Opportunistic, Receiver-Initiated Data-Collection Protocol*

Stefan Unterschütz, Christian Renner, and Volker Turau

Hamburg University of Technology, Hamburg, Germany
{stefan.unterschuetz,christian.renner,turau}@tu-harburg.de

Abstract. This paper presents and evaluates ORiNoCo, a novel data-collection and event-reporting protocol for sensor networks. ORiNoCo is built upon the asynchronous duty-cycle protocol RI-MAC and breaks with the tradition of exchanging extensive neighborhood information, a cornerstone of many competing collection protocols and one of their major source of communication overhead and energy expenditure. The merit of this venture is an opportunistic, energy-efficient, latency-reducing, and self-stabilizing protocol. ORiNoCo comes at virtually no extra costs in terms of memory demand and communication overhead compared to RI-MAC. We derive theoretical boundaries for the improvements in radio efficiency, latency, and energy-consumption. ORiNoCo is verified with these findings via simulation and compared with CTP. ORiNoCo achieves lower energy-consumption while reducing end-to-end delays.

1 Introduction

Data collection is one of the major driving applications of sensor networks. In this scenario, all data sampled on the individual nodes is transported to a data sink. Beside sampling data, sensor nodes act as data forwarders, since networks generally cover areas too large for direct communication with the sink. There exists a diversity of solutions for obstacles in the way of data collection; among these obstacles being low computing and memory resources, limited bandwidth, poor and fluctuating radio connectivity, and restricted energy supplies.

Link-quality estimation techniques [6,17] have been proposed to tackle the problem of the fluctuating, unstable, and lossy nature of the wireless channel [22] in order to identify reliable communication partners. Low-power medium access protocols were designed to buy extended network lifetime at the cost of increased latency [12,20,5]. These techniques were used to devise both generic [1,7,15] and special purpose [3,10,13] collection protocols. Despite the success of these protocols, they impose some extra cost on the network by requiring up-to-date knowledge of state information; particularly neighborhood tables have to be exchanged—including node IDs and link-quality metrics (e.g., LQI, RSSI, or packet success rates). This results in energy-expensive protocol overhead.

We believe that in many cases there is a cheaper option, which we will substantiate in this paper: By adding a set of modifications and integrating a path metric, we turn the Receiver-Initiated MAC [20] (RI-MAC) from a

* The original publication is available at www.springerlink.com

simple point-to-point communication protocol to an end-to-end routing mechanism for data collection. These changes are accompanied by enhancements and fine-tunings of the protocol. The result is ORiNoCo, an opportunistic, resource- and energy-efficient collection protocol for sporadic data generation. ORiNoCo introduces virtually no overhead compared to RI-MAC. It provides self-stabilizing properties for adaptive route correction in the face of changing network connectivity, while it avoids exchanging neighborhood information. Moreover, ORiNoCo can be extended with energy-aware techniques—such as adaptive duty cycling [8]—making it compatible to the evolving field of energy-harvesting sensor networks [9].

The idea behind our approach is as easy as it is functional. Instead of building an explicit routing tree, thus forcing nodes to select a single parent, we generate an implicit and loose tree-like structure. A path metric is used to enable nodes to identify the minimum cost for sending data to the sink. By attaching this cost to RI-MAC beacons, nodes advertise an implicit path to the sink. Nodes willing to transmit data decide upon reception of a beacon whether to commit to the offer, i.e., whether they forward their data to the sender of the beacon. This scheme is inspired by stock-market trading, where bids and asks are matched. Accepting an offer with slightly higher costs than the latest minimum offer allows a sender to deliver its packet earlier and thus switch off its radio sooner. Reduced waiting time entails lower energy consumption and decreases end-to-end latency.

The paper makes the following contributions. We (i) introduce a light-weight and efficient data-collection protocol, (ii) assess theoretical limits for the expected waiting time of the protocol, showing its potential to reduce latency and energy-consumption, (iii) evaluate these findings by simulation and compare them with CTP, and (iv) end with a thorough discussion of the results.

2 State-of-the-Art Data-Gathering

The design of collection protocols for sensor networks has a long tradition. Kim et al. presented the reliable bulk transport protocol Flush [10]. Like ORiNoCo, Flush is receiver-initiated but can handle only one active flow (per sink) at a time. Flush produces extra control traffic to set up routes, ORiNoCo prevents control traffic and selects routes opportunistically. The Koala protocol is designed for long-term, infrequent bulk-data collection at ultra-low duty cycles of less than one percent [13]. The network is sleeping most of the time; the sink wakes up the network using low-power probing (cf. Sect. 3.2) to initiate bulk data download. After the wake-up, nodes collect neighborhood information and report it to the sink; the protocol does hence not try to keep track of network states during sleep times. Routing paths are established by the sink using this neighborhood information. The centralized approach and bulk data transfer are major differences to ORiNoCo. In contrast to Koala and other bulk transfer protocols, ORiNoCo offers continuous availability for low-latency event reporting.

The Collection Tree Protocol (CTP) is probably the routing mechanism most frequently used for multi-hop collection in sensor networks [7]. The strengths of

CTP are its ability to quickly discover and repair path inconsistencies and its adaptive beaconing, which reduces protocol overhead and allows for low radio duty cycles. Arbutus [15] is a protocol similar to CTP, but promises higher performance in weakly connected networks with poor node connectivity. Unlike CTP and Arbutus, ORiNoCo avoids control packets by adding minimal information to the beacons produced by the underlying low-power probing MAC protocol.

The benefits of exploiting unstable but bursty links have been discussed in [1]. The authors show that employing short-term link-quality estimation to identify and subsequently use these links can improve the performance of multi-hop routing, e.g., by adding bursty-link usage to CTP. However, explicit link-quality estimation is required, which sets this approach apart from ORiNoCo.

A similar concept to ORiNoCo was presented in [11]. The authors argue that link-quality estimation poses large overhead on network traffic and nodes' resources, while estimates may be outdated at the time they are used for routing decisions. They suggest using a low-power probing MAC and integrating a routing metric into the probes, but only evaluate link availability.

3 Energy-Efficient Medium Access Control

This section explains established approaches on energy-conserving medium access control in sensor networks for point-to-point communication. The underlying principles build the foundation for our collection protocol.

The radio of a sensor node is its main energy consumer, yet is only needed upon transmission and reception. Since these are infrequent events in sensor networks, there has been considerable effort to decrease the power consumption of a node by duty-cycling its radio while still maintaining low latency of packet delivery. The main challenge of low-power communication is posed by the dilemma that a node cannot be contacted from another node, if its radio is switched off; there hence must be a kind of coordination. If data and therefore radio packets are not generated based on a schedule (e.g., when using TDMA), an intended receiver cannot know, when it has to enable its radio for packet reception. The two predominant solution concepts to this problem are explained in the following.

3.1 Low-Power Listening

The idea of low-power listening (LPL) is to divide time into cycles of length T_{slp} and switch on the radio only once per cycle to check briefly if there is activity on the channel. Nodes willing to transmit data must occupy the channel via sending long preambles (B-MAC [14]), repeated preambles (X-MAC [2]), or repeated data packets (BoX-MAC [12]). Channel occupation must be sufficiently long to ensure that a receiver wakes up and switches on its radio before the transmitter stops its activity. On average, this implies that the transmitter congests the channel for half a cycle, possibly causing interference on other nodes, thus preventing packet reception and affecting throughput. Choosing the parameter T_{slp} is critical for network latency, throughput, and lifetime. While a small value of T_{slp} offers low latency and high throughput, energy expenditure is heavy, since nodes are awake frequently. However, the authors of [18] report that choosing T_{slp} too large may also result in poor energy-efficiency. Here, sending a packet induces long waiting times and energy costs for the sender, since the receiver wakes up infrequently.

3.2 Receiver-Initiated Medium Access Control (RI-MAC)

RI-MAC [20] is one of the latest low-power MAC protocols and inverts the concept of low-power listening. RI-MAC does not rely on repeatedly sent packets (or preambles) by the sender. A node willing to send data switches on its radio and waits for the receiver to also enable its radio. The receiver signals that it is ready to receive by sending a beacon—this concept is called low-power probing (LPP)—and the sender transmits its data packet immediately upon reception of this beacon, if the channel is idle. The receiver acknowledges successful packet reception with another beacon, containing the ID of the sender. Having received the (acknowledging) beacon, the sender either switches off its radio, if there are no more packets left, or transmits additional packets to the (same) receiver. The receiver node dwells (waits) for a short period and switches off its radio, if no additional data packet arrives. To decrease the chance of accidental and undesired node synchronization—which results in beacon-collisions—RI-MAC chooses inter-beacon time randomly from the interval $[0.5 \cdot T_{\mathrm{slp}}, 1.5 \cdot T_{\mathrm{slp}}]$.

The dwelling time T_{dwell} serves two purposes. Firstly, it enables additional nodes to send data to the receiver. Secondly, every beacon contains the current dwelling time, which defines a backoff window for all potential senders to reduce the probability of collisions. If a receiving node identifies collisions, it adapts the backoff window—e.g., by increasing the dwelling time using binary exponential backoff—and distributes the value to the contending senders via the following beacon. Since senders choose a random backoff using the dwelling time, the likelihood of collisions is reduced. Initial beacons sent out directly after a node's wake-up contain zero dwelling time. The authors argue that this approach is beneficial in networks with low traffic and that many collisions are prevented by the capture effect (cf. [21]). If the receiver reaches the maximum backoff time and still experiences collisions, it switches off its radio without notice and sleeps one cycle. A sender manages a retry counter and cancels a transmission, if a predefined limit of retries is exceeded. The retry counter is incremented, if no beacon is received from the intended receiver within a given number of sleep periods—i.e., there appears to be no connection to the receiver. If the sender has received an initial beacon and thus sent its data but has not received an acknowledging beacon within the dwelling time provided in the latest beacon, the retry counter is also incremented.

The choice of T_{slp} is also critical for RI-MAC; the same observations as explained in Sect. 3.1 apply. Unlike LPL, RI-MAC avoids channel congestion for large values of T_{slp} by not sending repeated preambles. Yet, RI-MAC will produce congestion for too small values of T_{slp} due to an increasing number of beacons: The more frequent a node wakes up, the more beacons it produces.

4 Opportunistic, Receiver-Initiated Data Collection

This section presents ORiNoCo, an opportunistic, resource-efficient, and energy-conserving collection protocol for sensor networks. ORiNoCo is based on the RI-MAC protocol (cf. Sect. 3.2). We add a path weight ϕ that is used to implicitly build a tree-like routing structure to the RI-MAC beacons. The value of ϕ reflects the costs for delivering data to the sink. Upon reception of a beacon from neighbor $\mathbf{b}$, node $\mathbf{v}$ calculates the path weight $\phi_{\mathbf{v},\mathbf{b}}$ via $\mathbf{b}$ using the cost metric

$$\phi_{\mathbf{v},\mathbf{b}} \leftarrow \phi_{\mathbf{b}} + \kappa_{\mathbf{v},\mathbf{b}} \,. \tag{1}$$

Here, $\phi_{\mathbf{b}}$ is the weight carried by the beacon from node $\mathbf{b}$ and $\kappa_{\mathbf{v},\mathbf{b}}$ is a measure of the cost (calculated by $\mathbf{v}$) for sending a packet to $\mathbf{b}$. If ϕ is the hop count metric, then $\kappa_{\mathbf{v},\mathbf{b}} = 1$. Each node $\mathbf{v}$ only stores the best known weight $\phi_{\mathbf{v}}$, which is updated only if a better path is identified. Initially, $\phi_{\mathbf{v}} = \infty$ for all nodes $\mathbf{v}$ except the sink $\mathbf{s}$, which always maintains $\phi_{\mathbf{s}} = 0$.

In the following we describe ORiNoCo in more detail. We distinguish between the *initialization phase*, which runs only once for connected graphs and can be started asynchronously, and the *packet transmission phase*. The latter contains a *recovery mode* in case of topology changes.

4.1 Initialization Phase

A node $\mathbf{v}$ that has not yet received a path weight from a neighbor is in the initialization phase: $\mathbf{v}$ switches on its transceiver and starts listening to the wireless channel, but does not send any beacons itself. If $\mathbf{v}$ receives a beacon from a neighbor $\mathbf{b}$, $\mathbf{v}$ calculates an initial weight using (1). Once node $\mathbf{v}$ has accepted a beacon and assigned a value to $\phi_{\mathbf{v}}$, $\mathbf{v}$ proceeds to the *packet transmission phase* and starts the RI-MAC protocol. The initial assignment of weights shapes an implicit tree-like routing structure that is not necessarily optimal in terms of ϕ. Further optimization and maintenance of routing information are achieved during the *packet transmission phase*.

4.2 Packet Transmission Phase

Each node broadcasts beacons with an inter-beacon time chosen randomly from the interval $[(1-\alpha)\cdot T_{\text{slp}}, (1+\alpha)\cdot T_{\text{slp}}]$, where $0<\alpha<1$. This interval is a generalization of the inter-beacon time used with RI-MAC (cf. Sect. 3.2). We analyze the influence of α in Sect. 5.1.

A node willing to transmit data starts listening to the channel and waits for a beacon from an arbitrary node $\mathbf{b}$ satisfying

$$\phi_{\mathbf{v},\mathbf{b}} = \phi_{\mathbf{b}} + \kappa_{\mathbf{v},\mathbf{b}} \le \phi_{\mathbf{v}} \,, \tag{2}$$

which guarantees routing progress towards the sink. If (2) holds, node $\mathbf{v}$ transmits its packet to $\mathbf{b}$ and waits for an acknowledging beacon from $\mathbf{b}$. Only if that

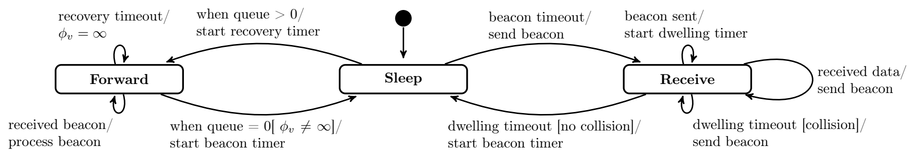

Fig. 1: ORiNoCo's state machine (without initialization phase) in UML notation

beacon is received, $\mathbf{v}$ updates its weight $\phi_{\mathbf{v}}$ with $\phi_{\mathbf{v},\mathbf{b}}$. The goal is to prevent $\mathbf{v}$ from adapting to a path with low cost but poor connectivity. If no acknowledging beacon is received, $\mathbf{v}$ waits for the next beacon (possibly from a different node).

In order to favor links with high delivery probability, link-quality estimates, such as LQI or RSSI, should be used to ignore beacons from neighbors offering a poor link [19]. We do not consider logical link metrics based on packet statistics, since they cannot be implemented in a stateless fashion, i.e., nodes must explicitly keep track of their neighbors and the corresponding link metric [17].

ORiNoCo enables $\mathbf{v}$ to forward packets to the first available parent satisfying (2), so that the expected waiting time per node is decreased. Hence, radio usage and therefore energy-consumption are reduced, while packet latency is decreased.

4.3 Recovery Strategy

If node $\mathbf{v}$ does not receive a beacon satisfying (2) within a period equal to the maximum sleep cycle $(1+\alpha) \cdot T_{\mathrm{slp}}$, it overwrites its weight $\phi_{\mathbf{v}}$ with the minimum $\phi_{\mathbf{v},\mathbf{b}}$ of all rejected beacons during that time (excluding those ignored due to low link properties, cf. Sect. 4.2), and starts a new transmission cycle. If $\mathbf{v}$ received one or more beacons yielding sufficient values $\phi_{\mathbf{v},\mathbf{b}}$, but did not receive any acknowledging beacon, $\mathbf{v}$ also increases its value $\phi_{\mathbf{v}}$ to the minimum $\phi_{\mathbf{v},\mathbf{b}}$ of all rejected beacons. If no beacon is received for $(1+\alpha) \cdot T_{\mathrm{slp}}$, $\mathbf{v}$ resets $\phi_{\mathbf{v}}$ to ∞.

4.4 State Machine

ORiNoCo can be described as state machine (see Fig. 1). Note that the initialization phase and implementation details are omitted in the figure for better visualization. A node is either sleeping, receiving, or forwarding data. Most of the time is spent in the first of these states, while the radio is only on in the latter two. The state machine shows that receiving and forwarding are non-interleaving:

– A node does not send beacons, while it is forwarding. This strategy prevents that node from accidentally failing to transmit a packet (after receiving a beacon from a neighbor), if its beacon timer fires simultaneously.
– A node stays in the receive state, until no more packets are received, i.e, there is no packet left to be received. This strategy enables data aggregation and saves energy, as average waiting time is reduced for subsequent transmissions.

The state machine illustrates the recovery mechanism of ORiNoCo: upon expiration of the recovery timeout, $\phi_{\mathbf{v}}$ is reset and the node does not leave the forwarding state, until it has accepted a weight $\phi_{\mathbf{b}}$.

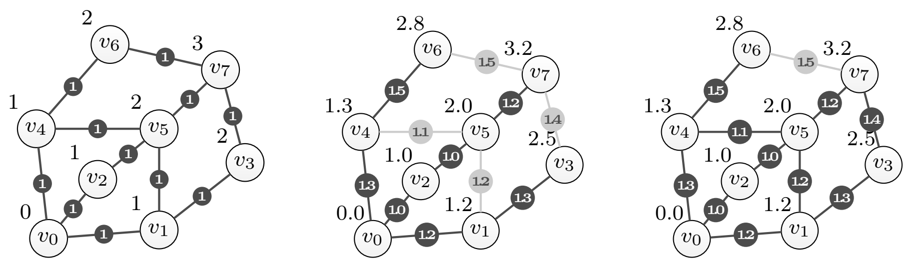

(a) Hop Count, $\Theta = 0$ (b) ETX, $\Theta = 0$ (c) ETX with $\Theta = 0.9$

Fig. 2: Comparison of hop count and ETX metric w.r.t. available parents. Links not available with the ETX metric are light gray. Link weights $\kappa_{\mathbf{v},\mathbf{b}}$ and node weights $\phi_{\mathbf{v}}$ are attached to the corresponding entities. $\mathbf{v}_0$ is the sink

4.5 Choosing a Weight Metric

The granularity of the weight influences the number of available parents, cf. (2). For the hop count metric, the number of parents is high, as many nodes $\mathbf{v}$ have the same weight $\phi_{\mathbf{v}}$. This situation is shown in Fig. 2a. For a more fine-grained metric (e.g., ETX values [4]) the following situation is likely: The network stabilizes, i.e., all nodes $\mathbf{v}$ have identified their minimum $\phi_{\mathbf{v}}$. Since the values $\phi_{\mathbf{v}}$ are fine-grained, the chance of a node having multiple neighbors satisfying (2) is decreased. Since $\mathbf{v}$ will eventually update $\phi_{\mathbf{v}}$ to the minimum in its vicinity, only the nodes offering this minimum weight are available parents w.r.t. (2). In the worst case (see Fig. 2b), each node will have only a single parent.

For that reason, we suggest a threshold Θ enabling a node $\mathbf{v}$ to accept another node $\mathbf{b}$ as its parent, if the following equation holds:

$$\phi_{\mathbf{b}} + \kappa_{\mathbf{v},\mathbf{b}} \leq \phi_{\mathbf{v}} + \Theta \qquad (0 < \Theta < \kappa^*) \,. \tag{3}$$

Here, κ^* is the minimum cost measure for one-hop communication, e.g., $\kappa^* = 1$ for ETX. In contrast to a fixed value, Θ may also increase with the time t_{wait} that node $\mathbf{v}$ has already been waiting on its search for a suitable parent. Figure 2c shows the effect of employing (3) for the ETX metric with a constant Θ. Note that in order to use ETX values, nodes must estimate and possibly store link weights $\kappa_{\mathbf{v},\mathbf{b}}$. This could be done by converting RSSI or LQI values or by counting the number of transmissions vs. received acknowledging beacons per parent.

4.6 Properties of ORiNoCo

ORiNoCo is self-stabilizing and will end up in a *legitimate state* in connected topologies. Here, a legitimate state means that each packet is routed loop-free to the sink in a finite number of steps. However, a transient fault may put the system in an illegitimate state. In this situation, the proposed algorithm eventually returns to a legitimate state and is thus fault-tolerant. In the following, we prove the correctness of ORiNoCo's routing mechanism and the self-stabilizing property by showing convergence and closure regarding the legitimate state.

Lemma 1. *(correctness) Starting from any system state, each packet is routed to the sink.*

Proof. Only the sink s consumes packets, whereas all other nodes forward data to a destination with a lower weight. If a node v is a local minimum, i.e., $\phi_v \leq \phi_n, n \in N_v$, its weight is increased according to (1). Assume that a packet never reaches the sink. Since a packet never rests at one node there must be a set of nodes $F \subseteq V \setminus \{s\}$ among which the packet is relayed. Since each packet transmission ends up in a local minimum, which is increased, the minimum weight $\phi_F := min\{\phi_v | v \in F\}$ of F increases over time. Due to the connectivity property of the graph there are nodes adjacent to F and whose weights remain constant. A packet will be forwarded to an adjacent node v when $\phi_F > \phi_v$. This contradicts with the assumption that F is limited and doesn't contain the sink. Hence, each packet transmission must eventually end up in the sink. Due to the page limitation the extension of the proof for multiple packets is omitted.

Lemma 2. *(convergence) Starting from any state, a repetitive packet transmission leads to a legitimate state.*

Proof. During repetitive packet transmissions all weights are gradually adapted, so that all packets can be loop-free forwarded to the sink. Let S_i be the set of nodes with a distance of i from the sink. First all nodes in S_1 build up valid paths to the sink s according to (1). Note that depending on the weighting metric, a successor of S_1 can be a node of S_1 itself. A packet sent by a node S_1 to the sink can still run into a local minimum, but after some time all such local minima are removed (cf. Lemma 1). After this all paths from S_1 with decreasing weights end up in the sink. The statement of the lemma follows by induction.

Lemma 3. *(closure) A legitimate state is never left under error-free system execution.*

Proof. According to Lemma 2 a legitimate state is reached. In this state all possible paths are leading to the sink. This state is not affected by the algorithm any longer and the system thus stays in this state.

ORiNoCo differs from most self-stabilizing approaches in its execution, which is carried out during the routing process and thus no dedicated state announcement is needed. Self-stabilization is beneficial if the algorithm is applied in networks with large link-quality dynamics. ORiNoCo inherently supports multi-sink routing. Here, neither modification nor additional communication is necessary since sinks must only transmit beacons of zero path weight. Moreover, it is even possible to move sinks. It has to be noted that the initialization phase is not mandatory for the execution of the algorithm. It's sole purpose is to provide a fast convergence after an initial deployment.

The self-stabilizing property of the algorithm in case of a fault is exemplarily illustrated in Fig. 3. The left image depicts a ring graph with the hop-count weight shown next to the nodes ($\mathbf{v}_0$ is the sink). Due to environmental changes, the link between $\mathbf{v}_5$ and $\mathbf{v}_0$ breaks. When trying to transmit its next packet, $\mathbf{v}_5$

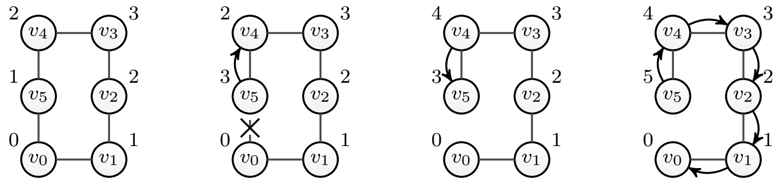

Fig. 3: Example showing the self-stabilizing property of ORiNoCo

does not receive any beacon satisfying (2) and starts the recovery strategy by forwarding the packet to an available neighbor, which is in this case node $\mathbf{v}_4$. The updated weights are shown in the second image. Node $\mathbf{v}_4$ tries to forward the packet, but subsequently starts the recovery mechanism, since it does not receive a beacon of weight 1. In the worst case, the packet is sent back to node $\mathbf{v}_5$ as shown in the third figure. $\mathbf{v}_5$ will have to update its weight again and sends the packet back to $\mathbf{v}_4$. Finally, the implicit routing structure is valid again and can be used for routing the packet to the sink (right image).

4.7 Sleep Time Adaptation

Introducing energy-awareness to ORiNoCo is only one step away. Nodes may adapt T_{slp} according to their energetic condition. Since senders pick the first available parent for relay, a node with larger T_{slp} sends less beacons and is thus less frequently used as parent. In traditional sensor networks, nodes may consider their residual energy for sleep time adaptation. In modern, energy-harvesting sensor networks [9], it is even possible to consider future energy prospectives: Energy intake predictions [16] may improve network throughput, reliability, and lifetime: However, sleep time adaptation is a possible extension only, and choosing the right adaptation scheme is a challenging task. Due to space constraints, we do not explore this complex field in this paper.

5 Theoretical Analysis

This section develops theoretical models for ORiNoCo to quantify and assess the simulation results presented in our evaluation.

5.1 Expected Average Waiting Time

The probability density function for receiving a beacon from node $\mathbf{b}$ with sleep times chosen randomly from $[(1-\alpha)\cdot T_{\text{slp}}, (1+\alpha)\cdot T_{\text{slp}}]$ at time t is

$$f_T(t) = \begin{cases} \frac{1}{T_{\text{slp}}} & 0 \leq t \leq (1-\alpha)\cdot T_{\text{slp}} \\ \frac{1}{2\alpha}\left(\frac{1+\alpha}{T_{\text{slp}}} - \frac{t}{T_{\text{slp}}^2}\right) & (1-\alpha)\cdot T_{\text{slp}} \leq t \leq (1+\alpha)\cdot T_{\text{slp}} \\ 0 & \text{else} \end{cases} . \tag{4}$$

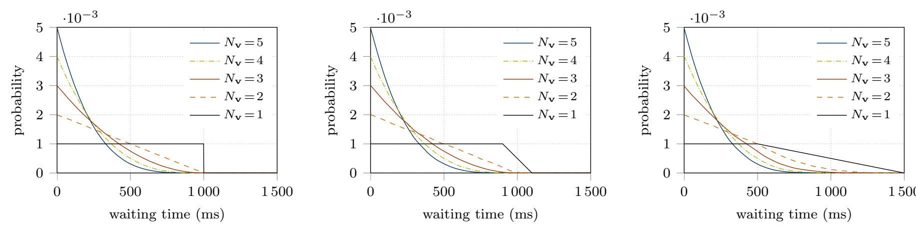

(a) $T_{\mathrm{slp}} = 1\,000\,\mathrm{ms},\ \alpha = 0.0$ (b) $T_{\mathrm{slp}} = 1\,000\,\mathrm{ms},\ \alpha = 0.1$ (c) $T_{\mathrm{slp}} = 1\,000\,\mathrm{ms},\ \alpha = 0.5$

Fig. 4: Simulated ideal distribution (no collisions) of beacon waiting times

The expected average waiting time is $E[T] = \left(0.5 + \alpha^2/6\right) \cdot T_{\mathrm{slp}}$. The fact that $E[T] > T_{\mathrm{slp}}/2$ is known as the *hitchhiker's paradox*.

For small α we find $E[T] \to T_{\mathrm{slp}}/2$ and can thus derive a lower bound for the expected average waiting time for $N_{\mathbf{v}}$ parents satisfying (2) with common T_{slp}:

$$E[\min(T_0, \ldots, T_{N_{\mathbf{v}}-1})] \geq \frac{T_{\mathrm{slp}}}{N_{\mathbf{v}}+1}\,. \tag{5}$$

Allowing a node to choose from multiple parents decreases beacon waiting time, energy consumption, and transmission latency. Figure 4 shows the actual waiting-time distributions for different values of $N_{\mathbf{v}}$ and α. The plots indicate that for $N_{\mathbf{v}} \geq 2$, α has no visible influence on the distribution. Hence, (5) is used as an approximation.

5.2 Power Consumption

In the following, we derive a simplified model of the power consumption of ORiNoCo. The model does not include packet loss, collisions, and backoff, and it assumes low traffic rates, i.e., no packets are queued.

Base-load power consumption $\bar{P}_{\mathrm{idle}}$ is defined by averaging the power for sleeping (P_{slp}), sending a beacon (P_{tx}), and waiting for a reply (P_{rx}):

$$\bar{P}_{\mathrm{idle}} = \frac{T_{\mathrm{slp}} \cdot P_{\mathrm{slp}} + T_{\mathrm{beac}} \cdot P_{\mathrm{tx}} + T_{\mathrm{dwell}} \cdot P_{\mathrm{rx}}}{T_{\mathrm{slp}} + T_{\mathrm{beac}} + T_{\mathrm{dwell}}} \tag{6}$$

using their average timings T_{slp} (cf. Sect. 5.1), T_{beac}, and T_{dwell}.

Let $\lambda_{\mathbf{v}}^{\mathrm{in}}$ denote the average incoming data rate of node $\mathbf{v}$ (incoming refers to packets received from other nodes). The mean power consumption $\bar{P}_{\mathrm{rx}}^{\mathbf{v}}$ of $\mathbf{v}$ for receiving packets is given by

$$\bar{P}_{\mathrm{rx}}^{\mathbf{v}} = \lambda_{\mathbf{v}}^{\mathrm{in}} \cdot (T_{\mathrm{pkt}} \cdot P_{\mathrm{rx}} + T_{\mathrm{beac}} \cdot P_{\mathrm{tx}})\,. \tag{7}$$

Note that the acknowledging beacon sent after packet reception (cf. Sect. 3.2) serves as invitation for subsequent transmissions in the same cycle.

The outgoing data rate $\lambda_{\mathbf{v}}^{\mathrm{out}}$ is the sum of the rate of self-generated packets $\lambda_{\mathbf{v}}^{\mathrm{sense}}$ and $\lambda_{\mathbf{v}}^{\mathrm{in}}$. The mean transmission power $\bar{P}_{\mathrm{tx}}^{\mathbf{v}}$ for node $\mathbf{v}$ is

$$\bar{P}_{\mathrm{tx}}^{\mathbf{v}} = \lambda_{\mathbf{v}}^{\mathrm{out}} \cdot \left(T_{\mathrm{pkt}} \cdot P_{\mathrm{tx}} + T_{\mathrm{beac}} \cdot P_{\mathrm{rx}} + \frac{T_{\mathrm{slp}}}{N_{\mathbf{v}}+1} \cdot P_{\mathrm{rx}}\right). \tag{8}$$

The right-most summand reflects the mean energy consumption before beacon reception, cf. (5). Overall power consumption is the sum of $\bar{P}_{\text{idle}}$, $\bar{P}^{\mathbf{v}}_{\text{rx}}$, $\bar{P}^{\mathbf{v}}_{\text{tx}}$. More traffic implies a larger impact of $\bar{P}^{\mathbf{v}}_{\text{tx}}$. From (8) it follows that having more than one parent effectively cuts down the waiting cost, because $T_{\text{slp}} \gg T_{\text{beac}}, T_{\text{pkt}}$. We evaluate the impact in Sect. 6.2 and 6.4.

6 Evaluation

This section evaluates the performance of ORiNoCo using the OMNeT++ network simulator. We compare the results with the findings of Sect. 5, CTP, and a strict tree routing protocol. The latter is a modified version of ORiNoCo and will be referred as 1-parent routing in the following. This protocol selects a fixed parent for packet forwarding. A node $\mathbf{v}$ updates this parent only if this decreases the weight $\phi_{\mathbf{v}}$ or if a recovery is needed (cf. Sect. 4.3).

The evaluation accounts for the key concerns of low-rate data-collection: energy efficiency, responsiveness, and reliability. To assess energy efficiency, we measured the energy consumption during the execution of the protocols by tracking the states of the transceiver hardware. Packet delays are indicators for the responsiveness. We distinguish between *1-hop delay*, which is the amount of time between the begin and the successful end of a packet transmission using RI-MAC, and the *end-to-end delay*, which is the time needed for sending a packet from the source to the sink. We recorded the number of lost packets to assess reliability.

6.1 Simulation Environment and Setup

Power consumption is modeled after an ATmega128RFA1 transceiver: $P_{\text{slp}} = 6\,\mu\text{W}$ in low-power mode, $P_{\text{rx}} = 25\,\text{mW}$ when listening/receiving, and $P_{\text{tx}} = 29\,\text{mW}$. Beacons and data packets have a size of 25 and 72 byte leading to $T_{\text{beac}} = 0.78\,\text{ms}$ and $T_{\text{pkt}} = 2.25\,\text{ms}$ for a transfer rate of 250 kbit/s. We set $T_{\text{slp}} = 2\,500\,\text{ms}$, $\alpha = 0.1$, and $T_{\text{dwell}} = 10\,\text{ms}$. The latter is the average dwell time for the backoff algorithm.

The following network parameters were chosen according to typical sensor network deployments. The connected topologies consist of 40 to 200 nodes, randomly dispersed in an area of $150 \times 150\,\text{m}^2$. The average node degree rises linearly with the number of nodes. Traffic generation follows a Poisson distribution with an inter-arrival-time of 5 s for the whole network—e.g., for 200 nodes the packet rate per node is $\lambda_v^{\text{sense}} = (5\,\text{s} \cdot 200)^{-1}$. For each network size, we executed 50 runs (10 topologies with 5 different random seeds) with a simulation time of 10 000 s.

For the simulations we used two different propagation models: A simplified physical layer, which does not consider packet loss, uses a unit-disk graph with 40 m communication range. Additionally, we used the IEEE 802.15.4 physical layer from the MiXiM framework. We chose a path loss exponent of 3.5, a log-normal shadowing mean of 0.5 dB, and a standard deviation of 0.25 dB. Reliable communication is possible up to 40 m, the transitional region [22] stretches additional 10 m. Signal interference and packet loss is considered with high accuracy. For both physical layers we calculated an average node density of 6.7 up to 34 for network sizes of 40 and 200 nodes respectively.

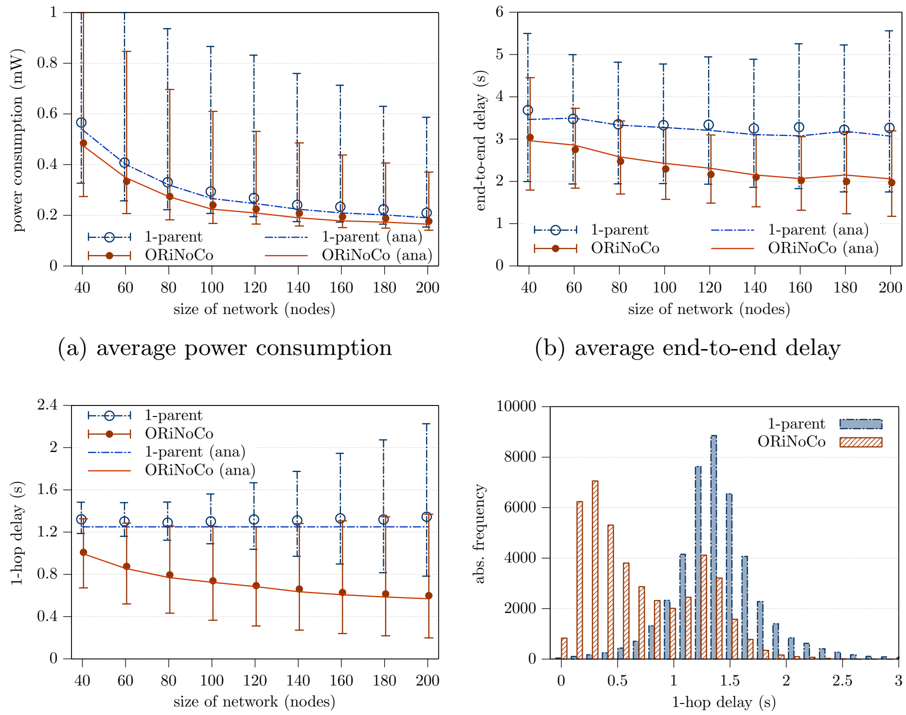

(a) average power consumption (b) average end-to-end delay

(c) average 1-hop delay (d) 1-hop delay distribution (200 nodes)

Fig. 5: Comparison of theoretical analysis and simulation. Error bars are averaged upper and lower deviations of single simulation runs.

6.2 Validation with Theoretical Analysis

A validation of ORiNoCo is done by a comparison with the analysis of Sect. 5. The power consumption, waiting time, and the end-to-end delay were calculated for the settings and topologies given in Sect. 6.1. The restriction to a fixed parent gives the approximation for 1-parent routing. In the simulation the simplified physical layer and the hop count metric are used. For each single simulation run the average value, upper, and lower standard deviation are calculated. These results are averaged over all runs with same network settings.

Figure 5a shows the average power consumption per node. Although ORiNoCo and 1-parent routing are using the same RI-MAC layer, power consumption of ORiNoCo is lower. The decrease of the power consumption for a growing network is due to the reduced load per node. For 200 nodes, ORiNoCo and 1-parent routing consume 0.18 mW and 0.21 mW, respectively. The large standard deviation supports that neighbors of the sink have to cope with the highest load and thus have to spend the most energy. The theoretical analysis is similar to the simulation result due to the fact that no packet loss is simulated.

The responsiveness of the protocols is evaluated using 1-hop and end-to-end delays depicted in Fig. 5b and 5c. ORiNoCo achieves considerably lower latencies than 1-parent routing. The reduced 1-hop delay, i.e., waiting time, is in tune with the decreased power consumption of ORiNoCo. Figure 5c also shows that for 1-parent routing the 1-hop delay is slightly higher than the value of the analytical approximation, since hitchhiker's paradox is not considered (cf. Sect. 5.1).

Figure 5d shows the distribution of 1-hop delays for the simulations of networks of size 200. ORiNoCo has two peaks: The one around 1.3 s is due to the fact that each packet has to be transmitted to (exactly) one sink, so that waiting time cannot be reduced on the last hop. The second peak at 0.4 s discloses the possible performance of ORiNoCo in multi-sink environments. As expected from Fig. 5c 1-parent routing has its mean around 1.3 s. However, delay times may become greater than T_{slp}. The reason is that a node, which currently forwards data, defers broadcasting beacons.

6.3 Comparison with CTP under Realistic Conditions

This section compares ORiNoCo with CTP under realistic conditions using MiXiM's physical layer. CTP uses a CSMA leading to a power consumption of at least 25 mW with our hardware. Hence, we replaced it with RI-MAC. This permits a fair comparison of the two protocols in terms of power consumption.

In CTP each node broadcasts control packets (network beacons) in order to build up a routing tree. The initial rate is given with 64 ms and is later reduced according to the Trickle algorithm [7]. This approach is poorly collaborating with low-power MAC protocols, e.g., BoX-MAC, RI-MAC. For the latter, a broadcasting node stays listening for an interval of $(1 + \alpha) \cdot T_{\mathrm{slp}}$ time units and replies with a broadcast to all received beacons. Especially in dense networks this leads to high congestion. Based on the results of several simulations, we set the control message rate of CTP to 50 s and applied a warm-up period of 200 s before generating any traffic. CTP uses the expected transmissions (ETX) as cost metric. The actual ETX values are pre-calculated to remove the influence of link-quality estimation. ORiNoCo filters beacons according to their LQI value.

The results are shown in Fig. 6a, 6b, and 6c. In general, the results for ORiNoCo and 1-parent routing comply with the characteristics depicted in Fig. 5. Nevertheless, a small increase of power consumption is measured for all topologies. Also the delay times are slightly higher with the MiXiM physical layer. For 200 nodes the 1-hop delay of ORiNoCo is 0.69 s, whereas 0.60 s is measured for the simplified physical layer. This degradation is much smaller compared to 1-parent routing which experiences an increase from 1.44 s to 1.81 s. As a result, ORiNoCo is doing fine in error-prone, dense networks.

CTP shows similar results as 1-parent routing in terms of 1-hop delay, since both are exploiting the same RI-MAC layer. However, the performance of the end-to-end delay of CTP drops significantly. The reason is the non-optimal selection of the routing tree due to packet loss and the non-optimization for low-power MAC protocols. It has turned out that despite of the few control messages, the power consumption of CTP is much higher than of ORiNoCo.

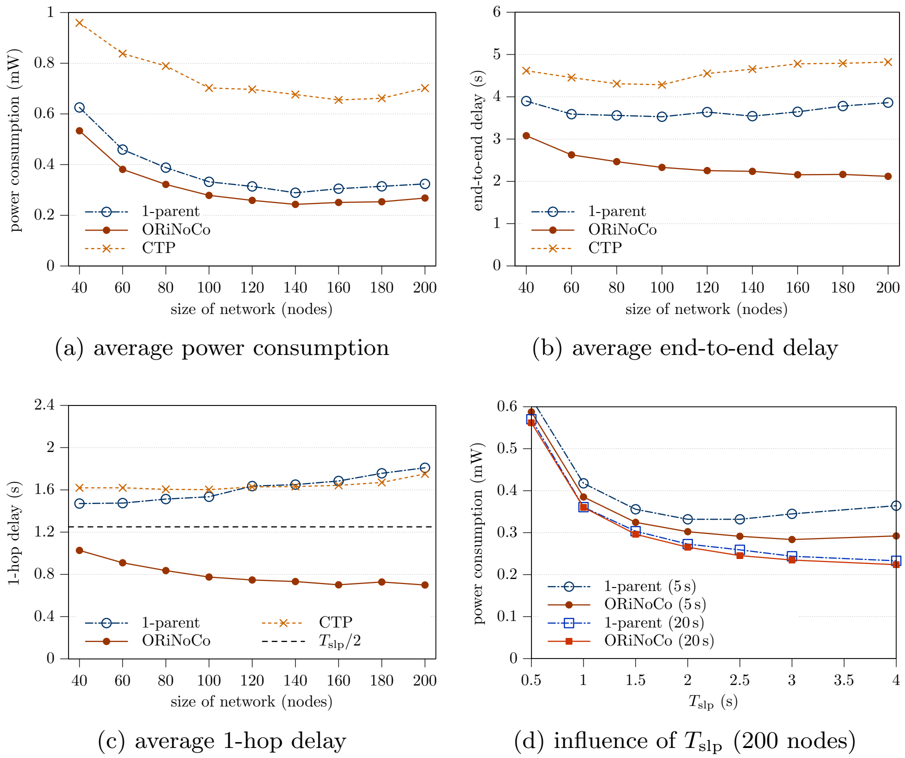

Fig. 6: Comparison of ORiNoCo, CTP, and 1-parent routing using MiXiM

We observed no end-to-end packet loss in the simulations. The RI-MAC layer uses acknowledgment beacons for ensuring successful transmissions. However, in heavy loaded networks packet loss may occur due to limited queue sizes.

6.4 Dimensioning of T_{slp}

Finally, we examined the influence of T_{slp}. For the simulation the MiXiM physical layer is used and the inter-arrival-time for the network traffic is set to 5 s and 20 s. The value of T_{slp} is crucial since it determines the beacon rate and thus the energy consumption as well as the responsiveness of ORiNoCo.

The result is depicted in Fig. 6d. For small values of T_{slp}, the overhead caused by beaconing is the main energy consumer. For large values of T_{slp}, beaconing consumes less energy, but the costs of data transmission rises as shown in (8). These costs are proportional to the traffic rate. Because traffic and beaconing costs are showing a contrary behaviour, there is an optimal T_{slp} depending on the network parameters. A proper explanation of this effect can be found in [18]. As expected, ORiNoCo shows a lower energy consumption than 1-parent routing. This difference is larger in networks with higher load.

6.5 Discussion

ORiNoCo reveals itself as a promising alternative to existing collection protocols. Our evaluation shows that an opportunistic approach—in comparison to a fixed routing tree—significantly reduces latency and slightly decreases power consumption. However, we examined low data rates only. If running ORiNoCo in networks of high traffic load, power consumption is expected to rise. Yet, waiting time is expected to decrease, as multiple queued packets are sent in a burst transmission. Since an increased load leads to a higher collision probability, an extended evaluation is needed for a final assessment. In dense networks, ORiNoCo will suffer from beacon collisions; it is thus necessary to adapt T_{slp} to the density. This will not necessarily decrease agility, because an increased T_{slp} may cancel out with additional parents. In large networks, multiple sinks will increase throughput.

The results show that ORiNoCo is competitive with state-of-the-art collection protocols. It is more flexible, because it allows for selecting among multiple parents, increasing reliability and decreasing energy consumption and latency. Another advantage of ORiNoCo is its waiving of explicit broadcasts for route and neighborhood updates. Running CTP with RI-MAC—CTP natively uses CSMA—pollutes the channel with broadcasts, hindering the construction of a valid routing tree due to collisions. For this reason, even 1-parent routing could easily outplay CTP. Furthermore, CTP requires link-quality estimation—which is expensive in networks of high density—to select the most suitable parent. In contrast, ORiNoCo uses the LQI value of beacons as metric for link benchmarking. However, the accuracy of this metric depends on the hardware.

7 Conclusion

This paper introduced ORiNoCo, a novel data-collection protocol. Its tight integration into RI-MAC avoids additional overhead by exploiting the LPP mechanism to set up tree-like routes. ORiNoCo finds a good trade-off between power-efficiency and responsiveness, while maintaining high reliability due to its self-stabilizing property. In simulations of single-sink data-gathering scenarios it outperforms CTP and passes the tests with an average power consumption of less than 0.3 mW while providing an end-to-end latency below 2.5 s for a network of 200 nodes. This makes ORiNoCo highly suitable for event-reporting applications. Even though there is room for improvement. We have scheduled a deployment of ORiNoCo for the near future to verify its real-world practicability. Furthermore, optimizations on the backoff strategy and sleep time adaptation will be carried out. The latter will particularly focus on energy-harvesting sensor networks.

8 Acknowledgments

This research has been partially funded by the German Research Foundation under contract number TU 221/4-1. The authors would like to thank the anonymous reviewers for their valuable feedback.

References

1. Alizai, M.H., Landsiedel, O., Bitsch Link, J.A., Götz, S., Wehrle, K.: Bursty Traffic Over Bursty Links. In: SenSys (2009)
2. Buettner, M., Yee, G.V., Anderson, E., Han, R.: X-MAC: a Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks. In: SenSys (2006)
3. Burri, N., von Rickenbach, P., Wattenhofer, R.: Dozer: Ultra-low Power Data Gathering in Sensor Networks. In: IPSN (2007)
4. De Couto, D.S.J., Aguayo, D., Bicket, J., Morris, R.: A High-Throughput Path Metric for Multi-Hop Wireless Routing. In: MobiCom (2003)
5. Dutta, P., Dawson-Haggerty, S., Chen, Y., Liang, C.J.M., Terzis, A.: Design and Evaluation of a Versatile and Efficient Receiver-Initiated Link Layer for Low-Power Wireless. In: SenSys (2010)
6. Fonseca, R., Gnawali, O., Jamieson, K., Levis, P.: Four-Bit Wireless Link Estimation. In: HotNets VI (2007)
7. Gnawali, O., Fonseca, R., Jamieson, K., Moss, D., Levis, P.: Collection Tree Protocol. In: SenSys (2009)
8. Hsu, J., Zahedi, S., Kansal, A., Srivastava, M., Raghunathan, V.: Adaptive Duty Cycling for Energy Harvesting Systems. In: ISLPED (Oct 2006)
9. Jiang, X., Polastre, J., Culler, D.: Perpetual Environmentally Powered Sensor Networks. In: IPSN (2005)
10. Kim, S., Fonseca, R., Dutta, P., Tavakoli, A., Culler, D., Levis, P., Shenker, S., Stoica, I.: Flush: a Reliable Bulk Transport Protocol for Multihop Wireless Networks. In: SenSys (2007)
11. Landsiedel, O., Johansson, M.: Poster Abstract: Towards a Life without Link Estimation. In: SenSys (2010)
12. Moss, D., Levis, P.: BoX-MACs: Exploiting Physical and Link Layer Boundaries in Low-Power Networking. Tech. rep., Stanford Information Networks Group (2008)
13. Musaloiu-E., R., Liang, C., Terzis, A.: Koala: Ultra-Low Power Data Retrieval in Wireless Sensor Networks. In: IPSN (2008)
14. Polastre, J., Hill, J., Culler, D.: Versatile Low Power Media Access for Wireless Sensor Networks. In: SenSys (2004)
15. Puccinelli, D., Gnawali, O., Yoon, S., Santini, S., Colesanti, U., Giordano, S., Guibas, L.: The Impact of Network Topology on Collection Performance. In: EWSN (2011)
16. Recas Piorno, J., Bergonzini, C., Atienza, D., Simunic Rosing, T.: Prediction and Management in Energy Harvested Wireless Sensor Nodes. In: VITAE (2009)
17. Renner, C., Ernst, S., Weyer, C., Turau, V.: Prediction Accuracy of Link-Quality Estimators. In: EWSN (2011)
18. Renner, C., Unterschütz, S., Turau, V.: Power Management for Wireless Sensor Networks Based on Energy Budgets. Tech. Rep. urn:nbn:de:gbv:830-tubdok-11065, Hamburg University of Technology, Hamburg, Germany (Jul 2011)
19. Srinivasan, K., Levis, P.: RSSI is Under Appreciated. In: EmNets (2006)
20. Sun, Y., Gurewitz, O., Johnson, D.B.: RI-MAC: A Receiver-Initiated Asynchronous Duty Cycle MAC Protocol for Dynamic Traffic Loads in Wireless Sensor Networks. In: SenSys (2008)
21. Whitehouse, K., Woo, A., Jiang, F., Polastre, J., Culler, D.: Exploiting the Capture Effect for Collision Detection and Recovery. In: EmNetS-II (2005)
22. Zuniga, M., Krishnamachari, B.: Analyzing the Transitional Region in Low Power Wireless Links. In: SECON (2004)