips [Zuniga and Krishnamachari 2007].

4.1 Physical Layer (PHY)

Goal and metrics. Our objective is to investigate the performance of the bare-bones physical layer, i.e., in the absence of factors such as collisions and retransmissions that are normally dealt with by higher layers of the network stack. To this end, each node sends periodic broadcast messages and tracks messages received from other nodes. In each experiment, we input the duration of the test (DT), and compute the following quantities as output:

- the *packet delivery ratio* ($PDR_{i \rightarrow j}$) between any pair of nodes i, j , that is, the ratio of packets received by node j over the number of packets sent by node i or, equivalently, the *packet error ratio* $PER_{i \rightarrow j} = 100\% - PDR_{i \rightarrow j}$;
- the average RSSI (\overline{RSSI}) and LQI (\overline{LQI}) between any pair of nodes, computed over the packets received;
- the average *temperature* and *relative humidity* observed during the test.

PDR , \overline{RSSI} , and \overline{LQI} are fundamental to investigate the characteristics of network links and the accuracy of link quality estimators, as illustrated in Section 5, where we also assess how these are affected by environmental conditions.

Implementation. To avoid collisions, periodic transmissions are scattered over time as illustrated in Figure 5(a). For the entire duration of an experiment, each node broadcasts a packet every $N \times \delta$ time instants, where N is the number of participating nodes and δ is a known time interval. The transmissions are scattered at the beginning of each experiment by a time offset, different for each node, computed by multiplying the node identifier by δ . To minimize clock drifts, the nodes are time synchronized at the beginning of each experiment by using a form of reference broadcast [Elson et al. 2002], where the reference

message includes the duration of the experiment DT . Unless otherwise specified, we use $N=20$ and $\delta=300$ ms: the latter is within the accuracy of our time synchronization and of the back-off timers employed by the radio transceiver. At the end of each experiment, all nodes transmit their results to the gateway using network-wide flooding, for maximum reliability. The nodes can also be configured to dump to flash memory the packets they received during the experiment, and later communicate these to the gateway along with the aggregated results. Although this drastically increases the time required to perform an experiment, it allows us to study the variability of the above metrics also across packets.

4.2 MAC Layer (MAC)

Goal and metrics. To investigate the impact of collisions on the behavior of the TinyOS MAC layer—especially w.r.t. the effectiveness of the CCA functionality—we devised experiments where we control the *probability of concurrent transmission* P_{ct} , i.e., the probability that any two nodes broadcast a packet concurrently. The value of P_{ct} , along with the number N of participating nodes, the test duration DT and, if LPL is used, the sleep interval SI , are given as input to the experiments. As output, we collect the same quantities of PHY tests. If LPL is used, we also compute the number of *radio activations* (R_{act}) at each node. Here, however, \overline{RSSI} , \overline{LQI} , and the environmental readings are used mostly to “match” PHY and MAC experiments. This allows us to investigate how the behavior of the physical layer impacts the MAC performance under similar conditions, as discussed in Section 5.

Implementation. At the beginning of each experiment, and based on the number N of participating nodes, the desired probability P_{ct} and the maximum duration of a packet transmission, the gateway computes a time interval ρ that nodes use to proceed in rounds after a time-synchronization phase similar to PHY experiments, as shown in Figure 5(b). Within ρ , a node decides randomly when to send a packet. Therefore, a shorter ρ corresponds to a higher probability of concurrent transmissions. The value of ρ is communicated to the nodes as part of the reference broadcast message, along with DT and SI (if any), whereas the results of the experiment are collected as in PHY tests.

4.3 Routing Layer (RTN)

Goal and metrics. To study how the tunnel environment affects the performance of the routing layer, we developed a dummy data collection application atop CTP, where each node sends messages with a given inter-message interval (IMI), for an overall duration DT . The experiments return:

- the *total message delivery ratio* at the CTP root (node 0 in every deployment), i.e., the number of application messages received over the total number of application messages generated;
- the number of *duplicates* dropped inside the network (DUP);
- the *parent in the tree* at every node at the time of sending a message, along with the total number of *parent changes* (PT) observed throughout the experiment;
- the number of failed *link-layer acknowledgments* at every node (ACK_{failed});
- the number of *beacons* transmitted at each node to compute ETX.

The metrics above are used to study both the overall performance of CTP, as in the case of message delivery ratio, and to assess the effectiveness of specific mechanisms. For

Observation at the physical layer	Section	Effects and implications
Links are long-range and stable, grey areas and noise from far transmitters are relevant.	5.1	[MAC] CSMA has poor reliability and high energy consumption, TDMA-like techniques are more suited.
LQI is a good indicator of link quality, RSSI is not except for very good links.	5.2	[Routing] LQI can be used in place of the more costly ETX, as they build similar routes.
Link asymmetries are mostly permanent and related to the position of nodes.	5.3	[Routing] Link-layer acknowledgments work poorly, but pathological links could be identified a priori.
Vehicular traffic affects the physical layer with temporary link disruptions.	5.4	[Routing] Need specific reliability mechanisms to avoid high resource consumption.

Table I. Summary of the main findings about the tunnel environment.

instance, PT and the number of ETX beacons are useful to investigate the behavior of ETX-based link quality estimates.

Implementation. We instrumented the CTP implementation to provide our test application with the necessary hooks to gather the above statistics. We embed this information within the application messages to report them to the gateway. The tests are started as in the PHY case, but this time the generation of messages is randomly scattered over time, as in most data collection applications [Mainwaring et al. 2002].

5. RESULTS AND DISCUSSION

The ultimate objective of our experiments is to understand if and how the characteristics of the physical layer in the tunnel environment affect the performance of MAC and routing protocols. To this end, instead of discussing the results separately at each layer, we draw first some observations on the behavior of the physical layer based on the results of PHY tests. For every such finding, we investigate its impact on the performance of the higher layers by analyzing data obtained from MAC and RTN tests in comparable conditions. This process is the stepping stone to derive a set of guidelines for the design of networking techniques suited to the scenarios we consider.

Table I summarizes the observations we draw at the physical layer and their implications on MAC and routing protocols, while also providing a road-map for the rest of this section. First, however, we discuss the impact of environmental conditions on the physical layer, which provides the rationale to match PHY tests to MAC and RTN experiments.

Impact of environmental conditions. We studied the impact of environmental conditions on the network performance using the on-board temperature and humidity sensors. To allow the measurement of these parameters, we pierced four holes on the bottom of the (otherwise air-tight and water-proof) boxes in which nodes were packaged. Based on the results of PHY experiments, we can state that:

- (1) in both tunnels, humidity varies the most at the entrance, whereas the inner portions are barely affected by changes in the outside environmental conditions;
- (2) throughout our study, temperature did not affect the performance of the physical layer, while humidity had an impact only above a given threshold.

Figure 6(a) provides quantitative support for the first observation, by plotting the average relative humidity and its standard deviation at some sample nodes, recorded in both tunnels at the same time. The plot shows that the excursion in both cases is significantly higher at nodes closer to the entrance. Similar trends are observed also for temperature. In a sense,

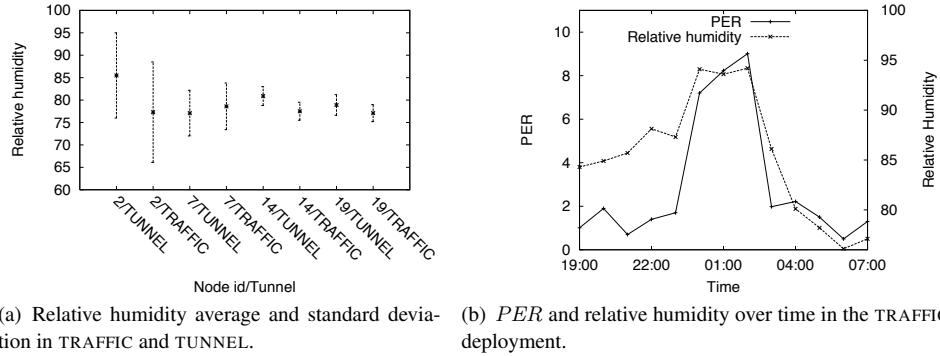


Fig. 6. Impact of humidity.

the tunnel acts as a “low-pass filter” for humidity and temperature, preventing short-term variations from propagating inside the tunnel. This confirms previous findings in the field of tunnel engineering [Sun et al. 2008] and similar phenomena in analogous contexts such as mine galleries [Nowak et al. 1997].

To motivate the second consideration, we consider a sample set of 192 PHY tests with duration $DT=30$ min evenly distributed over 4 days. Figure 6(b) shows the relative humidity and PER over time between node 1 and 2 in TRAFFIC¹. Figure 6(b) shows three distinct behaviors. Until 23:00, humidity was around 85% with $PER \approx 1\%$. During the night, humidity reached 93% and the PER increased abruptly to 9%. Later on, when humidity decreased to about 75%, the link performed similarly to when humidity was around 85%. Therefore, humidity appears to affect the PER only when over a threshold of about 90%. We observed the same trends for almost every pair of nodes in all deployments. We also noticed that, unlike humidity, temperature is not relevant, at least in the period of our experiments. To some extent, this is expected because temperature should affect the performance only when outside the hardware operating range ($[-40^\circ C, +60^\circ C]$ in our case), whereas we always observed average temperatures in the $[-9^\circ C, +11^\circ C]$ range.

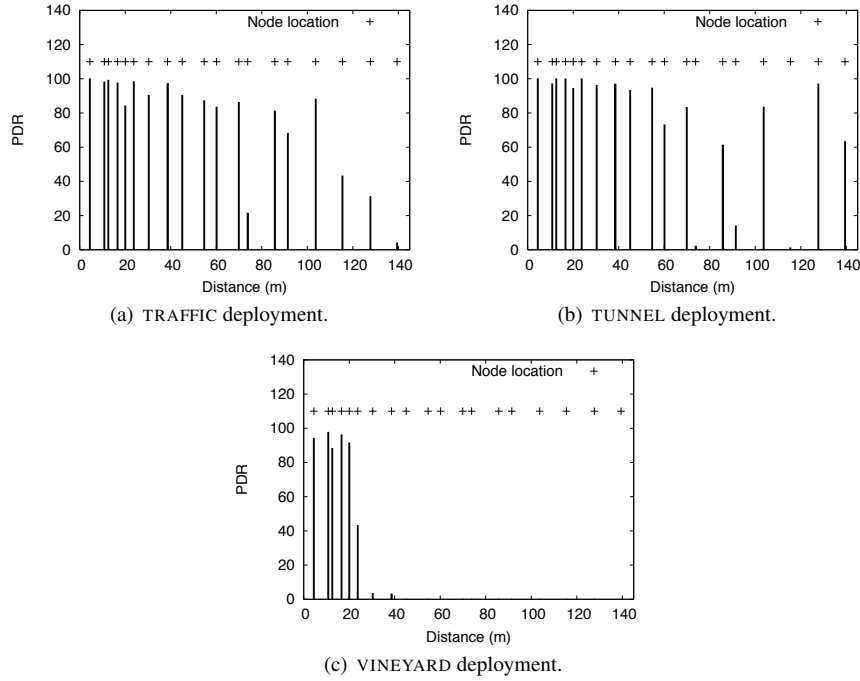
Based on the results we just described, the considerations presented next are derived on tests run under comparable environmental conditions.

5.1 Spatio-Temporal Characteristics of Links

Physical layer. To characterize the network links, we ran a total of 650 PHY tests in each deployment with varying DT . For now, we ignore periods with significant vehicular traffic in TRAFFIC, deferring the corresponding discussion to Section 5.4. By comparing the results across the three deployments, we draw the following observations:

- (1) the communication range measured in VINEYARD confirms previous results [Zhao and Govindan 2003; Srinivasan et al. 2010]; in contrast, in both tunnels the communication range is much greater than in VINEYARD;
- (2) similarly, the extent of the “grey areas” in VINEYARD matches known results [Zhao and Govindan 2003; Srinivasan et al. 2010]; instead, in both tunnels grey areas span a



¹The chart was obtained from tests run on non-working days. According to the data provided by the local authorities, during this time span the vehicular traffic inside TRAFFIC is very light. Thus, it does not constitute a relevant factor in these experiments.

Fig. 7. *PDR* against distance from node 0.

- significant fraction of the maximum communication range;
- (3) in contrast to most WSN deployments reported in the literature, the link performance in both tunnels is quite stable over time.

The first observation is supported by the *PDR* data we obtained from PHY tests. As an example, we consider the average *PDR* against distance, for packets sent by node 0 over 200 PHY experiments with $DT=10$ min, although the results we obtained for different DT are essentially the same. As Figure 7 illustrates, in both tunnels nodes more than 100 m apart have a *PDR* as high as 94%. Figure 7(b) shows how this phenomenon is even more evident in TUNNEL, where the link between node 0 and 18 has $PDR \approx 96.67\%$, over a distance of about 120 m. Conversely, in VINEYARD the *PDR* drops abruptly with distance: nodes barely communicate when more than 30 m apart.

Figure 8 offers an aggregate view on the behavior of *PDR*, based on all 650 PHY experiments. The trends show that in VINEYARD communication tends to be localized, while a relevant fraction of the transmitter-receiver pairs in both tunnels enjoys good reliability. The complementary cumulative distribution function (CDF) for *PDR* shows that about 60% of the transmitter-receiver pairs in the two tunnels observe a *PDR* higher than 90%, while in VINEYARD this percentage drops to only 25%. We conjecture that these trends are due to the peculiar shape of the tunnel environment and the corresponding waveguide effect, which yields significant multi-path effects. This may create constructive interferences resulting in extended communication range w.r.t. outdoor environments [Molina-Garcia-Pardo et al. 2009]. Although our results are in accordance to known propagation models for tunnels, to the best of our knowledge this is the first time that quantitative results obtained

Fig. 8. Complementary CDF for PDR .Fig. 9. Average PDR difference.

in a real-world setting are reported.

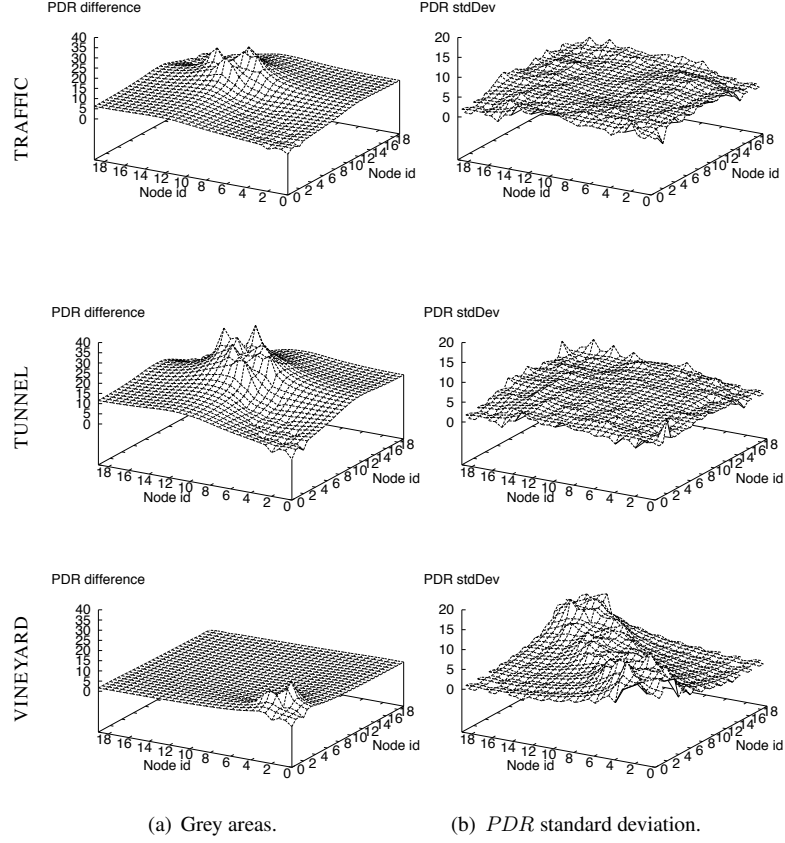

On the other hand, while constructive interference may appear in some areas, others may experience destructive interference. This phenomenon is confirmed in our PHY experiments. Indeed, we observed the existence of large “grey areas”, where physically close nodes experience large differences in PDR for packets sent by the same transmitter, as already shown, for instance, in Figure 7. Nevertheless, for a deeper understanding, we consider again transmissions from node 0 and, for each receiver i , we compute the absolute difference between $PDR_{0 \rightarrow i}$ and $PDR_{0 \rightarrow j}$, where j ranges among *each* of the four physically closest neighbors of i . The results are illustrated in Figure 9 and 10(a)². The former shows the average value over the four aforementioned neighbors. This value is markedly higher in both TRAFFIC and TUNNEL compared to VINEYARD, showing that grey areas are more severe in tunnel environments. Figure 10(a) shows an alternative view that makes the spatial extent of grey areas more evident. In both tunnels, these are larger than in VINEYARD, demonstrating that constructive/destructive interference has a greater impact in tunnels.

The temporal characteristics of packet reception in the two tunnels are also peculiar. We investigated this aspect both across and within experiments. As for the former, we ran 330 PHY tests with $DT=1$ min, randomly scattered throughout the entire time span of our study, to “sample” the links at different points in time. Figure 10(b) depicts the standard deviation of PDR between any transmitter-receiver pair across all these experiments. The performance of wireless links varies largely in VINEYARD, where the standard deviation of PDR is significant only for physically close nodes, as farther ones often do not communicate at all. On the contrary, wireless links in the two tunnels are quite stable, although some small changes are observed for almost every link. To the best of our knowledge, similar behaviors where the link performance remains stable over large time scales are rarely observed in real-world settings.

To study the link behavior in the tunnels within individual experiments, we took advantage of our access to TUNNEL and ran about 200 ad-hoc tests with only one node transmitting, using $DT=2$ s and varying δ , from 50 ms to 250 ms. This time, the other nodes logged every received packet on flash memory³. Based on this information, we

²In Figure 10(a) and 10(b), data points along the diagonal of the base plane are excluded from the analysis.

³The absence of a wired back-channel in TUNNEL, and the consequent use of flash memory to log received packets, prevented using $\delta < 50$ ms.

Fig. 10. Spatial and temporal characteristics of *PDR*.Fig. 11. Conditional packet delivery function (CPDF) in TUNNEL for a link with average *PDR* = 94.73%. Negative values on the *x*-axis indicate previously missed packets.

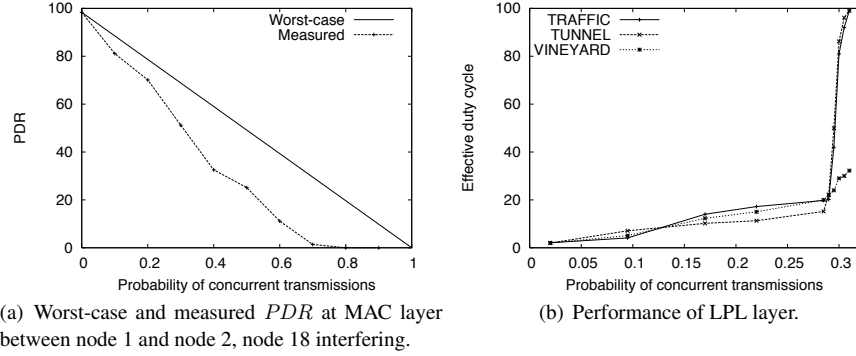


Fig. 12. Performance of a CSMA MAC protocol in the tunnel environment.

computed the Conditional Packet Delivery Function (CPDF) as done by Srinivasan et al. [2010]. This describes the probability that the current packet is received, based on whether the last k packets were received or missed. As shown by two sample charts in Figure 11, given that the behavior of wireless links in TUNNEL is either very good or very poor, most of the CPDFs turn out to be flat and approaching the average PDR on the same link. This indicates that losses are independent, confirming previous results [Srinivasan et al. 2010].

These observations have deep implications on the performance of common WSN MAC protocols, as described next.

Implications on MAC. To test the effectiveness of a CSMA MAC layer such as the one in TinyOS, we ran about 200 MAC tests with varying DT and P_{ct} . Half of the tests were conducted using LPL with different SI .

Not surprisingly, in a very connected scenario like the two tunnels, concurrent transmissions are very likely to generate collisions and thus packet losses. In quantitative terms, when P_{ct} approaches 40% almost no packet is correctly received. This performance was surprisingly bad, as one would expect the CCA mechanism to alleviate the risk of collisions, at least to some extent. To investigate this aspect further, we went back to TUNNEL and carried out some experiments with only some of the nodes, identifying some pathological configurations. Consider for instance a setting where only nodes 1, 2, and 18 are operating. Node 18 is very far from both node 1 and 2, as shown in Figure 4(a), and has no network-level visibility to either of them in either direction, i.e., $PDR_{1 \rightarrow 18} = PDR_{18 \rightarrow 1} = 0$ (similarly for node 2). In this configuration, one would expect node 18 to affect communication between 1 and 2 only to a very limited extent. Therefore, in the *worst case*, assuming that the CCA mechanism constantly fails and therefore each concurrent transmission results in a packet loss, the PDR between 1 and 2 in MAC experiments should obey the expression:

$$PDR_{worst-case:1 \rightarrow 2} = PDR_{1 \rightarrow 2} * (1 - P_{ct})$$

Nevertheless, Figure 12(a) indicates that this is not the case. The PDR measured with our MAC tests was lower than the worst-case, suggesting that node 18 did have significant impact even without network-level visibility. We conjecture that the additional noise caused by 18 was responsible for this behavior. Specifically, we maintain that the communication hardware at node 1 or 2 was not accurate enough to recognize the on-going transmission

from a far node, yet the noise did corrupt the ongoing transmissions. We verified that the observed behavior was not caused by malfunctioning hardware, by replacing node 1 and 2 with other two nodes. On the other hand, when we repeated the same experiment without node 18, the measured $PDR_{1 \rightarrow 2}$ was on average 61% higher than the expected worst-case.

In addition to poor reliability, the phenomena above also cause relevant energy consumption when asynchronous LPL layers are used, as in our MAC tests. Transmissions from very far nodes induce frequent attempts to receive messages that, however, turn out to be corrupted as they fail CRC checks. As these protocols usually employ some form of “delay after receiving”⁴, the radio may end up staying on much longer than necessary. To quantify this aspect, we compute the *effective duty-cycle*, i.e., the fraction of time the radio is on over the entire test duration DT , based on the number R_{act} of radio activations during the same test. Figure 12(b) depicts the corresponding trends against different P_{ct} , when LPL operates with $SI=250$ ms. When P_{ct} goes above 30%, in both tunnels the effective duty-cycle increases abruptly and the radio remains on most of the time. This behavior is not observed in VINEYARD, where the long-range noise effect is absent. Presumably, this time the communication hardware failed in the opposite way w.r.t. the MAC experiments above, recognizing an attempt to transmit when the transmitted power was too low to guarantee a successful reception.

Tunnel-like environments have thus two deep implications on the behavior of CSMA-like MAC protocols for WSNs: *i)* collisions and noise due to concurrent packet transmissions have a great impact on *reliability* even without link-level visibility of the nodes involved, and *ii)* asynchronous LPL techniques suffer from frequent failed attempts at receiving packets, yielding bad radio duty-cycling performance. Note how these findings implicitly suggest the use of TDMA schemes (e.g., [Rajendran et al. 2006]), as most of the issues above could be solved by controlling the transmission schedules. The temporal stability of links would also reduce the need for reconfiguring the schedules too frequently. However, the use of TDMA MAC protocols would come at the extra cost of time synchronization. Unfortunately, at the time we conducted our tests we could not find stable implementations of TDMA MAC protocols for our platform to verify the argument. Nevertheless, the results we presented can be an asset to MAC designers in their investigation and further motivation to refine the existing TDMA implementations.

5.2 Performance of Link Quality Estimators

Physical layer. A number of approaches to evaluate the quality of network links, i.e., their expected delivery performance, rely on information gathered at the physical layer, such as RSSI and LQI. To investigate their performance, we refer to PHY tests with $DT=1$ min as in Section 5.1, and look at the corresponding \overline{LQI} and \overline{RSSI} . Our results suggest that:

- (1) in all deployments, RSSI does not reflect accurately the link quality unless its value crosses a threshold; in this case, however, it is useful to identify only very good links;
- (2) contrary to most reported experiences and our own observations in VINEYARD, in both tunnels a small number of LQI measurements suffice to identify also links of intermediate quality.

To support the first observation, we consider about 125,000 data points $\langle \overline{RSSI}, PDR \rangle$ obtained from the aforementioned PHY tests where each node sends 10 packets with $DT=1$

⁴With this technique, MAC protocols keep the radio on after a receive attempt to favor packet bursts [TinyOS a].

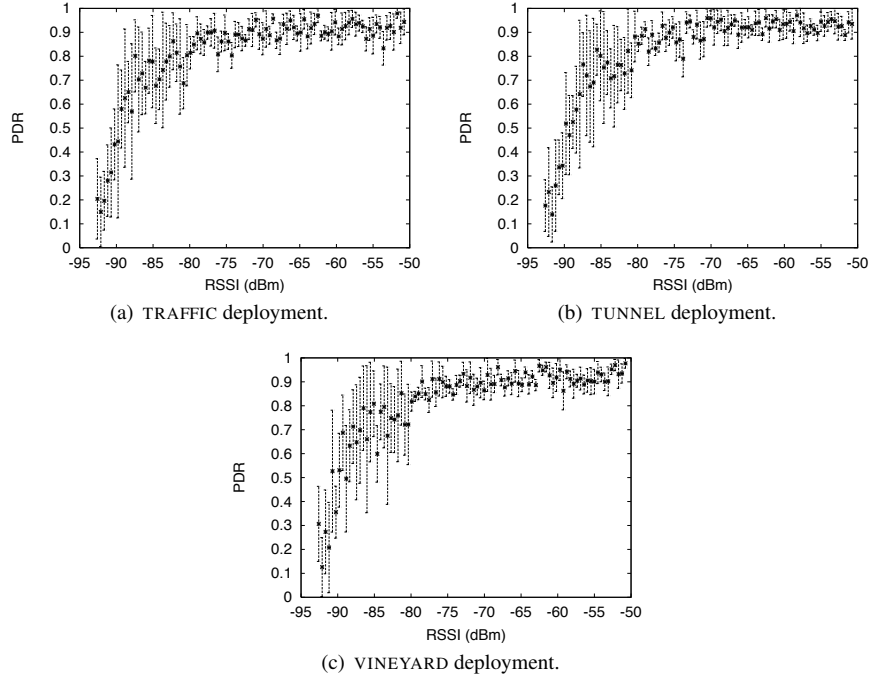


Fig. 13. Relation between $\lfloor \overline{RSSI} \rfloor$ and average PDR with standard deviation.

min. We round the former to the nearest integer $\lfloor \overline{RSSI} \rfloor$ and compute the average and standard deviation of all PDR values with the same $\lfloor \overline{RSSI} \rfloor$. Figure 13 plots the results of this analysis. In all deployments, values of RSSI over -81 dBm exhibit a sharp reduction in standard deviation, and usually indicate a good link, i.e., $PDR > 80\%$. However, they do not allow one to discern among different specific values of PDR in this range. Below -81 dBm, instead, the standard deviation of RSSI readings is larger and the same RSSI may be observed for good and bad links, which are therefore indistinguishable based solely on this metric. We conjecture that the threshold of -81 dBm we observed is due to a combined effect of the radio sensitivity threshold and of the environmental noise floor characteristic of our tunnels. A better understanding of this phenomenon, however, would require dedicated instrumentation.

Figure 14 supports the second observation by illustrating the relation between $\lfloor \overline{LQI} \rfloor$ and PDR , similarly to the previous analysis. The results for VINEYARD in Figure 14(c) are in line with the existing literature [Srinivasan et al. 2010]. In this scenario, LQI can be used only to identify good links, as it shows high variability below 110. However, both tunnels exhibit a different behavior, characterized by three areas. Below 85, LQI values are very variable, although not as much as in VINEYARD. Between 85 and 110, the variability is reduced, making LQI a reasonably good estimator also for intermediate links with $PDR > 50\%$. Above 110, LQI has a close correspondence with PDR .

These observations suggest that LQI should receive better consideration as an indication of the link quality in tunnel-like environments, as we discuss next.

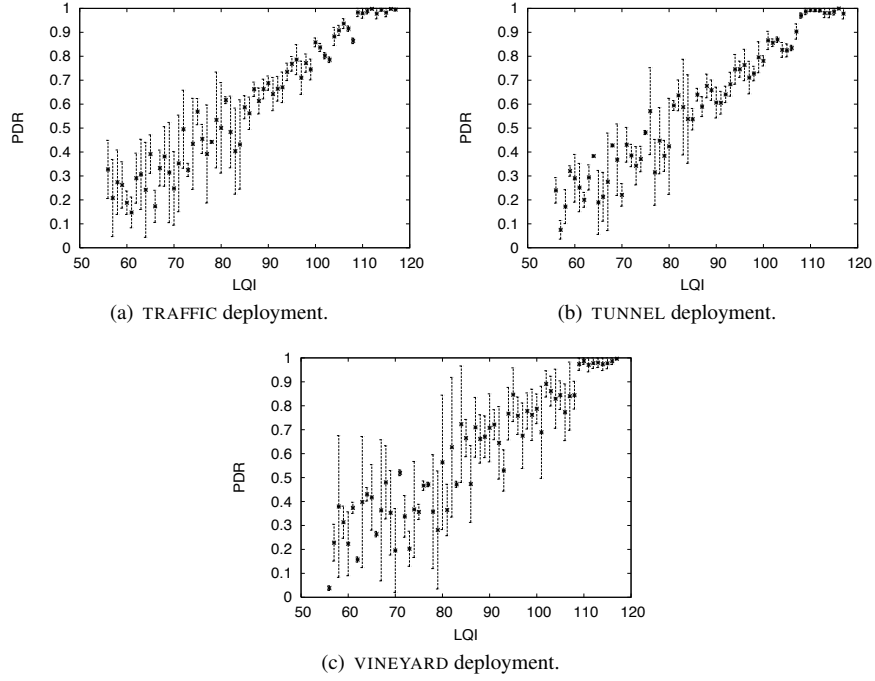


Fig. 14. Relation between $\lfloor \overline{LQI} \rfloor$ and average PDR with standard deviation.

Implications on routing. To study the impact of these findings on CTP, we ran 100 RTN tests with $DT=1$ hour and IMI ranging⁵ from 30 to 10 s. Based on the corresponding results, we assessed that the ETX estimator used in CTP captures reasonably well the quality of links in the tunnels. Indeed, as Figure 15(a) illustrates, CTP yields good message delivery in both tunnels. The performance in VINEYARD is not as good, however, presumably because of the poor connectivity pointed out in Figure 8. Furthermore, ETX recognizes the stability of network links in the tunnels, as Figure 15(b) highlights by plotting the av-

⁵CTP is not really designed to sustain $IMI < 30$ s. However, we wanted to stress its mechanisms to investigate how heavy network traffic impacts its performance.

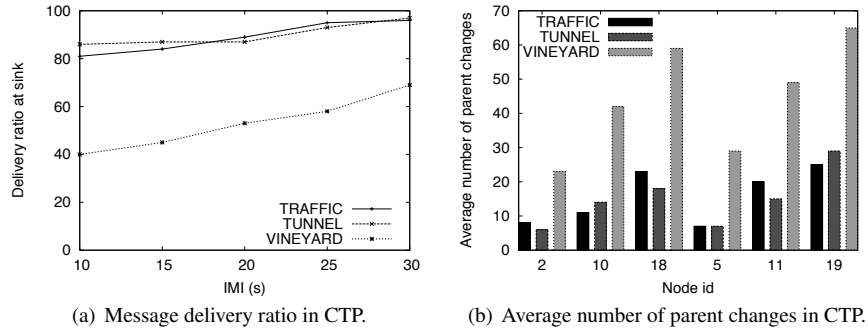


Fig. 15. CTP performance.




Fig. 16. Overlapping between ETX-based and LQI-based trees, in TUNNEL. The picture shows the links that are used more than 90% of the times in our CTP experiments and overlap with our topology based on LQI.

erage number of parent changes PT at some sample nodes. The experiments in TUNNEL and TRAFFIC indeed show a much smaller PT than in VINEYARD. Minimizing this figure is important in tree-based routing protocols, as frequent parent changes may cause temporary loops and thus redundant transmissions. Moreover, in every deployment, nodes farther from the root usually have a higher PT . This is due to the end-to-end nature of the ETX-based routing metric in CTP, which is likely to vary more as more physical links are travelled to reach the root, “accumulating” oscillations at each hop.

Nevertheless, we maintain that the properties shown by LQI in tunnels are an asset protocol designers should leverage off in this environment. In support of this claim, we carry out the following exercise: we take the values of \overline{LQI} in Figure 14 (from experiments with $DT=1$ min) and use them to compute off-line the values of a metric, alternative to ETX, used to derive the routing tree that maximizes end-to-end delivery. In a sense, we reconstruct off-line a routing topology similar to the one that LQI-based protocols such as MultiHopLQI [TinyOS b] would build⁶. Next, at every minute during a 1-hour long experiment, we log the shape of the CTP tree. Based on these data, we isolate the links used more than 90% of the time throughout an experiment. Figure 16 shows the links common to the two topologies, in TUNNEL. Surprisingly, there is a substantial overlapping: to make the topology connected, only 4 links are missing. A similar exercise for TRAFFIC experiments yields 5 missing links. However, our logs show that ETX requires at least an order of magnitude more messages per node to build such a topology.

5.3 Link Asymmetries

Physical layer. According to Srinivasan et al. [2010], a link between node i and node j is said to be asymmetric if $|PDR_{i \rightarrow j} - PDR_{j \rightarrow i}| > 40\%$. Existing works report that asymmetric links are mostly transient and presumably due to the different noise floor at different nodes, or to miscalibrated hardware [Zuniga and Krishnamachari 2007]. We analyze asymmetric links by looking at the results of PHY tests mentioned in Section 5.1, which lead to the following considerations:

- (1) as a consequence of the stability of network links in the tunnels, when asymmetric links are present they are mostly permanent;
- (2) in the two tunnels, link asymmetries seem to be related to node positioning; indeed,

⁶While [TinyOS b] considers single LQI values, in this case \overline{LQI} is the average of 10 LQI readings obtained during PHY tests with $DT=1$ min.

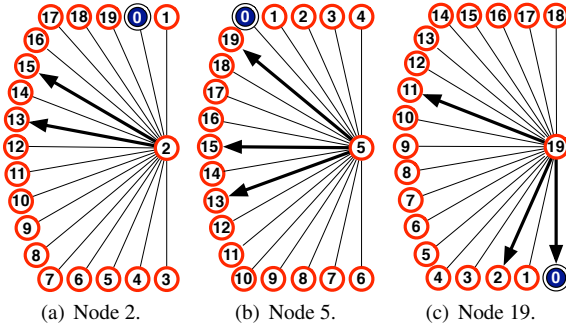


Fig. 17. Asymmetric links at example nodes. The node at the center is i such that $PDR_{i \rightarrow j} - PDR_{j \rightarrow i} > 40\%$. Thin lines denote bidirectional links, thick arrows denote asymmetric links.

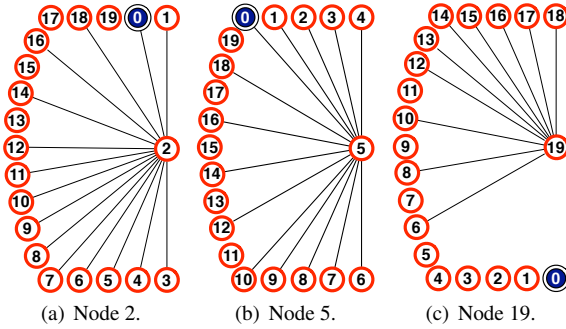


Fig. 18. Line-of-sight relations, represented by thin lines, at example nodes.

every asymmetric link was observed between nodes *not* in line-of-sight, although the opposite does not necessarily hold.

The first observation is supported by the results we already showed in Figure 10(b). Given the small variability of PDR among all transmitter-receiver pairs in the tunnels, an asymmetric link tends to remain so over time.

Figure 17 and 18 offer a graphical representation in support of the second consideration. For some sample nodes, Figure 17 depicts the links in TUNNEL found to be asymmetric in at least 90% of our PHY tests, whereas Figure 18 represents the line-of-sight relations between the chosen nodes and the remaining devices. By overlapping the two pictures we note that link asymmetries “imply” the *lack* of line-of-sight, but the vice versa does not hold. The PHY results we obtained in TRAFFIC confirm this phenomenon, in that all nodes are roughly in line-of-sight in this deployment and indeed we found no link asymmetries. This phenomenon, however, has serious implications on the performance of WSN routing protocols, as explained next.

Implications on routing. Asymmetric links are detrimental to the performance of CTP because of the link-level acknowledgement mechanism employed for reliable delivery. If the bidirectional-link assumption is not valid, the protocol ends up repeating the transmission a number of times, ultimately causing network congestion.




Fig. 19. Average ACK_{failed} in CTP.

A closer look at the RTN tests discussed in Section 5.2 indeed reveals that some nodes suffer a much higher number of failed acknowledgments ACK_{failed} compared to others, as Figure 19 illustrates. In particular, node 19 confirms our hypothesis that link asymmetries are most likely the cause of failed acknowledgments. Our logs show that at the beginning of the experiments node 19 often insisted in selecting node 2 as its parent—connected to 19 only through an asymmetric link, as shown in Figure 17(c)—before eventually settling on node 16. Instead, nodes 2 and 5, shown in Figure 17 to have asymmetric links only towards nodes very far from the CTP root, never selected any of those as their parent.

These behaviors also evidence how ETX incurs a significant latency—up to several minutes in our experiments—before eventually recognizing and excluding asymmetric links. From the perspective of routing protocol designers, these findings highlight that when reliability is provided with acknowledgment-like mechanisms, dedicated solutions are required to recognize asymmetric links. In addition, if the relation between link asymmetry and node placement were to be systematically confirmed, WSN designers could devise mechanisms to prevent routing through links that are, based on the shape of the scenario at hand, likely to be asymmetric.

5.4 Impact of Vehicular Traffic

Physical layer. The availability of TRAFFIC was a unique opportunity to investigate the impact of vehicular traffic on the system behavior. We carried out this analysis by looking at the results of 100 PHY tests with $DT=30$ min run during Jan. 12–16, 2009. As these span the 24 hours, they include periods with both *light traffic* (late night, <100 vehicles/hour) and *heavy traffic* (early morning, >800 vehicles/hour). Information on the flow of vehicles inside the tunnel was provided by the local authorities based on the observations at a monitoring station along the road leading to TRAFFIC. The results of our experiments form the basis for the following considerations:

- (1) vehicular traffic *does* impact the reliability of the network, especially for links of intermediate quality;
- (2) the links most affected are those between nodes on the opposite sides of the tunnel;
- (3) as expected, link disruptions are largely transient.

Figure 20(a) provides quantitative support for the first observation. We plot the complementary CDF for PDR , with light and heavy vehicular traffic. There is a substantial gap between the two curves until the PDR reaches 90%. Particularly, the largest gaps are ob-

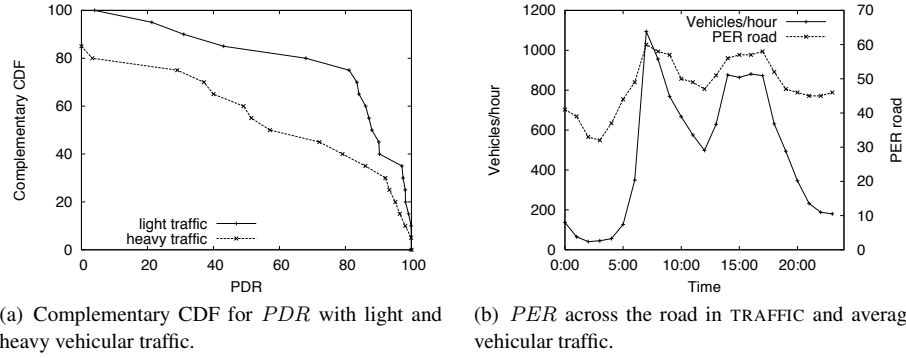


Fig. 20. Impact of vehicular traffic.

served for links whose PDR is between 40% and 80%. These links are more susceptible than others to temporary disruptions. On the other hand, when the transmission between two nodes is very reliable, vehicles have only limited impact.

To investigate the spatial characteristics of link disruptions, we compute the measured probability that no packet is delivered through any of the links across the road, according to the following formula:


$$PER_{road} = \prod_{i \in odd} \prod_{j \in even} (PER_{i \rightarrow j} \cdot PER_{j \rightarrow i})$$

where *odd* and *even* are the nodes on either side of the tunnel, as shown in Figure 2. In statistical terms, PER_{road} represents the aggregate PER between nodes on opposite sides of the tunnel.


Figure 20(b) plots PER_{road} and the average hourly vehicular traffic against time of day. Remarkably, the former closely follows the traffic trends. On the other hand, our data do not show relevant changes in PDR for links between nodes on the same side of the tunnel. This is further confirmed by looking at the variability of PDR over time. To this end, Figure 21 illustrates different views on the PDR standard deviation across all PHY experiments in TRAFFIC. Figure 21(a) considers only transmissions from node 0. The PDR is an order of magnitude more when the receiver is on the side opposite to where node 0 is deployed, compared to situations where the receiver is on the same side of node 0. Figure 21(b) and 21(c)⁷ complement this analysis by considering transmissions between all pairs of nodes, therefore illustrating the spatial trends. Figure 21(b) shows that heavy traffic causes the links connecting nodes on opposite sides to experience transient disruptions. Instead, links between nodes on the same side are largely unaffected by traffic: the trends in Figure 21(c) are comparable to the average ones in Figure 10(b) for TRAFFIC. Therefore, our results quantitatively confirm the last two observations above.

To the best of our knowledge, we are the first to investigate the impact of vehicular traffic on WSN transmissions and to provide quantitative data to characterize it. Next, we provide some insights into how these observations affect the behavior of routing protocols like CTP.


⁷Data points along the diagonal of the base plane are excluded from the analysis.




(a) Transmissions from node 0.



(b) Nodes on opposite sides. (All pairs).



(c) Nodes on the same side. (All pairs).

Fig. 21. PDR standard deviation in TRAFFIC, with heavy vehicular traffic.Fig. 22. DUP and ACK_{failed} in CTP with light and heavy vehicular traffic.

Implications on routing. To study this aspect, we ran 100 RTN tests with $DT=30$ min and $IMI=30$ s interleaved with the previous PHY tests. Generally, the temporary link disruptions caused by heavy traffic increase the dynamicity of the environment, yet CTP is able to cope with this scenario. The message delivery ratio is indeed almost the same as with light traffic. However, this occurs at the cost of a drastic increase in resource consumption. An indication of this comes from the higher ACK_{failed} and corresponding retransmissions, illustrated in Figure 22. Most likely, these are a direct result of the temporary link disruptions discussed above. Moreover, these also affect the ETX metric, yielding frequent parent changes. In turn, this causes temporary routing loops, resulting in a higher number of duplicates DUP in the presence of heavy traffic, also shown in Figure 22.




Fig. 23. Links used at least 90% of the times by CTP in TRAFFIC.

It is also interesting to analyze *which* links are most used by CTP under heavy traffic. Figure 23(a) depicts the links used more than 90% of the times⁸ under light traffic. Figure 23(b) shows the same metric under heavy traffic. In the former case these links are more numerous, and many connect nodes on opposite sides. In the latter case, not only there are fewer links used often, but almost all frequently-used links connect nodes on the same side, unless a node is very close to the CTP root. This behavior was expected, as it is the most natural way for ETX to react to the characteristics of the physical layer we discussed above. The decrease in the number of links used often is a consequence of the increased variability of the link performance, whereas their particular location is the effect of temporary link disruptions between nodes on opposite sides of the tunnel.

The observations discussed here, especially those regarding the resource consumption of CTP to sustain good performance in this setting, suggest that phenomena like link disruptions caused by heavy vehicular traffic deserve more attention. As part of our future work, we plan to investigate a link quality metric immune to these phenomena, allowing us to establish routes independently of vehicular traffic, and to design reliability mechanisms that take into account this kind of temporary link disruption.

6. RELATED WORK

The scientific literature related to our work can be divided in three classes: *i)* analytical propagation models for tunnel-like environments obtained from, or verified with, short-term experimental results gathered with specialized equipment; *ii)* experiences from real-world WSNs deployments; and *iii)* systematic studies on low-power wireless transmissions, mostly performed in controlled testbeds. The pragmatic, system-oriented perspective we took distinguishes our work from those in *i)*. Indeed, these works focus on issues

⁸These links do not necessarily form a fully-connected topology. For instance, node 11 is not connected by any of these links: its data will be relayed through links (not shown here) that are used less than 90% of the times.

such as deciding the optimal modulation scheme in a tunnel-like environment, while in our case these aspects are dictated by the WSN platform we chose. On the other hand, our work can be regarded as a combination of the approaches taken in *ii*) and *iii*), with the goal to increase the intrinsic value of our results to WSN designers. However, it should be noted that the works in *iii*) focus mostly on the physical layer with a few exceptions considering also non-duty-cycling MAC protocols. In contrast, our study extends up to the routing layer, which bears a more direct impact on the application performance.

Propagation models. Because of their ability to act as an “oversized waveguide” [Molina-Garcia-Pardo et al. 2009], tunnel-like environments have been quite extensively studied in the field of wireless communications. In their recent work, Molina-Garcia-Pardo et al. [2009] investigate the electromagnetic field statistics in an arched tunnel across frequency bands from 2.8 GHz to 5 GHz to develop a detailed channel model. Sun and Akyildiz [2008] present a multi-mode channel model for road tunnels and subways, able to determine the path loss and delay spread at any point along the tunnel. Earlier works (e.g., [Wang and Yang 2006]) focus on understanding the behavior of wireless propagation at frequencies used by mobile telephony. As already mentioned, these works have different and complementary goals w.r.t. ours. However, the results reported in this paper could be used to verify the existing analytical propagation models against experimental findings.

Application deployments. A number of application deployments witness behaviors similar to the ones we observed in VINEYARD. For instance, Mainwaring et al. [2002] report poor system performance in outdoor environments with significant vegetation, using WSN platforms operating in the 400/900 MHz band. However, the authors also acknowledge problems with their implementation, which makes their results difficult to compare with ours. Thelen et al. [2005] report about a deployment in potato fields where the link layer performance improves with high humidity, contrary to our findings in all deployments. We conjecture that this is due to the radio technology they employed (400 MHz band) and the placement of nodes close to the ground.

Grape monitoring is investigated, for instance, in [Beckwith et al. 2004], where the authors report a partial failure of the WSN due to connectivity issues. However, no quantitative data is gathered to analyze the corresponding causes. The use of WSNs is also reported in environments analogous to our tunnels, such as metropolitan subways [Cheekiralla 2005], underground mines [Li and Liu 2009], and service pipes [Sabata and Brossia 2005]. However, to the best of our knowledge, a characterization of the performance of low-power wireless transmissions in these scenarios is largely missing.

Systematic studies. Srinivasan et al. [2010] carry out a comprehensive study on IEEE 802.15.4 radios in the absence of concurrent transmissions. Their objective is different from ours in that they mostly aim to study the link-layer behavior on tiny time scales, to provide guidelines on fine-grained design decisions such as scheduling re-transmissions at the MAC layer. Because of our practical need to develop a real-world, efficient WSN system, we were more interested in the long-term behaviors affecting macro design decisions, e.g., whether to choose a CSMA or a TDMA MAC protocol.

The symmetry of WSN links is the subject of several studies [Srinivasan et al. 2010; Zuniga and Krishnamachari 2007; Reijers et al. 2004], yet general conclusions are still to come. Srinivasan et al. [2010] notice that asymmetric links are mostly transient and due to differences in the noise floor. Zuniga and Krishnamachari [2007] analyze link asymmetries due to imperfections in the hardware that may cause variations in the power output during

transmissions. Both Zhou et al. [2004] and Reijers et al. [2004] investigate the effect of antenna orientation on the generation of asymmetric links. In contrast to these works, in TUNNEL we observed asymmetric links to be permanent and related to node placement.

Zhao and Govindan [2003] study the link characteristics using low-power radios in the 800/900 MHz band. They verify that grey areas may extend to up to one third of the communication range, and conjecture that multi-path effects are the cause. The analytical models defined by Zuniga and Krishnamachari [2004] confirm this hypothesis. The results we presented in Section 5.1 also concur, although we use a different wireless technology. Zhao and Govindan [2003] also hint at a MAC-level neighbor filtering mechanism to address this issue. This solution, however, is unlikely to work in a tunnel because of the influence of far transmitters without link-layer visibility, as we illustrated in Section 5.1.

Several experimental studies exist whose objective is to derive synthetic models to use in simulations. For instance, Son et al. [2006] focus on understanding the effect of concurrent transmissions and on their modeling in simulation. Anastasi et al. [2005] study factors such as distance from the ground and transmission data rate, and derive corresponding analytical models. These works are complementary to ours. By making our data publicly available [D3S Group], we aim to provide further input to similar investigations.

Understanding the characteristics of WSN links in indoor environments has received particular attention. For instance, Lal et al. [2003] present an experimental study on wireless link quality in office environments to derive a link cost metric that minimizes the necessary channel measurements. Tang et al. [2007] perform extensive experimental studies to understand the effect of obstacles and moving objects in factory environments. As already mentioned, road tunnels bring together elements from both indoor and outdoor deployments. For instance, despite being delimited by walls, there are no fixed obstacles in the tunnels. Thus, they represent a unique setting for which we are the first to report a quantitative characterization based on extensive experimental studies.

Finally, some works explored the use of external antennas in WSNs. An example is the work by Raman et al. [2006], where different external antennas are tested in a narrow-road and a dense-foliage environments. Using a parabolic antenna, the authors reach communication ranges up to 800 m when the nodes are in line of sight, leading to the conclusion that 1-hop topologies and centralized routing may be a viable alternative to multi-hop topologies and distributed routing. However, as already mentioned, in our scenario the use of external antennas would drastically increase costs and complicate the installation, while in the tunnels the on-board antenna already enables communication ranges above 100 m.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we reported about an extensive experimental study on the behavior and performance of WSNs in tunnels. We analyzed the behavior of the physical layer in real road tunnels, and its impact on the operation of commonly-employed MAC and routing protocols. We showed that some popular networking techniques are ill-suited to the tunnel environment, and provided guidelines to support the design of more efficient solutions.

The results we presented are an asset for the community at large, although they were originally motivated by the needs of the TRITon project. The complete data set gathered during our experiments is publicly available [D3S Group]. Our immediate plans are to apply our findings in the TRITon project, beginning with the design of a dedicated MAC layer providing an efficient foundation for higher-level protocols.

ACKNOWLEDGMENTS

The authors wish to thank the Autonomous Province of Trentino, and in particular the Department of Road Management (Servizio Gestione Strade) that partially funded the TRITon project and therefore this work, for making available to us the tunnels and related traffic data described in this paper. The authors also wish to thank the Edmund Mach Foundation, in S. Michele all'Adige, Italy, for making their vineyards available to us, and for their logistical support with the experimental setup.

REFERENCES

- ANASTASI, G., BORGIA, E., CONTI, M., GREGORI, E., AND PASSARELLA, A. 2005. Understanding the real behavior of mote and 802.11 ad hoc networks: An experimental approach. *Elsevier Pervasive and Mobile Computing Journal* 1, 2, 237–256.
- BECKWITH, R., TEIBEL, D., AND BOWEN, P. 2004. Unwired wine: Sensor networks in vineyards. In *Proc. of IEEE Sensors*. IEEE, 561–564.
- CHEEKIRALLA, S. 2005. Wireless sensor network-based tunnel monitoring. In *Proc. of the RealWSN Workshop*.
- CHIPCON TECH. CC2420 Datasheet. focus.ti.com/docs/prod/folders/print/cc2420.html. Accessed on 1/2010.
- COSTA, P. ET AL. 2007. The RUNES middleware for networked embedded systems and its application in a disaster management scenario. In *Proc. of the 5th Int. Conf. on Pervasive Communications*. IEEE, 69–78.
- COUTO, D. D., AGUAYO, D., BICKET, J., AND MORRIS, R. 2005. A high-throughput path metric for multi-hop wireless routing. *Wireless Networks* 11, 4, 419–434.
- D3S GROUP. Comparative Study on Tunnels - Data Sets. d3s.disi.unitn.it/tunnelvineyard.
- ELSON, J., GIROD, L., AND ESTRIN, D. 2002. Fine-grained network time synchronization using reference broadcasts. *SIGOPS Operating Systems Review* 36, SI, 147–163.
- GNAWALI, O., FONSECA, R., JAMIESON, K., MOSS, D., AND LEVIS, P. 2009. The collection tree protocol. In *Proc. of the 7th Int. Conf. on Embedded Networked Sensor Systems (SenSys)*. ACM, 1–14.
- HILL, J., SZEWCZYK, R., WOO, A., HOLLAR, S., CULLER, D., AND PISTER, K. 2000. System architecture directions for networked sensors. In *Proc. of the 9th Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*. ACM, 93–104.
- LAL, D., MANJESHWAR, A., HERRMANN, F., UYSAL-BIYIKOGLU, E., AND KESHAVARZIAN, A. 2003. Measurement and characterization of link quality metrics in energy constrained wireless sensor networks. In *Proc. of the Int. Global Telecommunications Conf. (GLOBECOM)*. IEEE, 446–452.
- LI, M. AND LIU, Y. 2009. Underground coal mine monitoring with wireless sensor networks. *ACM Transactions on Sensor Networks* 5, 2, 1–29.
- MAINWARING, A., CULLER, D., POLASTRE, J., SZEWCZYK, R., AND ANDERSON, J. 2002. Wireless sensor networks for habitat monitoring. In *Proc. of the 1st Int. Workshop on Wireless Sensor Networks and Applications*. ACM, 88–97.
- MOLINA-GARCIA-PARDO, J., LIENARD, M., AND DEGAUQUE, P. 2009. Propagation in tunnels: Experimental investigations and channel modeling in a wide frequency band for MIMO applications. *EURASIP Journal on Wireless Communications and Networking* 2009, 3, 1–9.
- NOWAK, B., FILEK, K., AND SINHA, A. 1997. Temperature and humidity of cooled air in mining galleries. In *Proc. of the 6th Int. Mine Ventilation Congress*. Society for Mining, Metallurgy, and Exploration.
- POLASTRE, J., SZEWCZYK, R., AND CULLER, D. 2005. Telos: Enabling ultra-low power wireless research. In *Proc. of the 5th Int. Conf. on Information Processing in Sensor Networks (IPSN)*. IEEE, 48.
- RAJENDRAN, V., OBRACZKA, K., AND GARCIA-LUNA-ACEVES, J. J. 2006. Energy-efficient, collision-free medium access control for wireless sensor networks. *Wireless Networks* 12, 1.
- RAMAN, B. AND CHEBROLU, K. 2008. Sensor networks: a critique of “sensor networks” from a systems perspective. *SIGCOMM Computing Communication Review* 38, 3, 75–78.
- RAMAN, B., CHEBROLU, K., MADABHUSHI, N., GOKHALE, D., VALIVETI, P., AND JAIN, D. 2006. Implications of link range and (in)stability on sensor network architecture. In *Proc. of the 1st Int. Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*. ACM, 65–72.

- REIJERS, N., HALKES, G., AND LANGENDOEN, K. 2004. Link layer measurements in sensor networks. In *Proc. of the Int. Conf. on Mobile Ad-hoc and Sensor Systems (MASS)*. IEEE, 224–234.
- SABATA, A. AND BROSSIA, S. 2005. Remote monitoring of pipelines using wireless sensor network. US Patent Application 20050145018.
- SON, D., KRISHNAMACHARI, B., AND HEIDEMANN, J. 2006. Experimental study of concurrent transmission in wireless sensor networks. In *Proc. of the 4th Int. Conf. on Embedded Networked Sensor Systems (SenSys)*. ACM, 237–250.
- SRINIVASAN, K., DUTTA, P., TAVAKOLI, A., AND LEVIS, P. 2010. An empirical study of low power wireless. *ACM Transactions on Sensor Networks*. To appear.
- SRINIVASAN, K. AND LEVIS, P. 2006. RSSI is under-appreciated. In *Proc. of the 3rd Int. Workshop on Embedded Networked Sensors (EmNets)*.
- SUN, Y.-J., LI, J., CAO, J., LU, S.-L., AND SHI, Y. 2008. Small variation of annual temperature in deep tunnel can produce annual variation in tilt and strain. *Acta Seismologica Sinica* 21, 5, 464–473.
- SUN, Z. AND AKYILDIZ, I. 2008. Channel modeling of wireless networks in tunnels. In *Proc. of the Int. Global Telecommunications Conf. (GLOBECOM)*. IEEE, 158–162.
- TANG, L., WANG, K.-C., AND GU, Y. H. F. 2007. Channel characterization and link quality assessment of IEEE 802.15.4-compliant radio for factory environments. *IEEE Transactions on Industrial Informatics* 3, 2, 99–110.
- THELEN, J., GOENSE, D., AND LANGENDOEN, K. 2005. Radio-wave propagation in potato fields. In *Proc. of 1st Workshop on Wireless Network Measurements*.
- TINYOS. TEP 105 - Low Power Listening. www.tinyos.net. Accessed on 1/2010.
- TINYOS. TEP 119 - Collection. www.tinyos.net. Accessed on 1/2010.
- TINYOS. TEP 123 - Collection Tree Protocol. www.tinyos.net. Accessed on 1/2010.
- TINYOS. TEP 126 - CC2420 Radio Stack. www.tinyos.net. Accessed on 1/2010.
- WANG, T.-S. AND YANG, C.-F. 2006. Simulations and measurements of wave propagations in curved road tunnels for signals from GSM base stations. *IEEE Transactions on Antennas and Propagation* 54, 9, 2577–2584.
- ZHAO, J. AND GOVINDAN, R. 2003. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. of the 3rd Int. Conf. on Embedded Networked Sensor Systems (SenSys)*. ACM, 1–13.
- ZHOU, G., HE, T., KRISHNAMURTHY, S., AND STANKOVIC, J. 2004. Impact of radio irregularity on wireless sensor networks. In *Proc. of the 2nd Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*. ACM, 125–138.
- ZUNIGA, M. AND KRISHNAMACHARI, B. 2004. Analyzing the transitional region in low power wireless links. In *Proc. of the Int. Conf. on Sensors and Ad-hoc Communications and Networks (SECON)*. IEEE, 517–526.
- ZUNIGA, M. AND KRISHNAMACHARI, B. 2007. An analysis of unreliability and asymmetry in low-power wireless links. *ACM Transactions on Sensor Networks* 3, 2, 63–81.