– Technical Report 377/10 –

On the Behavior of Black Bursts in Tick-Synchronized Networks

Dennis Christmann (christma@cs.uni-kl.de)

Department of Computer Science
University of Kaiserslautern

June 14, 2010

Abstract

In wireless networks, many existing protocols suffer from destructive collisions of simultaneously sent frames, influencing their quality-of-service and their application in real-time scenarios. With *black bursts*, this paper investigates a relatively unexplored method for the collision-protected transmission of information in wireless networks. Although such transmissions are slower and less efficient than transmissions with regular frames, their robustness enables advanced and deterministic protocols, e.g., synchronization with deterministic convergence delay, deterministic medium arbitration, or negative-acknowledgments of broadcast transmissions. This paper explores black burst-based transmissions theoretically as well as experimentally, and analyzes their requirements regarding time. In particularly, we extend previous work by presenting detailed derivations of formulae used by two black burst-based transfer protocols, called *cooperative transfer* and *arbitrating transfer*. By means of two experiments, we evaluate the detection of black bursts based on the Clear Channel Assessment (CCA) mechanism and show their applicability under real-world conditions. Finally, we discuss differences in the transmission ranges of regular MAC frames and black bursts and show how to bring them in line.

Contents

1 Introduction

In wireless ad-hoc networks, many protocols suffer from the problem of destructive frame collisions, destroying information and wasting bandwidth as well as energy. Especially in dense networks, problems due to collision-prone transmissions become crucial and make guarantees regarding real-time behavior and quality-of-service almost impossible. When looking at multi-hop networks, the situation gets even worse. Several mechanisms – like CSMA/CA, several classes of inter frame spacings, or RTS/CTS handshakes [14] – have been devised that indeed help to decrease frame collisions at contention-based medium access. Nevertheless, they cannot avoid collisions completely. Thereby, end-to-end delays in contention-based networks are hardly predictable, and deterministic guarantees cannot be given without further measures.

Compared to wired networks, detection of and dealing with frame collisions differs in wireless networks completely. In wired networks, on the one hand, listening to the medium while transmitting data is easily possible and e.g., implemented in CSMA/CD. On the other hand, wired networks provide also techniques enabling concurrent, non-destructive transmissions. E.g., in the successful CAN bus [19], which is particularly found in the automotive field, identifiers are transmitted concurrently and used for deterministic medium arbitration. Thus, the medium access can be granted based on priorities of the identifiers, so that only the station with the highest priority is allowed to transmit the actual data. Since in general, stations of wireless networks are not able to transmit and listen to the medium simultaneously, the principle of CAN can not be transferred to the wireless domain in a straightforward way.

With *arbitrating transfer* and *cooperative transfer*, this paper proposes two black burst-based protocols on MAC layer with deterministic transfer delays. Arbitrating transfer enables network-wide, priority-based medium arbitration similar to CAN but in the domain of wireless networks. Cooperative transfer, on the other hand, provides the network-wide propagation of a bit sequence within guaranteed time bounds, usually initiated by a single node. Since in both protocols, transmissions can proceed concurrently, the transfer delay does not depend on the number of nodes but on the maximal network diameter only. Both protocols rely on tick synchronization with deterministic guarantees regarding tick offset as e.g., provided by *Black Burst Synchronization* (BBS) [10].

Cooperative as well as arbitrating transfer are primarily designed for wireless sensor networks. Their behavior was specified with the Specification and Description Language (SDL) [13] and implementations of black bursts exist for MICAz [7] and Imote2 [8] platforms, both equipped with TexasInstruments' CC2420 transceiver [5]. But the implementations are not limited to those platforms, since the only hardware requirement is the availability of Clear Channel Assessment (CCA) that is necessary for CSMA/CA and required by all 802.15.4 compliant nodes [12].

The remainder of the paper is structured as follows: First, Sect. 2 presents an extensive survey of related work. Then, Sect. 3 introduces the concept and realization of black bursts. In Sect. 4, we present cooperative and arbitrating transfer and derive their timing behavior. Thereafter, Sect. 5 provides insights in the experimental evaluation of black bursts. Finally, Sect. 6 draws conclusions.

2 Related Work

In recent years, novel wireless MAC protocols have been developed, exploiting the fact that information may not be encoded in the regular payload of data frames only. Instead, information can also be encoded in start and duration of medium occupancy, which can also be decoded if several transmissions overlap, enabling the transmission of information in a collision-protected way. Especially MAC protocols aiming for quality-of-service and real-time support are found in this field of research. Most of the protocols are related to the arbitrating transfer protocol (Sect. 4.1); for cooperative transfer (Sect. 4.1), no related work could be found.

In the following, we give a survey of some of those wireless protocols and discuss the origins of black bursts. The protocols are categorized in schemes based on *buzz* signals of variable length and binary encoding-based approaches. A third category contains all proposals that do not fit in neither of the two categories.

2.1 Jamming-based Protocols

To our knowledge, the term black burst was first introduced into protocol terminology by Sobrinho and Krishnakumar [29]. In their work, black bursts represent jamming signals of different length and are used to guarantee bounded access delays to real-time traffic in IEEE 802.11 networks [11]. Their approach is compatible with existing IEEE 802.11 implementations and distinguishes between extended real-time stations and regular data stations complying with the standard IEEE 802.11 behavior. Real-time stations are assumed to send data frames in periodical intervals of equal length. Frames are scheduled a constant duration after a previous data frame successfully obtained medium access. Thus, after an initialization phase, different stations always schedule transmissions at different points in time. The arbitration is performed by jamming the medium with black bursts consisting of several *black slots*, whose number is chosen proportionally to the time a station with data to send has been waiting for the channel to become idle. The station sending the longest black burst is allowed to transmit the data in a regular transmission subsequently. Stations sending shorter black bursts detect stations with longer black bursts by performing CCA after their transmission is finished. This results in a fair global scheduling following a round-robin order for real-time nodes.

The presented black burst-based medium access was extended in [28] and marginally discussed for multi-hop networks without hidden stations. Since the approach guarantees a bounded access delay as long as the number of real-time nodes does not exceed a computable threshold, it can also guarantee a bounded end-to-end delay. To reduce overhead due to black burst-based arbitration, the authors propose a chaining mechanism that is a sequence of stations, where each station invites a next station for transmission. Although, this measure reduces contention with black bursts within one communication range, it does not address multi-hop communication.

Stability of the resulting approach is formally proven both in [29] and with chaining in [28]. The simulative evaluations demonstrate high improvements compared to the IEEE 802.11 standard in terms of throughput and real-time packet delay. But it is assumed that the proposed protocol does not perform well in multi-hop networks or if real-time stations want to send data in irregular intervals. Especially in multi-hop networks with hidden stations, collisions can occur, since nodes compete for medium access at different points in time. Furthermore, priorities, i.e. the length of black bursts, depend only on the waiting time and can not be assigned according to application-specific metrics.

In [20], Pal et al. present two algorithms for single-hop transmissions in IEEE 802.11 networks, called *elimination by sieving DCF* (ES-DCF) and *deadline bursting DCF* (DB-DCF). Both algorithms distinguish between non-real-time nodes and real-time nodes executing different protocols that are non-compatible to the IEEE 802.11 standard. Before accessing the medium with ES-DCF for transmission of a real-time frame, a node must observe the channel to be idle for a duration that depends on the packet's deadline, the associated priority (two levels are supported: hard vs. soft real-time), and a generated random number. If the channel is still idle after waiting, the node initiates an RTS/CTS handshake, which must be completed successfully before transmitting the actual data.Since waiting times of contending stations may be equal, collisions of RTS/CTS frames can not be precluded, leading to a failure of the RTS/CTS handshake. In such conflict situations, black burst-based arbitrations are applied, where the length of black bursts is set according to unique node ids to guarantee a single winner node. Instead of waiting a random backoff interval, DB-DCF applies an adapted version of the black burst-based arbitration of Sobrinho and Krishnakumar before a station is allowed to initiate an RTS/CTS handshake. In the adapted version, the length of black bursts is set inversely proportional to the deadline of the real-time packet to be sent. Thereby, real-time packets in urgency are preferred, yet there is no guarantee for a unique winning node, leading to possible collisions of subsequent RTS/CTS frames. To solve this conflict, competing nodes execute black burst-based arbitration a second time with the length of black bursts proportional to the node's unique id.

The proposed protocols describe a very optimistic approach, in which collisions are not avoided a priori. Thus, energy is wasted and worst case delay is increased. Although RTS/CTS frames are used to alleviate the hidden station problem, it has to be assumed that the algorithms perform badly in multi-hop scenarios, because the proposed arbitration with black bursts is intended for single-hop networks only.

In [27], the authors present a distributed MAC protocol for single-hop networks, in which access to the medium is granted according to static, non-unique priorities. The medium arbitration consists of two phases: First, nodes compete with the black burst-based protocol presented in [29], where the length of a black burst depends on the node's priority. As result, a set of nodes holding the highest priority remains in contest. Secondly, conflicts between equal-prioritized nodes are solved by enforcing a round robin schedule. Thereby, all winner nodes of the first phase are allowed to access the medium before lower priority nodes may compete for the medium again. The round robin schedule requires an ordering of stations that is achieved by a modified version of the randomized ID initialization protocol [18]. Since the randomized initialization protocol requires transmitting nodes to simultaneously detect collisions, which is hardly possible in wireless networks, the proposed MAC protocol depends on an elected leader node acting as virtual collision detector. Because the randomized initialization protocol is a probabilistic approach, the termination of the second arbitration phase comes without guaranteed time bounds.

In [4], a decentralized MAC protocol for unsynchronized networks called *Real-Time Chain* is presented. It extends existing MAC protocols based on a CSMA/CA medium access by the support of periodical real-time traffic. The protocol is based on the presented black burst-based arbitration scheme [29] with extensions for multi-hop networks. The proposed protocol mainly addresses sensor networks and was experimentally evaluated on MICAz motes [7]. Before real-time traffic is transmitted, the source node must initiate a *chain opening* based on the black burst arbitration that informs all nodes on the route about the upcoming transmission. Chain opening as well as transmission of ordinary best effort traffic are done in one common radio channel. Real-time traffic is transmitted on dedicated radio channels, where every second node on the route from source to destination changes the channel when forwarding real-time data. Thereby, the probability of collisions of regular data frames are reduced and throughput is increased. To improve throughput further, real-time traffic is transmitted following the principle of pipelining. Based on the chain opening packet, a node determines on the one hand the receiving/transmission channel of real-time data. On the other hand, the chain opening packet is also used by a node to determine the priority of the black burst that is used to solve contention between nodes on the same dedicated radio channel. Thus, in contrast to [29], the length of black bursts does not depend on the waiting time but on the position within the route between source and destination. Thus, real-time data is forwarded smoothly within a chain without the need of long buffering times on intermediate nodes. If the source node has no further data to send, a *chain close* packet is sent along the route, letting nodes change back to the common radio channel.

Since Real-Time Chain targets multi-hop networks but does not provide synchronization at all, the authors report on problems about priority inversion and even starvation of high priority real-time traffic. They solve this problem by the transmission of a long jamming signal, if the channel is continuously observed busy. This indeed disposes nodes with low-priority traffic to temporarily stop further access to the medium, but also wastes bandwidth and energy. Furthermore, nodes can only participate in one chain at the same time, i.e. they are not available for further chain openings. Additionally, bandwidth is wasted if a long-living chain with low throughput is established.

The *Busy Tone Priority Scheduling* (BTPS) protocol [32] is slightly different to the already presented black burst-based approaches, since contention between high-priority nodes is not solved by the encoding of priorities in the length of jamming signals. It was designed as improvement of IEEE 802.11 for multi-hop networks and distinguishes between low and high priority traffic. It utilizes two types of narrow-band busy tone signals to prevent nodes with low priority data from transmission. One type of busy tone ($BT1$) is periodically sent by all stations with high priority data to send and the second one ($BT2$) is periodically sent by stations receiving the first type of busy tone. Thus, also hidden stations are informed about an upcoming high priority transmission. Busy tones are sent on dedicated sub channels of a regular IEEE 802.11 channel, thereby, reducing the overall transmission rate of the actual payload. Since stations with low priority data must not access the medium when overhearing a busy tone, contention takes only place between nodes with high priority data. To avoid collisions between high priority nodes, they choose a random backoff interval in which the data channel must be observed idle. Since nodes periodically send the busy tone, the actual data channel is only sensed if currently no busy tone (either $BT1$ or $BT2$) is to be sent. Thus, the data channel is not observed all the time but periodically. Data frames are sent with preceding RTS/CTS frames and confirmed with an ACK frame on success. Since an intended receiver can not listen to the data channel and transmit a busy tone simultaneously, the sending nodes transmit a black burst of constant length to let all nodes listen for incoming RTS frames. After transmission is complete (either by observing an ACK frame or by waiting for timeout), all stations contend for medium access again if data frames are available.

Because BTPS follows a probabilistic approach by choosing random backoffs, collisions can still occur between stations transmitting real-time traffic. Thus, BTPS can not provide time bounds for the delivery of real-time data. In addition, the introduced busy tones require special hardware support to separate the three sub channels. Since the proposed busy tones are only aimed for prioritizing single-hop transmissions, arbitration must be performed on every hop if data should be transmitted over several hops.

2.2 Binary Encoding-based Protocols

Synchronized MAC (SYN-MAC, [31]) was contemplated for IEEE 802.11 technology. It requires synchronization with high accuracy to divide time into equal-sized *time frames*. The authors do not address synchronization in detail and assume synchronization to be obtained by GPS or the cellular infrastructure. SYN-MAC is based on the *binary countdown scheme*, whose general idea is to assign a binary number of fixed length to each station ready to transmit and to declare the node with the highest number as winner. The binary countdown scheme is applied in the contention interval at the beginning of each time frame to perform medium arbitration. The binary numbers used as arbitration sequence are generated randomly by each node with data to be sent and are transmitted bit-wise at aligned points in time. Each binary 1 is realized by the transmission of a contention signal (regular MAC frame) that contains the destination's MAC address; each binary 0 corresponds to no transmission. Since multiple stations can process a 1 concurrently, collisions can occur and the payload of the contention signals gets lost. Nevertheless, colliding contention signals can still be detected by CCA. Each station that is processing a 0 and observes a busy medium – either by CCA in case of collisions or by the reception of a contention signal – gives up arbitration in this contention slot. In topologies with hidden stations or when multiple nodes generate the same random number, several stations can complete their arbitration sequence and rate themselves as winner. The first source of conflict is solved in a hidden station elimination interval subsequent to the contention interval. Conflicts due to identical arbitration sequences are not caught by SYN-MAC, justified by the low probability of equal arbitration sequences, if the size of the sequence is large.

In [21], the authors present an enhanced version of SYN-MAC. Different to the initial version, a station receiving a contention signal while contending with a 0, does not always stop contending directly. Instead, the node looks at the destination address specified in the received contention signal, and stops medium arbitration only if it is directly connected to the specified destination node. Otherwise, it continues with arbitration, leading to a possible better exploitation of spatial reuse.

Since time frames have constant length, data frames with variable length lead to a waste of bandwidth. Moreover, the external synchronization requires additional hardware and is not always applicable, e.g., in indoor environments. Although the authors argue that no network-wide synchronization is required and that synchronization within one collision domain is sufficient [21], it is questionable how the hidden station problem is alleviated when hidden station elimination intervals may not overlap in time. Another drawback comes when looking at multi-hop transmissions, because medium arbitration must be performed on every hop, increasing protocol overhead and end-to-end delay. Furthermore, the authors of [22] discovered the binary countdown arbitration to be employed such that collisions can cause misses of deadlines.

With Carrier Sense Media Access with ID countdown (CSMA/IC) [33], the authors present a MAC protocol for multi-hop networks following the binary countdown scheme, too. Synchronization is periodically achieved between neighboring nodes by the transmission of beacons. No global synchronization is required. For arbitration, every station contends bit-by-bit with a unique bit sequence by transmitting a *buzz* signal if the current bit is 1 and by listening to the medium if the current bit is 0. When receiving a 1 while listening to the medium, a node stops competing. Thus, compared to SYN-MAC [31], a binary 1 carries no regular payload. The unique bit sequence consists of the priority of the packet, favoring important data, and a unique node id, ensuring a single winner. Hidden station problems are ignored completely. The authors justify this decision by the simplified assumption of a carrier sensing range with doubled radius compared to the transmission range. Thereby, *hidden stations* (in terms of transmission range) with low priority lose during arbitration and do not cause collisions.

In multi-hop scenarios, the protocol suffers from successive losing of nodes over several hops. Thus, bandwidth is wasted, since spatial reuse is decreased. In addition, the beacon-based synchronization protocol does not offer guarantees regarding accuracy or convergence delay. Therefore, nodes can run out of synchronicity especially in dense networks. Though the authors report on a 100% collision-free medium access in their simulative evaluation, this result will probably not hold under real-world conditions, where hidden stations can not be precluded.

BitMAC [26] is aimed to guarantee deterministic convergence delays in typical Wireless Sensor Networks (WSNs) consisting of spanning trees, in which data flows from the sensors (leafs) to the sink (root) over multiple hops. To avoid interferences between different sub trees, BitMAC assumes transceivers to support a sufficient number of channels to allow the assignment of different channels to adjacent sub trees. BitMAC provides synchronization with deterministic accuracy and convergence delay, and applies the principle of an On-Off-Keying modulation similar to CSMA/IC (1 realized as transmission, 0 realized as no transmission at all). By doing so, a set of nodes within the same communication channel holding different values can agree on an aggregated value, e.g., obtained by an OR/AND/MIN/MAX operation, within a deterministic time bound. Among other things, this principle is utilized for synchronization, the allocation of channels, and the assignment of parent nodes. The main drawback is the limitation to tree structures, avoiding the application in ad-hoc networks without a designated root node.

WiDom [22] is also based on the binary countdown scheme and is designed to transfer the principle of the CAN bus [19] into the domain of wireless networks. In contrast to SYN-MAC [31], WiDom considers sporadic messages and does not rely on an external synchronization. In its initial version, it addresses single broadcast domains only. Ahead to every medium arbitration that is called *tournament phase*, WiDom performs a *synchronization phase* that brings all nodes in synchronicity. In the tournament phase, every node with data to transmit sends a unique, static, and message-dependent priority of fixed length, where the lowest number corresponds to the highest priority. The priority is sent bitwise in a way that allows stations to listen to the medium and detect a dominant 0 while processing a recessive 1. Stations stop competing as soon as a node with higher priority is detected. At the end of the tournament phase, the unique winner node can transmit the actual data frame with a conventional transmission. To face inaccuracy of the synchronization and hardware characteristics, constant guard times are considered in the tournament phase to enable a clear assignment of received bits to bit times.

In [23], WiDom was changed slightly to increase robustness, achieved by the rebroadcasting of dominant bits by receiving nodes. Although this increases the duration of the tournament phase, WiDom still guarantees deterministic time bounds for medium access depending on node priorities.

An extension of WiDom to multi-hop networks can be found in [24]. The presented multi-hop synchronization protocol is not intended to establish network-wide synchronization, yet it synchronizes at least all nodes in 2-hop range. With this cluster-based synchronization approach, the extended version of WiDom is independent of the network diameter and exploits spatial reuse, enabling simultaneous transmissions in different network areas. The protocol can not avoid the network-wide propagation of the synchronization tick, i.e. the convergence delay of synchronization is unknown. To face the hidden station problem, the extended version provides two hop arbitration that enables the subsequent one-hop transmission of the regular data frame. Thereby, arbitration must be performed on every hop for multi-hop transmissions, leading to an increased end-to-end delay. Although spatial reuse is one of the key goals of WiDom, the number of parallel transmissions is not maximized, since a node can lose arbitration against neighbored nodes which again lose against neighbored nodes. Similar problems were also reported in CSMA/IC [33].

2.3 Summary and Assessement

The outlined protocols can be summarized and categorized according to the following criteria:

- **Guarantees for medium access:** Some protocols (like WiDom [22]) provide deterministic time bounds according to the medium access. In other protocols following probabilistic approaches (like SYN-MAC [31]), no guarantees are given.
- **Encoding of information in a collision-protected way:** All protocols have in common that information is encoded in a collision-protected way. In some protocols (e.g., DB-DCF [20]), information is encoded in the length of a busy/buzz tone, where other protocols like BitMAC utilize binary encodings [26].

  Binary encoding schemes benefit from the logarithmic growth of their application's duration, when enlarging the number of allowed values to be transmitted. In contrast, the duration of busy tone approaches growth linearly with the number of allowed values. But with binary encoding schemes, additional overhead is introduced due to hardware-dependent characteristics (e.g., switching times from receive to send mode). Thus, binary encoded protocols are not *better* a priori.
- **Range of propagation:** Protocols differ in terms of the assumed topology. Protocols like the black burst-based arbitration introduced by Sobrinho [29] consider only single-hop networks. Other protocols (e.g., CSMA/IC [33]) include multi-hop networks but without hidden stations.

However, multi-hop networks with hidden stations are also considered, e.g., by BTPS [32]. Nevertheless, except Real-Time Chain [4], no protocol optimizes multi-hop transmissions of actual data frames.

- **Traffic pattern:** Many presented protocols are optimized for a specific traffic pattern. E.g., periodical traffic is assumed in [29], where WiDom can also deal with sporadic traffic [22]. Real-Time Chain [4] represents a special case, since it deals with short-time traffic with high throughput.
- **Assignment of priorities:** Priorities used for medium arbitration affect the length of black bursts and binary numbers, respectively. Protocols like [20] determine the priority based on node ids, where in WiDom [22], the priority is associated with messages. In other protocol, e.g., in [29], priorities are assigned according to deadlines.
- **Synchronization:** The surveyed protocols require at least one-hop synchronization, since decisions are based on time-related information. Some protocols (like SYN-MAC [31]) leave this problem completely out and assume external synchronization by GPS or the cellular infrastructure. Other protocols (like CSMA/IC [33]) provide local synchronization. Network-wide synchronization is only provided by BitMAC [26].

3 Principles of Black Burst Transmissions

This section outlines the foundations of cooperative transfer and arbitrating transfer, two protocols based on black bursts. In the first subsection, we give our definition of black bursts, which slightly differs from the original definition of Sobrinho et al. [29]. Since cooperative transfer as well as arbitrating transfer require an accurate and deterministic tick synchronization, Sect. 3.2 investigates this requirement in detail. Finally, we discuss the detection and decoding of black bursts in Sect. 3.3. Some of this section builds on previous work [16, 17].

3.1 Definition of Black Bursts

Black bursts are sent without prior medium arbitration and carry no regular payload [17]. Nevertheless, they include two important pieces of information: The length, represented by the duration of caused medium occupancy, and the starting point of transmission. Both pieces of information are also retained if several black bursts are transmitted (almost) simultaneously without changing the length of medium occupancy significantly. Thus, information can be encoded in a collision-protected way by the start and the duration of medium occupancy. In contrast to the initial definition of black bursts [29], one black burst following our definition encodes one bit only.

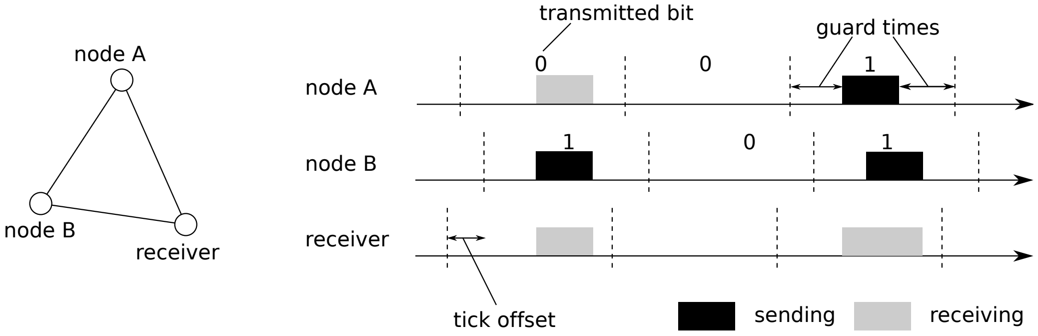

Figure 1: Principle of black bursts: Two senders (nodes A and B) with tick offset transmit sequences of black bursts concurrently. A receiver in sensing range detects the OR of both sequences.

Fig. 1 gives an example of the transmission of two bit sequences encoded with black bursts. Such bit sequences are also called *black burst frames*. In the example, node A transmits 001 and node B 101, respectively. Binary 0s are encoded recessively with black bursts of length 0 ("no transmission") and binary 1s are realized by dominant transmissions of black bursts with length > 0. Thus, nodes receive the OR of both concurrently transmitted black burst frames[1]. This includes also nodes transmitting their own bit sequence (like node A in Fig. 1), because they are able to sense the medium during processing recessive 0s.

Since two nodes intended to transmit black burst frames concurrently must observe a common starting point, synchronization is required. Due to tick offsets caused by inprecise clocks and inaccuracy of the synchronization, sending nodes must consider guard times to allow receiving nodes the unambiguous assignment of received black bursts to the corresponding time/bit interval. To ensure the correct assignment, tick offsets must be bounded deterministically. In Sect. 3.2, we provide a detailed discussion on tick synchronization. Guard times must also consider hardware characteristics (e.g., transceiver switching times) that are discussed in Sect. 3.3.

In our implementation, a black burst is realized with customary transceivers as a regular MAC frame of minimal length. With TexasInstrument's CC2420 transceiver [5], a single black burst representing

---

[1] If bits would be encoded vice versa (0s dominant and 1s recessive), nodes would receive the AND of both black burst frames.

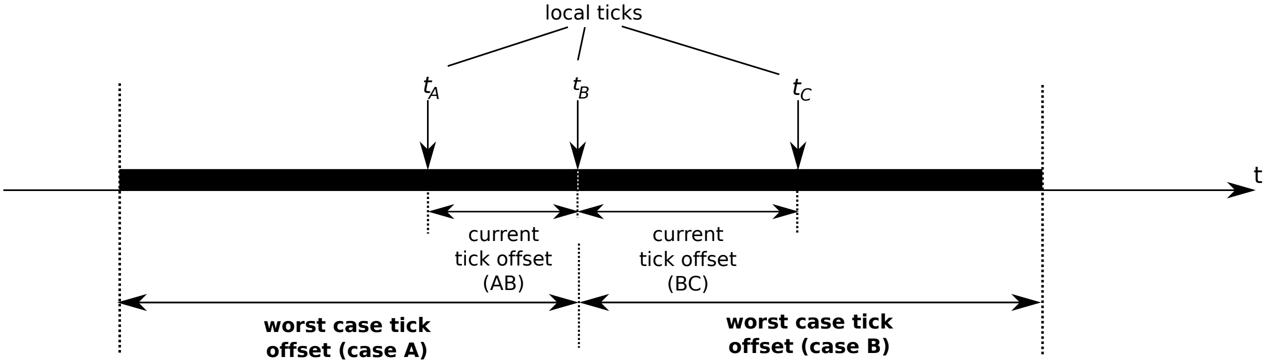

Figure 2: Node B's view of tick offset.

one logical bit consists of *five* bytes in our implementation[2]. Due to additional overhead added by the transceiver (e.g., switching times between sending/transmitting mode, ...), black burst frames are clearly less efficient than transmissions of bit sequences in regular frames. Nevertheless, the protection against collisions justifies their existence and clears the way for deterministic protocols. With advanced hardware, efficiency of black bursts can also be improved enormously in the future. E.g., switching times of Atmel's AT86RF230 transceiver (16 μs /32 μs) [1] are clearly smaller than the switching times of CC2420 (192 μs /192 μs) [5].

3.2 Tick Synchronization

The crucial key role for the successful application of black burst-based protocols lies in the synchronization of nodes. Since regular payload is lost when collisions occur, "receiving" stations can only determine when and for which duration the medium is occupied. Without any synchronization, nodes would indeed recognize medium occupancy for some duration, but starting and end point of occupancy would not provide any useful information, because two nodes would recognize those events on a locally different time base. Furthermore, distinction between periods in which collision-protected transmissions are allowed and periods in which regular frames are permitted would not be possible without any synchronization. Therefore, it is essential that all nodes build on the identical time base to interpret all time-related information in the same way.

For many application, e.g., medium slotting or duty cycling, it is sufficient to synchronize all nodes *relatively* to a common tick. Thus, nodes perceive the same relative time since the last tick occurred without associating points in time to absolute values. In comparison to *time synchronization* that adjusts absolute clock values, a *tick synchronization* is weaker but less expensive.

While tick synchronization is easy to achieve in single broadcast domains, it becomes challenging in multi-hop topologies. Many proposals aiming for multi-hop tick synchronization suffer from unbounded convergence delays or an inaccurate synchronization. The most popular one is probably Reference Broadcast Synchronization [9] that provides indeed a very high accuracy but is affected by unbounded convergence delay. In the following, we assume a network-wide tick synchronization with deterministic accuracy, i.e. with a deterministically bounded tick offset, as established by Black Burst Synchronization (BBS, [10]).

In Fig. 2, tick synchronization is shown from the local point of view of node B, which is synchronized with nodes A and C. Due to inaccuracy, there is a *current tick offset* and nodes perceive their local ticks at different points in time. This current tick offset is variable, but it is bounded by the *maximal tick offset* that has to be guaranteed by the synchronization protocol. Thus, node B not knowing the

---

[2] 3 byte preamble (shortened), 1 byte SFD, 1 byte length field with value 0; note that such a frame is not compliant to IEEE 802.15.4. It is therefore not detected as valid frame during reception.

current tick offset has to use the maximal tick offset for all time-related calculations, e.g., to determine guard times of a TDMA based medium access. This maximal tick offset represents the worst case of synchronization inaccuracy and consists of one of the following two cases:

A Node B perceives the tick after all other nodes, corresponding to the left side of Fig. 2.

B Node B perceives the tick in front of all other nodes. This corresponds to the worst case situated on the right side of Fig. 2.

It is crucial to see that both cases can not occur simultaneously and the current tick offset is somewhere in-between. However, both cases must be considered when designing protocols based on accurate synchronization. We see the impact of this fact in Sect. 3.3 when discussing the detection and interpretation of black bursts.

In the rest of the paper, we allow black bursts to be transmitted at predefined points in time and assume nodes to be tick synchronized with a maximal offset of $d_{maxOffset}$. According to [10], $d_{maxOffset}$ consists of $d_{maxBaseOffset}$ that is present immediately after the application of the synchronization protocol, and $d_{maxClockSkew}$ caused by clock skew of the hardware quartz. The concrete value of $d_{maxClockSkew}$ depends both on the accuracy of the hardware quartz and the interval of resynchronization. To obtain the deterministic bound $d_{maxOffset}$ as required, periodical resynchronization is presumed in an interval sufficient to the application's demands.

3.3 Detection of Black Bursts

As discussed above, black bursts are implemented by regular MAC frames of minimal length. But there are two reasons that avoid the reception of black bursts as regular MAC frames: First, black bursts may be sent simultaneously and resulting collisions can cause the corruption of *Preamble* or *Start-of-Frame Delimiter (SFD)*. Secondly, black bursts are not necessarily implemented as valid IEEE 802.15.4 frames. If, for instance, the preamble is shortened as it is done in our implementation for the CC2420 transceiver, the transceiver may not detect a valid frame during reception, even if no collision occurs. Thus, the decoding of black bursts is based on Clear Channel Assessment (CCA) only. CCA can be found in all customary wireless transceivers providing basic CSMA/CA support. Current implementations of CCA, e.g., in CC2420 [5], offer several modes for deciding whether the medium is busy. For detection of black bursts, most frequently found energy detectors are sufficient [17]. However, in a "noisy" environment, false positives may become a problem, since noise may be interpreted as black burst. This problem was also reported in WiDom [22]. To lower this problem, CCA thresholds must be calibrated specific to the environmental conditions to reduce false positives.

Now, we discuss the detection of black bursts, which is mainly affected by hardware-related characteristics. In more detail, we show the factors to be considered when determining the transmission time of black bursts and additional guard times.

3.3.1 Recognition of Medium Occupancy

One crucial part of cooperative transfer as well as arbitrating transfer is the correct assignment of received black bursts to time intervals. Therefore, tick offset as discussed in Sect. 3.2 and several delays introduced by hardware must be taken into account when determining the point in time at which black burst reception starts. In the following, we omit propagation delay and processing delay caused by processor. Compared to tick offset and transceiver-related delays, those delays are usually small and less significant, but in a real implementation, they must be considered, too.

Below, we consider two nodes A and B. Both nodes are tick synchronized within a guaranteed maximal offset $d_{maxOffset}$. Assume node A initiates the transmission of a black burst at its locally perceived tick t_A. This tick is perceived by node B at t_B. In the following, we derive the earliest and latest locally perceived point in time, at which node B recognizes the start of the black burst. This corresponds to the minimal and maximal delay that node B has to wait after t_B before “seeing” the black burst transmitted at t_A on the medium.

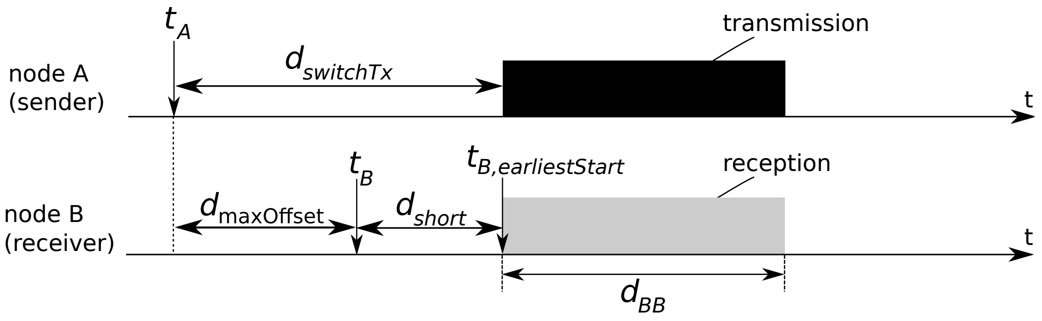

Figure 3: Earliest start of a black burst reception from node B's point of view.

First, we derive the earliest point in time by means of Fig. 3. In the figure, node B perceives its tick t_B with maximal delay $d_{maxOffset}$ after node A's tick. Thus, node A initiates the transmission before t_B and B recognizes the black burst on medium at:

$$\begin{aligned} t_{B,earliestStart} &= t_A + d_{switchTx} && (1)\\ &= (t_A + d_{maxOffset}) + d_{short} && (2)\\ &= t_B + d_{short} && (3) \end{aligned}$$

Here, $d_{switchTx}$ is the hardware-dependent switching time from receiving to sending mode. Note that $(d_{switchTx} - d_{maxOffset})$ could also be negative if the maximal tick offset is larger than the transceiver's switching time. In this case, node B would see the start of the black burst earlier than it observes its local tick t_B. Nevertheless, node B could also associate the black burst with its local tick t_B, since $(d_{switchTx} - d_{maxOffset})$ depends on constant values only and is thereby known without global knowledge.

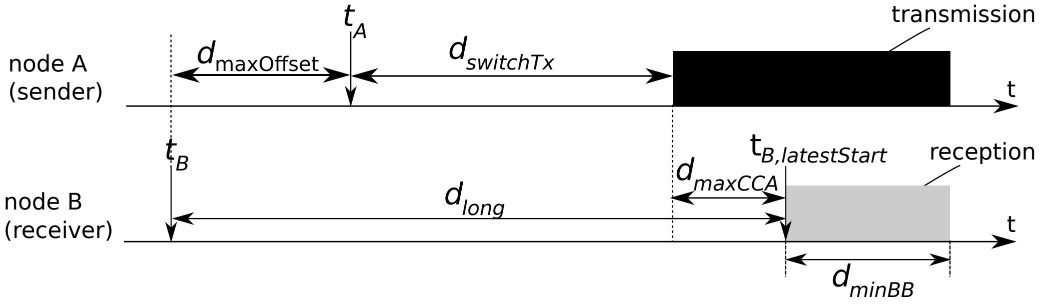

Figure 4: Latest start of a black burst reception from node B's point of view.

Now, we derive the latest point in time at which node B observes the black burst (see Fig. 4). In this case, node A's tick t_A is delayed with $d_{maxOffset}$ compared to t_B. In addition, the maximal recognition delay is also increased by the hardware-dependent execution of CCA. This additional CCA delay is actually variable, but an upper bound is provided by many transceivers, e.g., CC2420 has according to the data sheet an upper bound of $d_{maxCCA} = 128\,\mu\text{s}$ (8 symbols, each with $16\,\mu\text{s}$) [5]. Thus, the latest start of recognition is determined as:

$$\begin{aligned} t_{B,latestStart} &= t_A + d_{switchTx} + d_{maxCCA} && (4)\\ &= (t_A - d_{maxOffset}) + d_{long} && (5)\\ &= t_B + d_{long} && (6) \end{aligned}$$

Thus, node B associates a received black burst with its local tick t_B, if the start of reception is within

$$[t_{B,earliestStart}, t_{B,latestStart}] = [t_B + d_{short}, t_B + d_{long}] \tag{7}$$

$$\begin{aligned} d_{short} &= d_{switchTx} - d_{maxOffset} \quad (8) \\ d_{long} &= d_{maxOffset} + d_{switchTx} + d_{maxCCA}. \quad (9) \end{aligned}$$

3.3.2 Duration of Medium Occupancy

After discussing the association of reception start time to local perceived ticks, this section investigates the length of black bursts observed by receiving nodes. Although the transmission of black bursts is performed equally on each node with a constant duration, the duration of medium occupancy observed by receiving nodes can differ for hardware-dependent reasons and due to inaccuracy of synchronization. To distinguish black bursts from regular MAC frames, regular data frames should be significantly larger than black bursts.

The regular length of a black burst is shown in Fig. 3. In this case, sender and receiver coincide on the duration of medium occupancy. Since in our implementation, a black burst is realized as ordinary MAC frame, the regular duration of medium occupancy is calculated as

$$d_{BB} = \frac{s_{minFrame}}{r} \tag{10}$$

where $s_{minFrame}$ is the minimal length of a MAC frame in byte and r the transfer rate of the wireless transceiver.

The minimal duration of a received black burst is depicted in Fig. 4. Different from the regular duration, the receiver does not recognize the black burst immediately but with a delay according to the maximal CCA jitter. Thus, the minimal duration of medium occupancy caused by a black burst is determined as:

$$d_{minBB} = \frac{s_{minFrame}}{r} - d_{maxCCA}. \tag{11}$$

For the evaluation of the maximal duration of perceived medium occupancy, simultaneous transmissions have to be considered. The derivation is explained by means of Fig. 5.

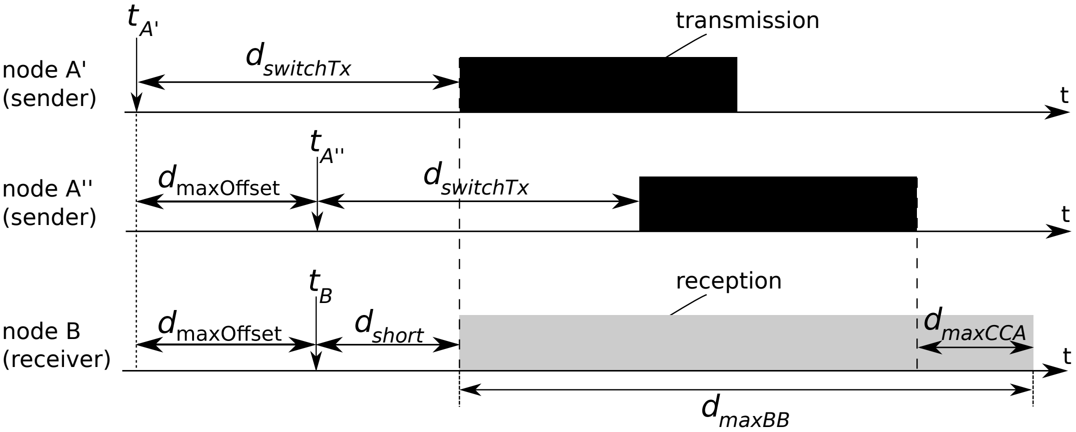

Figure 5: Maximal perceived duration of medium occupancy caused by simultaneous transmissions of black bursts.

According to the figure, we identify two sources of increase: First, the simultaneous transmissions cause a *real* increase of medium occupancy by $d_{maxOffset}$. Secondly, due to the delay introduced by

CCA jitter the perceived duration of medium occupancy is *virtually* increased by d_{maxCCA}, which is the duration the CCA mechanism requires to observe the medium idle again. This leads to the following maximal perceived length:

$$d_{maxBB} \quad = \quad \frac{s_{minFrame}}{r} + d_{maxCCA} + d_{maxOffset}. \tag{12}$$

Thus, a station reports the reception of a black burst if the medium occupancy is within the interval

$$[d_{minBB}, d_{maxBB}] = [\frac{s_{minFrame}}{r} - d_{maxCCA}, \frac{s_{minFrame}}{r} + d_{maxCCA} + d_{maxOffset}]. \tag{13}$$

By means of the reception start time and the perceived medium occupancy, a station can thereby determine *whether* a black burst is received and to which *tick* the received black burst is associated. Since black bursts are not compatible to error-prone regular frames, measures (e.g., time division schemes) must be applied to avoid collisions between black bursts and MAC frames with regular payload. When applying a time division scheme, checking the duration of medium occupancy to distinguish black bursts from regular frames seems to be dispensable in periods, in which only the transmission of black bursts is allowed. But by checking the duration, the probability of false positives is decreased and black burst-based transmissions become more robust against noise.

4 Black Burst-based Transfer Protocols

In this section, we investigate cooperative transfer and arbitrating transfer, two protocols for the network-wide propagation of bit sequences within deterministic time bounds. These time bounds do not depend on the number of nodes but on the maximal network diameter $n_{maxHops}$ only, where the maximal network diameter represents an upper bound and the actual network diameter may also be smaller. Both protocols are candidates for enhancing quality-of-service in wireless networks. Cooperative as well as arbitrating transfer are based on black bursts. They require an accurate tick synchronization with deterministic maximal tick offset (e.g., BBS [10]) and their application must start at common known points in time. Those points in time are determined either by preconfiguration or by dynamical agreement.

This section supplements previous work, in which both transfer protocols were introduced [6]. Theoretical discussions are based on [16].

4.1 Cooperative Transfer

Cooperative transfer is intended for the network-wide, collision-protected propagation of bit sequences with fixed length (*black burst frames*) within deterministic time bounds. This is, for instance, beneficial if several stations want to transmit an *identical* black burst frame network-wide or if a single node has to broadcast a valie. E.g., in [6], cooperative transfer supplements the synchronization protocol BBS by the master-based deterministic propagation of time values.

4.1.1 Protocol Behavior

For cooperative transfer, we introduce the following terminology:

- *Cooperative bit round:* Transmission of a single black burst over a single hop. Its constant duration is marked with $d_{coopBit}$.
- *Cooperative frame round:* Transmission of a black burst frame over a single hop, with a constant duration of $d_{coopRound}$. The duration includes the transmission time $d_{coopFrame}$ and an additional processing time $d_{processing}$.
- *Cooperative transfer:* Transmission of a black burst frame over $n_{maxHops}$, with a constant duration of d_{coop}.

The durations and all factors influencing $d_{coopBit}$ are derived in Sect. 4.1.2.

The mode of operation of cooperative transfer is explained by means of the scenario in Fig. 6. In the figure, node C transmits a black burst frame of 10 bits at a predefined point in time. By convention, the first bit of a black burst frame is always a dominant Start-of-Frame bit (SOF). To keep the figure simple, a binary 1, i.e. the transmission of a black burst, is depicted by a black rectangle, ignoring tick offset and all hardware-dependent influences.

With cooperative transfer, information is propagated frame-wise, where a black burst frame travels one hop within one cooperative frame round. Thus, node C sends the entire frame in the first frame round that is initially received by nodes A and D. After the transmission of all bits of a frame, we allow a processing delay $d_{processing}$. In the second frame round, the frame is forwarded by nodes A and D concurrently. Due to the collision-protected transmissions with black bursts, the frame is received correctly by nodes B, C, and E. Since the maximal network diameter is $n_{maxHops} = 3$, nodes B and E forward the received frame again, although all nodes are already informed. After $n_{maxHops}$ frame rounds, the propagation of black burst frames terminates, thus, all nodes in $n_{maxHops}$ range of node C are informed about the information carried by the black burst frame.

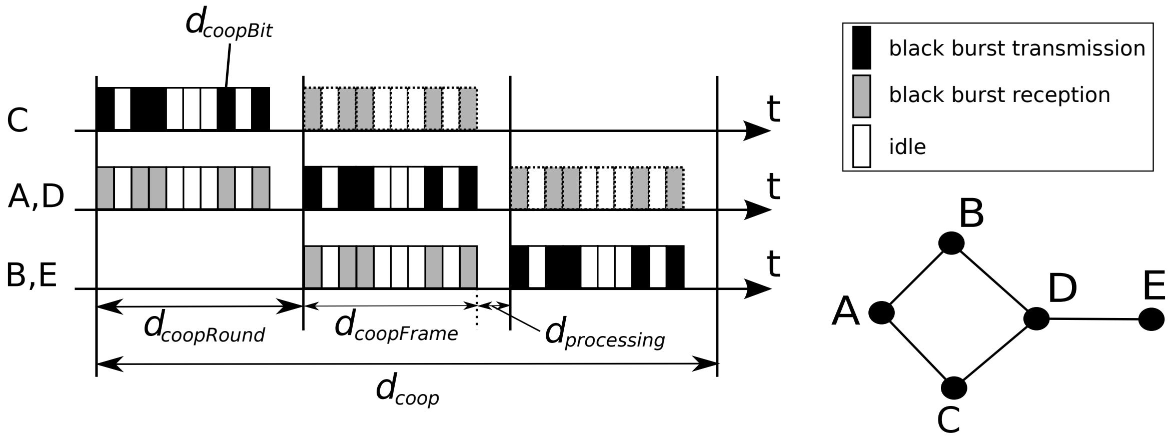

Figure 6: Cooperative transfer: Node C initiates the transmission of a black burst frame in a topology with $n_{maxHops} = 3$.

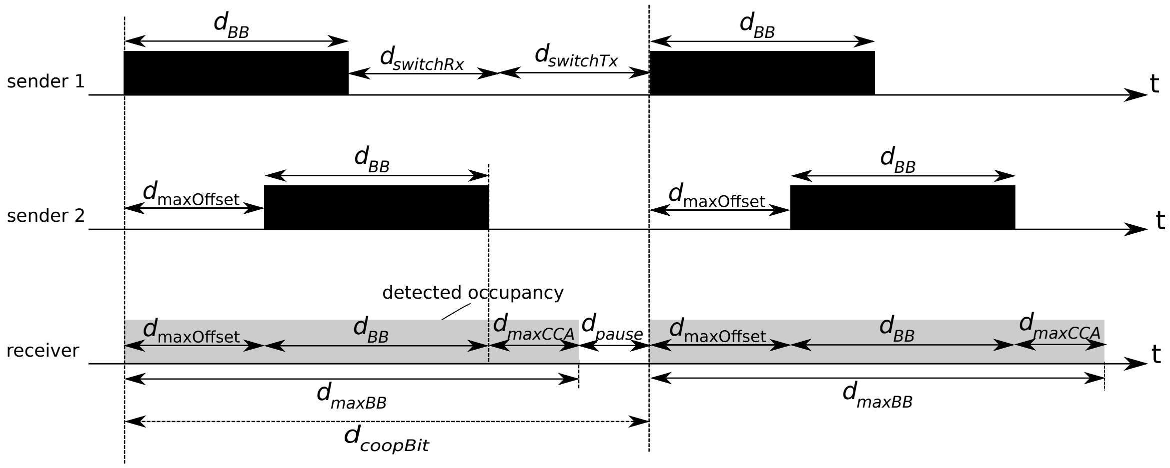

Figure 7: Derivation of $d_{coopBit}$. Constraints regarding receiving and sending nodes

4.1.2 Derivation of Convergence Delay

The crucial benefit of cooperative transfer is their deterministic bounded transfer delay. By means of Fig. 6, the transfer delay d_{coop} with n bits and a maximal network diameter of $n_{maxHops}$ is derived as follows:

$$\begin{aligned} d_{coop} &= n_{maxHops} \cdot d_{coopRound} && (14) \\ &= n_{maxHops} \cdot (d_{coopFrame} + d_{processing}) && (15) \\ &= n_{maxHops} \cdot (n \cdot d_{coopBit} + d_{processing}) && (16) \end{aligned}$$

To determine the duration $d_{coopBit}$ of a single cooperative burst, we have to look deeper into synchronization inaccuracy and hardware characteristics. Thereby, we utilize the fact that a receiving node never turns into a sending node within a frame round. In detail, the following two constraints must be fulfilled:

1. A transmitting station must be able to send two subsequent dominant bits of a frame. In this case, we must only look at the transmitter and can ignore the receiving node. This case is shown in Fig. 7 at sender 1. Here, the transmitter of the black burst requires the pure transmission time for the frame d_{BB}, the transceiver's switching time from send into receive mode $d_{switchRx}$, and the transceiver's switching time from receive into send mode $d_{switchTx}$. The switching times

Parameter	CC2420	Comment
$s_{minFrame}$	5 bytes	minimal frame size, used by black bursts
r	250 kBit/s	transmission rate
d_{BB}	160 μs	transmission time for black burst
$d_{switchTx}$	192 μs	switching time to send mode
$d_{switchRx}$	192 μs	switching time to receive mode
$d_{accessRx}$	320 μs	switching time to receive mode until CCA is valid
d_{maxCCA}	128 μs	maximal delay until medium status change is detected
d_{pause}	16 μs	minimal pause between subsequent black bursts
$r_{maxClockSkew}$	40 ppm	variation in clock rate of quartz
$d_{maxOffset}$	336 μs	maximal tick offset provided by BBS for $n_{maxHops} = 5$ and a resynchronization interval of 1 s [10]
$d_{processing}$	300 μs	processing delay after the reception of a black burst frame

Table 1: Hardware-specific parameters for CC2420 [5] transceiver.

seem to be unnecessary, since the mode is changed back and forth. However, most of common transceivers automatically switch into receive mode after transmission is finished. Therefore, we assume a transceiver with the same behavior.

Hence, the transmitter's constraint to $d_{coopBit}$ is

$$d^1_{coopBit} \geq d_{BB} + d_{switchRx} + d_{switchTx} \tag{17}$$

2. A receiving station must be able to distinguish two subsequent dominant bits of a frame, i.e. two subsequent black burst must always be separated by a minimal pause. Therefore, we consider the worst case, in which two stations with maximal tick offset transmit a black burst frame concurrently. This case is shown in Fig. 7 at the receiver. As already discussed in Sect. 3.3.2, the perceived medium occupancy is increased by the maximal tick offset $d_{maxOffset}$ and the maximal CCA delay d_{maxCCA}. To give the receiver a chance to see the medium idle, a short pause is added, resulting in the following receiver's constraint:

$$d^2_{coopBit} \geq d_{BB} + d_{maxOffset} + d_{maxCCA} + d_{pause} \tag{18}$$

According to Formulae 17 and 18, the concrete duration depends on the specific hardware and $d_{coopBit}$ is calculated as:

$$\begin{aligned} d_{coopBit} &= max\{d^1_{coopBit}, d^2_{coopBit}\} && (19)\\ &= max\{d_{BB} + d_{switchRx} + d_{switchTx}, d_{BB} + d_{maxOffset} + d_{maxCCA} + d_{pause}\} && (20)\\ &= d_{BB} + max\{d_{switchRx} + d_{switchTx}, d_{maxOffset} + d_{maxCCA} + d_{pause}\} && (21)\end{aligned}$$

4.1.3 Example

In this section, an example calculation is given with $n_{maxHops} = 5$ and $n = 16$ bits. Hardware-specific values are given for TexasInstruments' CC2420 transceiver [5] in Tab. 1. For frame processing, we allow $d_{processing} = 300\,\mu\text{s}$.

- *Cooperative bit round*:

$$d_{coopBit} = 160\,\mu\text{s} + max\{192\,\mu\text{s} + 192\,\mu\text{s}, 336\,\mu\text{s} + 128\,\mu\text{s} + 16\,\mu\text{s}\} = 640\,\mu\text{s} \tag{22}$$

- *Cooperative frame round*:

$$d_{coopRound} \quad = \quad 16 \cdot d_{coopBit} + 300\,\mu\mathrm{s} \ = 10.54\,ms \tag{23}$$

- *Cooperative transfer*:

$$d_{coop} \quad = \quad 5 \cdot d_{coopRound} = 52.7\,ms \tag{24}$$

4.2 Arbitrating Transfer

Arbitrating transfer is applicable to network-wide medium arbitration preceding to regular data transmissions, deterministic leader election, or to the network-wide agreement on a common value. As cooperative transfer, arbitrating transfer provides deterministic time bounds depending on the maximal network diameter and the number of transmitted bits. The difference to cooperative transfer is the decentralized nature of arbitrating transfer, since arbitrarily many stations can start the transmission of *different* black burst frames. As result, all stations have received the same unique black burst frame that corresponds to the highest value sent by all stations as soon as the protocol's application is finished. For instance, in [6], arbitrating transfer was utilized for the decentralized agreement on a common time value as part of a time synchronization protocol.

4.2.1 Protocol Behavior

For arbitrating transfer, we introduce the following terminology:

- *Arbitrating bit round:* Transmission of a single bit over one hop, with a duration of $d_{arbBRound}$. Although, $d_{arbBRound}$ looks similar to $d_{coopBit}$, both durations differ in length (see Sect. 4.2.2).
- *Arbitrating bit phase:* Proapgation of a single bit over $n_{maxHops}$ hops, with a duration of $d_{arbBPhase}$.
- *Arbitrating transfer:* Transmission of all bits over $n_{maxHops}$, with a duration of d_{arb}.

The relation between the durations and all factors influencing $d_{arbBRound}$ are derived in Sect. 4.2.2.

The functionality of arbitrating transfer is explained by means of the scenario in Fig. 8. In the figure, every station is initially in state *active*. Thus, they process their black burst frames and transmit a dominant bit (SOF) in the first arbitrating bit round. In contrast to cooperative transfer, arbitrating transfer processes black burst frames bit-wise, i.e. after transmitting a black burst, a station waits to allow the forwarding of the burst over $n_{maxHops}$ hops. Since the second bit is dominant in nodes B's and C's black burst frame only, only those two nodes transmit a black burst in the second arbitrating bit phase that is received by nodes A and D. A and D forward the received bit in the second bit round of the second bit phase and change their state into *repeating*. By means of their state, nodes decide whether they compete with their own black burst frame (state *active*) or whether they repeat received black bursts only (state *repeating*)[3]. So, node A and D stop competing with their own bit sequence to guarantee that only nodes holding the black burst frame with the highest value transmit their frames completely. The forwarded second bit is received by node E, changing node E's state into *repeating* and letting node E forward the black burst, too. Thus, after the second arbitrating bit phase, only nodes B and C are still in competition. After the third bit, node C remains as single node in competition, since node B processes a recessive 0 while C's third bit is dominant. Thereby, node B's state is also changed into *repeating* and only node C's black burst frame is sent completely. So at the end of arbitrating transfer, every station observes node C's sequence 1110 on the wireless medium and node C can rate itself as winner.

---

[3] Since arbitrating transfer is applied in tick synchronized networks, repeating nodes can locally decide whether a received black burst must be forwarded or whether a black burst has already travelled $n_{maxHops}$ hops. Thereby, termination of forwarding is always guaranteed.

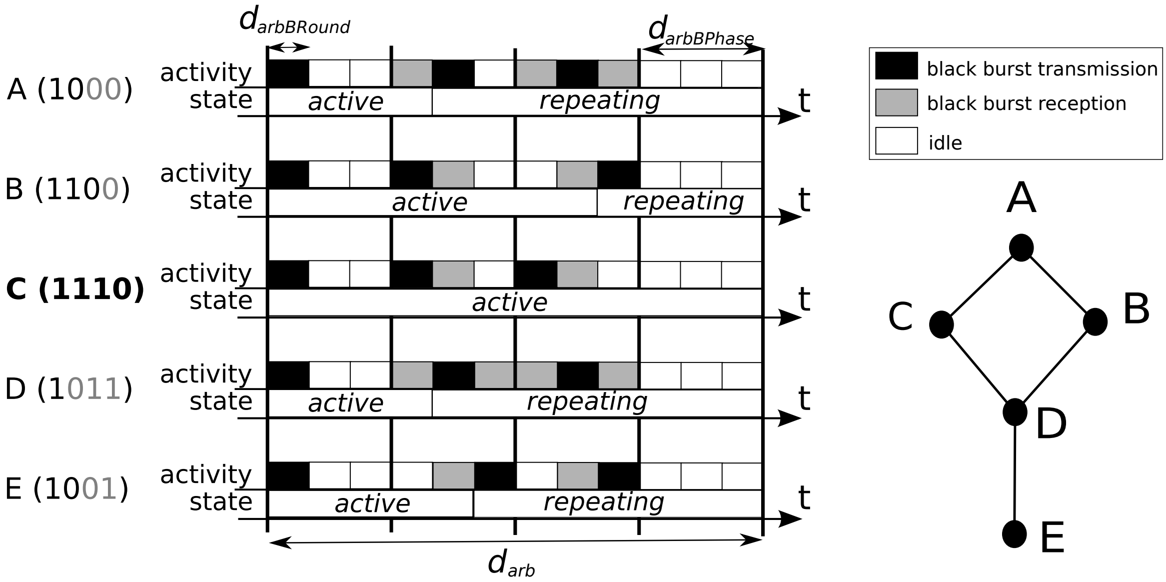

Figure 8: Arbitrating transfer: All nodes send different black burst frames in a topology with $n_{maxHops} = 3$. At the end, node C's bit sequence is received by all nodes.

4.2.2 Derivation of Convergence Delay

The deterministic transfer delay of arbitrating transfer depends on synchronization inaccuracy, hardware characteristics, the maximal network diameter $n_{maxHops}$, and the length n of the black burst frame to be transmitted and is derived as follows:

$$\begin{aligned} d_{arb} &= n \cdot d_{arbBPhase} && (25) \\ &= n \cdot n_{maxHops} \cdot d_{arbBRound} && (26) \end{aligned}$$

$d_{arbBRound}$ is again derived under constraints given by hardware times and tick offset. In particular, the following constraints must be considered:

1. A repeating station must be able to receive a dominant bit and to forward the bit immediately in the next arbitrating bit round. This case also covers the case that subsequent black bursts must not cause a single, continuous busy tone, i.e. that the maximal perceived occupancy of d_{maxBB} as derived in Sect. 3.3.2 is not violated.

   This case is outlined in Fig. 9, whereas nonrelevant durations are omitted: Due to both tick offset and CCA jitter, the repeating node perceives the medium idle at the locally latest point in time of the current bit round. Since the black burst must directly be forwarded in the next bit round, the bit round must additionally consider the hardware switching time $d_{switchTx}$, so that the repeating node is ready to transmit at the start of the next bit round. Hence, the following constraint must hold:

$$\begin{aligned} d^{1}_{arbBRound} &\geq d_{maxBB} + d_{pause} + d_{switchTx} && (27) \\ &= d_{BB} + d_{maxOffset} + d_{maxCCA} + d_{pause} + d_{switchTx} && (28) \end{aligned}$$

2. A repeating station must be able to forward a black burst in the last arbitrating bit round of a bit phase and to receive a black burst in the first bit round of the subsequent bit phase. The worst case for this constraint is shown in Fig. 10. The repeating node has maximal tick offset to the future transmitter of a black burst. To ensure that the repeater observes this future black burst completely, the repeater must be allowed to switch back to receiving mode and to obtain a valid CCA value before the black burst transmission starts.

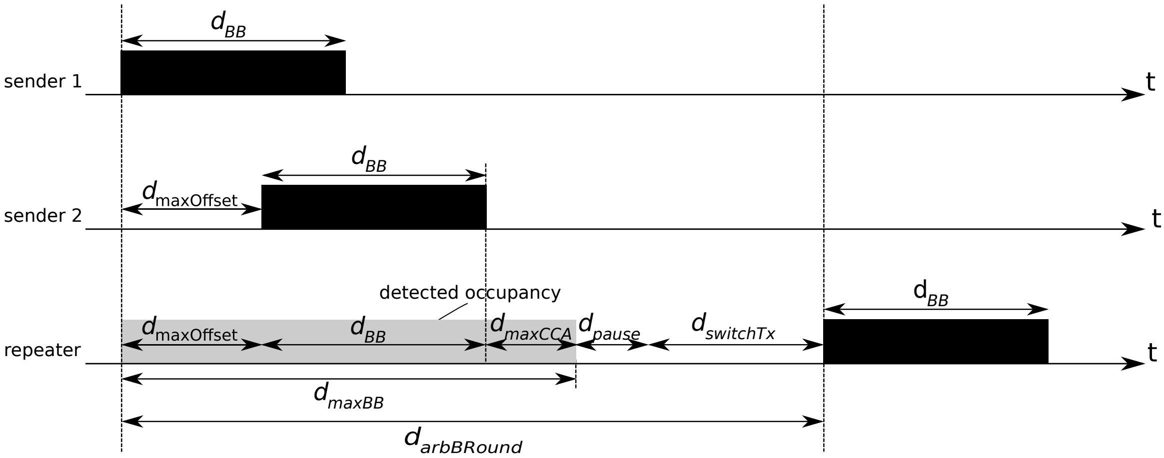

Figure 9: Derivation of $d^1_{arbBurst}$ covering constraints regarding repeating nodes: One arbitrating bit round must be sufficient to receive a "late" black burst and to prepare repeating the black burst in the next bit round of the same bit phase.

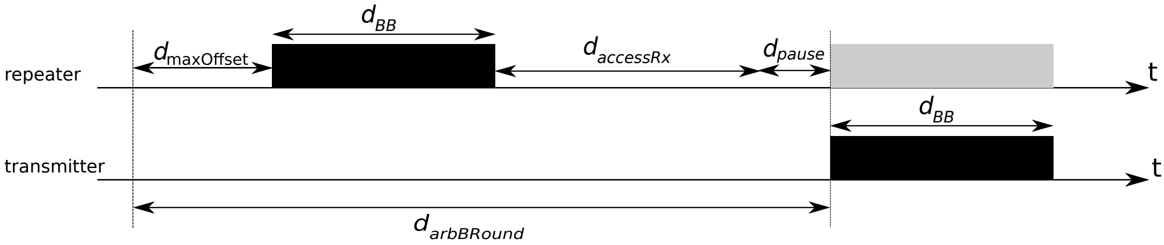

Figure 10: Derivation of $d^2_{arbBurst}$ covering more constraints regarding repeating nodes: Nodes repeating a black burst in the last arbitrating bit round of a bit phase must be able to receive a new black burst in the first bit round of the next bit phase.

Thus, the following inequation must hold:

$$d^2_{arbBRound} \quad \geq \quad d_{maxOffset} + d_{BB} + d_{accessRx} + d_{pause} \tag{29}$$

The duration for switching to receive mode and obtaining a valid CCA is marked with $d_{accessRx}$ and is in general larger than the pure switching time $d_{switchRx}$. On CC2420, it is composed of the switching time $d_{switchRx}$ and the maximal CCA jitter d_{maxCCA} [5].

3. If we do not assume $n_{maxHops} > 1$, a node must also be able to send two dominant bits in subsequent bit phases [4]. This case was already covered with cooperative transfer (see Fig. 7):

$$d^3_{arbBRound} \quad = \quad d^1_{coopBit} \geq d_{BB} + d_{switchRx} + d_{switchTx} \tag{30}$$

From Formulae 28, 29, and 30, the hardware-dependent duration $d_{arbBRound}$ is as follows:

$$\begin{aligned} d_{arbBRound} &= max\{d^1_{arbBRound}, d^2_{arbBRound}, d^3_{arbBRound}\} && (31)\\ &= d_{BB} + max\{d_{maxOffset} + d_{maxCCA} + d_{pause} + d_{switchTx}, && (32)\\ &\quad d_{maxOffset} + d_{accessRx} + d_{pause}, d_{switchRx} + d_{switchTx}\} \end{aligned}$$

$$= d_{BB} + max\{d_{maxOffset} + d_{pause} + max\{d_{maxCCA} + d_{switchTx}, d_{accessRx}\}, d_{switchRx} + d_{switchTx}\} \tag{33}$$

---

[4] If $n_{maxHops} > 1$ holds, this constraint is negligible, since a node has always a pause of at least one arbitrating bit round between the transmission/forwarding of two dominant bits.

In summary, $d_{arbBRound}$ is in general larger than $d_{coopBit}$, making the application of arbitrating transfer more expensive than cooperative transfer. Nevertheless, if several stations may transmit different values, the application of cooperative transfer is not possible and arbitrating transfer is the only alternative.

4.2.3 Example

With the CC2420's parameters in Tab. 1, transfer delay with arbitrating transfer for $n = 16$ bits and $n_{maxHops} = 5$ hops is calculated as follows:

- *Arbitrating bit round*:

$$d_{arbBRound} = 160\,\mu\text{s} + max\{336\,\mu\text{s} + 16\,\mu\text{s} + 320\,\mu\text{s}\,, 192\,\mu\text{s} + 192\,\mu\text{s}\} = 832\,\mu\text{s} \tag{34}$$

- *Arbitrating bit phase*:

$$d_{arbBPhase} = 5 \cdot 832\,\mu\text{s} = 4.16\,ms \tag{35}$$

- *Arbitrating transfer*:

$$d_{arb} = 16 \cdot 4.160\,ms = 66.56\,ms \tag{36}$$

5 Experimental Evaluation of Black Bursts

In this section, we evaluate black bursts on the Imote2 platform [8]. More experimental results can be found in [17, 16] with MICAz motes [7], which are equipped with the same TexasInstruments CC2420 [5] transceiver. The ollowing experiments can not evaluate the behavior of black bursts or even wireless signals in full detail, but they are rather intended to provide a proof-of-concept for the applicability of black bursts.

First, we consider the success rate of black burst transmissions that is affected by the accuracy of CCA and interferences/noise on the wireless channel. Thereafter, we investigate the transmission range of black bursts compared to the transmission range of regular MAC frames containing ordinary payload.

5.1 Proof-of-Concept - The Black Burst Reception Ratio (BBRR)

The great benefit of black bursts is their collision resistance in case of overlapping transmissions. This subsection provides a lightweight evidence of this advantage by provoking simultaneous transmissions of black bursts and by determining the resulting detection ratio.

5.1.1 Experimental Setup

Figure 11 introduces the single-hop topology on which the first experiment is based. The topology consists of three nodes: A master node, coordinating the experiment and recording received black bursts, and two sender nodes, transmitting black bursts concurrently. Samples, i.e. concurrently transmitted black bursts, are collected in a round-based fashion, where one round has a duration of 1 s. At the beginning of each round, synchronization is achieved by the master node by the transmission of a beacon. For this single-hop scenario, this simple synchronization is sufficient. After reception of a beacon, both sender nodes wait 100 ms before starting the transmission of 40 black bursts in an interval of 20 ms. Thereafter, nodes wait for the next round that starts again with the reception of a synchronization beacon. Beacons as well as black bursts are always transmitted on channel 15 ($\approx 2425\,MHz$, no overlap with 802.11 channels) and with the maximal output power of 0 dBm [5].

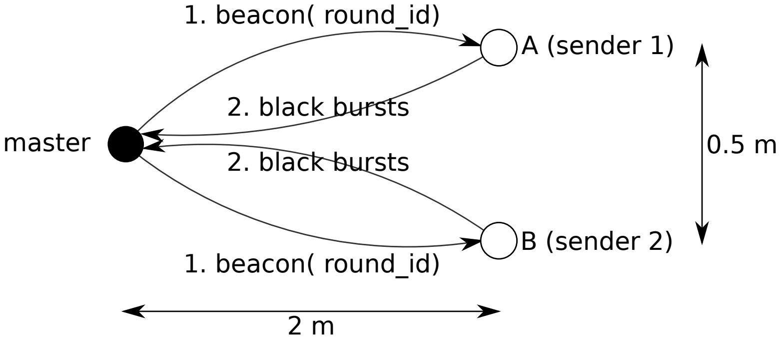

Figure 11: Toplogy of the first experiment, providing a proof-of-concept of collision-protected transmissions with black bursts.

To filter out rounds in which beacons get lost/corrupted, each beacon carries a round number that is logged by both sending nodes. Thereby, rounds, in which at least one node missed the beacon, are marked and omitted.

The initial synchronization inaccuracy after the transmission of a beacon is bounded by d_{maxCCA}, since both black burst senders can synchronize to the start of beacon reception. Because resynchronization is performed every second, the maximal tick offset $d_{maxOffset}$ at the end of each period is bounded by

$d_{maxCCA} + 2 \cdot 1\,s \cdot r_{maxClockSkew}$, where $r_{maxClockSkew}$ is the variation in the quartz clock rate. According to the values for the CC2420 transceiver in Tab. 1, $d_{maxOffset}$ is calculated as $d_{maxOffset} = 208\,\mu$s . By means of Formulae 11 and 12, this results in the following minimal and maximal perceived valid medium occupancy:

$$d_{minBB} = 32\,\mu\text{s} \tag{37}$$

$$d_{maxBB} = 496\,\mu\text{s} \ . \tag{38}$$

Thus, the master node interprets the busy medium as valid black burst, if the duration of perceived medium occupancy is within [$32\,\mu$s , $496\,\mu$s][5].

5.1.2 Results

The experiment ran for 5000 rounds with an overall duration of ≈ 90 min. Thus, in the ideal case, the master receives 200 000 black bursts. But due to loss of beacons, 3 rounds were rated invalid, resulting in 4997 remaining rounds.

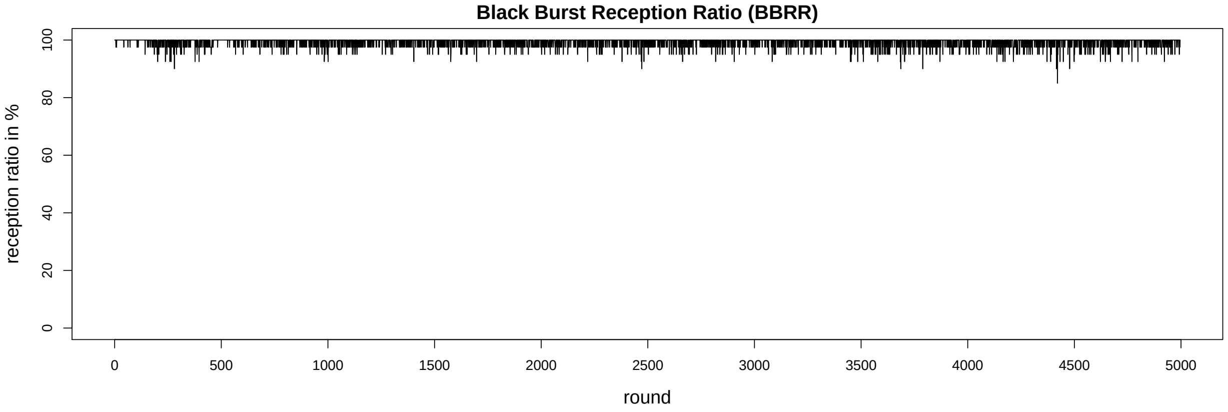

Figure 12: Black Burst Reception Ratio (BBRR).

The resulting Black Burst Reception Ratio (BBRR) is presented in Fig. 12. Although the ratio is close to 100%, rounds can be found, in which some transmitted black bursts are not detected correctly. We can trace those *false negatives* back to a noisy environment, since in most cases, the medium occupancy is in fact detected, but with an observed duration that does not fit into the interval of valid black bursts. Nevertheless, the average BBRR over all valid rounds is 99%, i.e. black bursts are transmitted with high reliability even in noisy environments without laboratory conditions.

Due to the noisy environment, the “mistaken” observed medium occupancy by the CCA mechanism of the transceiver does not only lead to an elongation of occupancy caused by valid black bursts, leading to the occurence of *false negatives*, but can also induce the detection of *false positives*. We investigate this problem by means of Fig. 13. The figure presents the number of received black bursts that are actually *not* transmitted per round. Those black bursts can be simply filtered out in this scenario, because the receiving node knows that black bursts are only sent in intervals of 20 ms, i.e. black bursts are sent at predefined points in time only. Nevertheless, false positives can become problematic, if black bursts are used to encode binary numbers as described in Sect. 3.1.

Although we observed the transceiver to detect medium occupancy due to noise multiple times (actually more than we expected), the number of false positives is very small (see Fig. 13). This is due to the fact that the duration of medium occupancy caused by noise (or interferences) is mostly outside the

---

[5] This interval does not include processing delay caused by the limited hardware speed. In our implementation, we allow $100\,\mu$s processing delay in addition.

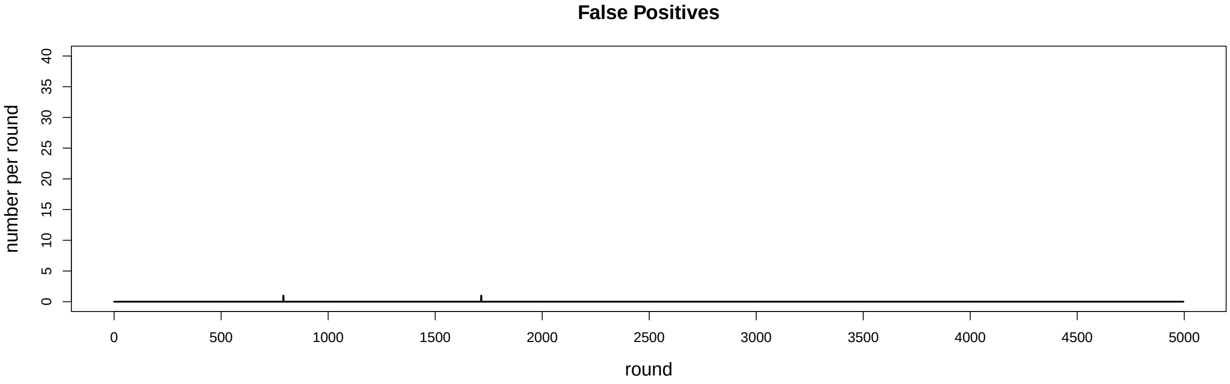

Figure 13: Reception of false positives.

interval of a valid black burst. In summary, in the environment of the first experiment, the probability of distorting a valid black burst is clearly higher than the generation of false positives. Nevertheless, this is no universal statement, but highly depends on the environmental conditions.

The problem of false positives and false negatives can be decreased when taking the received signal strength into account. Especially in topologies with small distances between senders and receivers, the received signal strength due to transmissions is in general significantly higher than noise. Thus, medium occupancy can be considered as relevant only when the received signal strength is above a threshold, which has to be determined empirically, depending on the specific topology. We investigate the received signal strength in the next section, when evaluating the transmission ranges of black bursts and regular MAC frames.

5.2 Transmission Range: Black Bursts vs. Regular MAC Frames

In this section, we investigate the differences between the transmission range of regular MAC frames compared to the transmission range of black bursts. First, we outline the origin of the actual problem, afterwards we provide a solution to this problem and an experimental evidence of its applicability.

5.2.1 Problem Statement

It is well-known that in wireless networks, links between nodes are divided into several classes. In [28], the authors identify the following categories:

- **Communication link**: Two nodes are connected with a communication link, if a frame transmitted by one node is correctly received by the second node.
- **Interfering link**: An interfering link exists between two nodes, if frames can not be transmitted free of errors on this link, yet, a transmission by one node disturbs the communication between the second node and an arbitrary third node.
- **Sensing link**: A sensing link exists, if a node observes a transmission started by another node, but the observed transmission can neither be received correctly nor it disturbs the communication to a third node.

In the famous paper “The mistaken axioms of wireless-network research” by Kotz et al. [15], this existence of different link types is also covered by the fourth mistaken axiom:

If I can hear you at all, I can hear you perfectly.

In their experiments, the authors disprove this axiom by demonstrating a high frame error rate between two nodes, although the nodes are able to sense the carrier of the transmission. Thereby, they show that on the one hand, sensing links are (in general) larger than communication links, and that on the other hand, communication links are not of fixed length but vary over time.

Justified by the larger sensing range, the authors of [33] even claim the hidden terminal problem to be overrated. Furthermore, they argue that the hidden terminal problem is solved automatically, since they assume sensing links to be larger than twice the communication links.

Though an increased sensing range is advantageous for several protocols (like CSMA/CA), it is often undesired for the application of black bursts. Particularly, the transmission range of black bursts is based on sensing links. Therefore, compared to regular MAC frames, black bursts can in general be transmitted over larger distances, which is a problem for many applications of black bursts. E.g., in [3], a quality-of-service routing protocol called BBQRt is presented, utilizing black bursts during route discovery. Without further measures, route discoveries can find routes that are actually not sufficient for regular MAC frames.

To face this problem, we consider the received signal strength (RSS) as further constraint when deciding whether a black burst is received. Therefore, the RSS indicator (RSSI) provided by the transceiver is evaluated beside the duration of medium occupancy. By a hardware dependent conversion, the RSSI can be transformed into RSS in dBm. Black bursts are only accepted if both the duration of medium occupancy is within $[d_{maxBB}, d_{minBB}]$ (see also Sect. 3.3) and the RSS exceeds a threshold, which we are going to determine in Sect. 5.2.3 empirically[6].

The approach of adjusting the transmission range of black bursts to the transmission range of MAC frames by means of the RSS implies a high correlation between the RSS and the MAC frame reception rate. Otherwise, no conclusions regarding the reception of MAC frames could be drawn from the RSS measured during a black burst reception. There already exist several publications confirming the assumption of a high correlation. E.g., in [30], the authors also utilize Texas Instruments CC2420 transceiver [5]. They do not only demonstrate that the RSSI represents a stable and significant factor, but also show a very high correlation between the RSSI and the packet error rate. Particularly, they determine a reception rate above 85%, if the RSS exceeds -87dBm. Similar results regarding correlation between RSS and packet error rate are also reported in [25], yet the authors report on a higher variability of RSS and packet error rate.

We should note that there are also publications showing the absence of any correlation between RSS and and reception rate (e.g., in [2]). It is therefore not adequate to see the correlation as universally valid law. Instead, the hardware as well as the concrete environmental conditions must be taken into account carefully, before relying on the assumption of a high correlation.

5.2.2 Experimental Setup

The experiment's topology evaluating the correlation between RSS and black bursts/MAC frames os shown in Fig. 14. The topology consists of two nodes, one sender transmitting black bursts with several output power levels and one receiver. Node are separated from each other with a distance of 0.5 m.

Because we are mainly interested in the transceiver's behavior near the sensitivity range, the sending node periodically transmits black bursts with the low output power levels from 1 to 9. According to [5], this corresponds to an output power of ≈-24 dBm to ≈-13 dBm. Due to the small distance between sender and receiver, higher levels would not provide any new insights. The sending node increases the

---

[6] The CC2420 transceiver normally concatenates the RSSI to each received MAC frame, but allows also to read the RSSI independently from frame receptions. The latter is necessary, since black bursts are probably not detected as regular MAC frames by the transceiver.

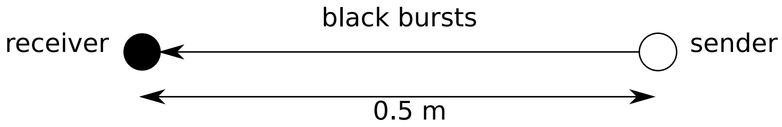

Figure 14: Toplogy of the second experiment, consisting of one sending and one receiving node.

output power every 15 minutes (10 minutes transmission, 5 minutes pause), starting with the power level 1 (-24 dBm). Thereby, the complete experiment lasts 135 minutes. Black bursts are transmitted in an interval of 20 ms for a duration of 10 minutes, thus, all in all, 30000 black bursts are transmitted per output level.

Up to now, a black burst was realized as regular MAC frame with shortened preamble and omitted checksum. By shortening the preamble, a black burst could not be received as regular MAC frame. In this experiment, a black burst must also be received as valid MAC frame, forbidding the shortage of the preamble. Therefore, we use an extended MAC frame format as implementation of a black burst in the following. Although the efficiency of black bursts is thereby decreased further, a black burst can now either be received both as black burst and regular MAC frame, or as black burst only (this happens if no preamble is detected or the MAC frame is corrupted), or as MAC frame only (this happens if the CCA mechanism fails or the observed duration of medium occupancy does not fit). Note that changing the frame format of a black burst requires the adaptation of the interval of valid black bursts, but does not change the concept behind black bursts.

5.2.3 Results

First, we investigate the MAC frame and black burst reception ratio compared to the output power levels. The result is shown in Fig. 15. As expected, the reception rate of black bursts is always higher than the reception rate of regular MAC frames, since the sensing range is larger than the transmission range. Especially, the case that a MAC frame was received correctly, but the corresponding black burst reception failed was observed very rarely (almost never). With transmission power level 2, both rates are almost identical but still less than 90%. By changing the transmission power level to 3, the reception rates of MAC frames as well as black bursts become satisfactory and almost all 30000 black bursts are received correctly. With higher power levels, reception ratios stay next to 100% (as expected) without bringing any new insight.

In a next step, we investigate the differences between MAC frame reception ratio and black burst reception ratio in more detail and compare reception ratios with the RSS.

In Fig. 16, the ratio of received MAC frames to received black bursts is plotted against the RSS. Formally, the figure plots

$$\Delta(rss) = \frac{\text{number of valid received MAC frames with signal strength } rss}{\text{number of valid received black bursts with signal strength } rss}. \tag{39}$$

Please note that the MAC frame (or black burst) reception ratio alone can not be plotted against the RSS, since all samples are recorded with the receiver node only, i.e. if the receiver node does not report medium occupancy, we can not assign an RSS to a transmitted black burst. Thus, in this case, we can only consider the number of received MAC frames and black bursts and have to ignore all "lost" black bursts.

Figure 16 underlines the fact that black bursts can still be received with a low signal strength, although regular MAC frames become corrupted or are not observed at all with such a low signal strength. E.g., at -97 dBm, the receiver "sees" some correct black bursts on the medium, but no corresponding MAC frame is received. However, with -90 dBm, the number of received black bursts corresponds to the number of received MAC frames .

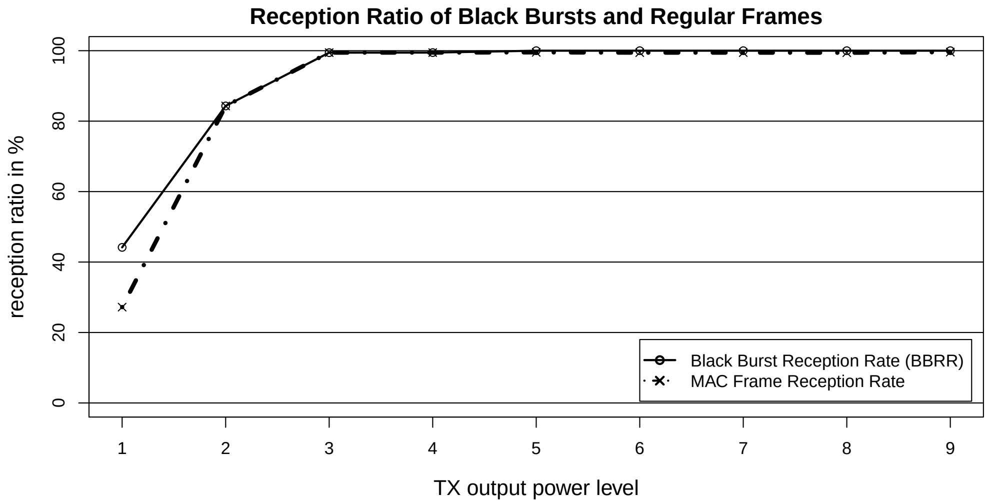

Figure 15: Reception ratio of MAC frames and black bursts for each transmission power level. With transmission power 1 only ≈28% of all MAC frames and ≈45% of all black bursts are received correctly. At transmission power 3, almost all MAC frames and black bursts are correctly received.

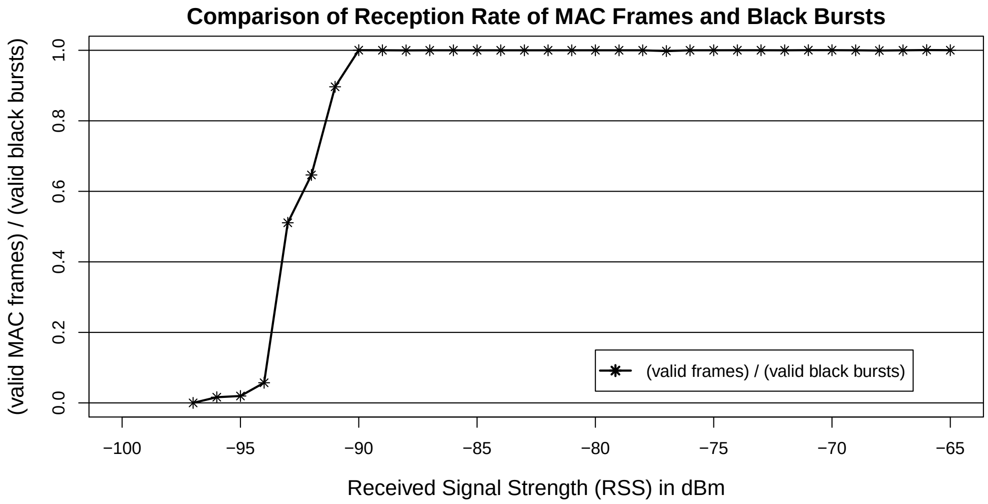

Figure 16: Δ(rss) showing the ratio between received black bursts and received MAC frames. If MAC frames are transmitted over the same link as black bursts, black bursts should only be accepted if their received signal strength is above -90 dBm. Otherwise, MAC frames may probably become corrupted.

As noted in Sect. 5.2.1, some protocols, like BBQRt [3], rely on the fact that the transmission range of black bursts and regular MAC frames must be equal. In our environment consisting of links with small distances, both ranges can be adjusted by either dropping all received black bursts with an RSS less than -90 dBm or by configuring the transceiver's CCA threshold to ignore medium occupancy with power levels less than -90 dBm. However, in many cases. it is not necessary to adjust both transmission ranges (e.g., in the synchronization protocol BBS [10]). Therefore, it is more beneficial to allow black bursts received with RSS less than -90 dBm on hardware driver level and to let higher-level protocols decide whether the black burst should be accepted or discarded.

6 Conclusions

In this paper, we havedescribed the advantages of collision-protected transmissions with black bursts. After an extensive survey of related work, we have discussed the implementation and detection of black bursts with customary hardware. In particular, all computations and experiments were based on the Imote2 platform, which is equipped with the common CC2420 transceiver. Although realizations of black bursts on customary transceivers are inefficient compared to transmissions with ordinary MAC frames, they are sufficient to prove the applicability of black bursts even with low priced hardware. Furthermore, the decreased efficiency is often compensated by the deterministic and collision-protected character.

Based on the principles of black bursts, we have introduced two protocols called cooperative and arbitrating transfer, where arbitrating transfer is more powerful but also more expensive in its application than cooperative transfer. Both protocols' advantage comes from their deterministic time bounds for transfer delay. In more detail, their convergence delay does not depend on the number of nodes, but on the maximal network diameter only. Thus, they are candidates for enhancing quality-of-service in wireless networks.

By means of two experiments, we have presented a proof-of-concept of the applicability of black bursts. On the one hand, we have shown the detection of overlapping black burst transmissions in a noisy real-world environment. All in all, the success rate was above 99%, showing the robustness and capabilities of black bursts. On the other hand, we have evaluated the difference between the transmission ranges of black bursts and regular MAC frames and have introduced a method to adjust both transmission ranges, as required by many protocols based on black bursts.

In future work, we are going to analyze concurrent transmissions in more detail. In particular, experiments are planned to determine the probability of destructive overlapping of simultaneously sent black bursts. Furthermore, we go for a multi-hop experiment consisting of multiple nodes and the implementation of complete protocols based on black bursts such as BBQRt. Since black bursts are a relatively new field of research, only few experimental evaluations can be found in the literature. To understand collision-protected transmissions with black bursts in detail, further experiments are indispensable.

References

[1] Atmel. AT86RF230 data sheet. `http://www.atmel.com/dyn/resources/prod_documents/doc5131.pdf`, 2009.

[2] Matthias Berning. Evaluation of effects in practical use of forward error correction in wireless sensor networks. Master's thesis, University of Kaiserslautern, 2009.

[3] Martin Birtel. Formal specification and evaluation of black burst quality-of-service routing. Master's thesis, Fachbereich Informatik, Technische Universität Kaiserslautern, 2009.

[4] Bach Duy Bui, Rodolfo Pellizzoni, Marco Caccamo, Chin F. Cheah, and Andrew Tzakis. Soft real-time chains for multi-hop wireless ad-hoc networks. In *IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 69–80. IEEE Computer Society, 2007.

[5] Chipcon / Texas Instruments. CC2420 datasheet. `http://focus.ti.com/lit/ds/symlink/cc2420.pdf`, 2007. Revision SWRS041b.

[6] Dennis Christmann, Reinhard Gotzhein, and Thomas Kuhn. Multi-hop Clock Synchronization in Wireless Ad-Hoc Networks. In *Proceedings of Workshops on Mobile Ad-Hoc Networks (WMAN 2009), Kassel, Germany*, March 2009.

[7] Crossbow. Micaz datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICAz_Datasheet.pdf`, 2007. Revision B.

[8] Crossbow. Imote 2 datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/Imote2_Datasheet.pdf`, 2009.

[9] Jeremy Elson, Lewis Girod, and Deborah Estrin. Fine-Grained Network Time Synchronization Using Reference Broadcasts. In *OSDI*. Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI 2002), Boston, MA, USA, 2002.

[10] Reinhard Gotzhein and Thomas Kuhn. Decentralized Tick Synchronization for Multi-Hop Medium Slotting in Wireless Ad Hoc Networks Using Black Bursts. In *Proceedings of the Fifth Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, SECON 2008, June 16-20, 2008, Crowne Plaza, San Francisco International Airport, California, USA*, pages 422–431. IEEE, 2008.

[11] IEEE Std. 802.11. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.* IEEE Computer Society, 1999.

[12] IEEE Std. 802.15.4. *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs).* IEEE Computer Society, New York, NY, USA, Oct. 2003.

[13] International Telecommunication Union (ITU). ITU-T Recommendation Z.100 (11/99): Specification and Description Language (SDL). `http://www.itu.int/ITU-T/studygroups/com10/languages/Z.100_1199.pdf`, 1999.

[14] Phil Karn. MACA - A New Channel Access Method for Packet Radio. In *Proceedings of the ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, 1990.

[15] David Kotz, Calvin Newport, and Chip Elliott. The Mistaken Axioms of Wireless Network Research. Technical Report TR2003-467, Dartmouth College, July 2003.

[16] Thomas Kuhn. *Model Driven Development of MacZ – A QoS Medium Access Control Layer for Ambient Intelligence Systems.* PhD thesis, University of Kaiserslautern, 2009.

[17] Thomas Kuhn and José Irigon de Irigon. An experimental evaluation of black burst transmissions. In Albert Y. Zomaya and Sherali Zeadally, editors, *MOBIWAC*, pages 163–167. Proceedings of the Fifth ACM International Workshop on Mobility Management & Wireless Access, MOBIWAC 2007, Chania, Crete Island, Greece, October 2007.

[18] Koji Nakano and Stephan Olariu. Randomized Initialization Protocols for Ad Hoc Networks. *IEEE Trans. Parallel Distrib. Syst.*, 11(7):749–759, 2000.

[19] International Standardization Organization. Controller Area Network (CAN), ISO 11898, 2004.

[20] Abhishek Pal, Atakan Dogan, and Füsun Özgüner. MAC layer protocols for real-time traffic in ad-hoc wireless networks. In *ICPP '02: Proceedings of the 2002 International Conference on Parallel Processing*, page 539, Washington, DC, USA, 2002. IEEE Computer Society.

[21] Lei Pan and Hongyi Wu. Performance Evaluation of The SYN-MAC Protocol in Multihop Wireless Networks. In *ICCCN*, pages 265–271. IEEE, 2007.

[22] Nuno Pereira, Björn Andersson, and Eduardo Tovar. WiDom: A Dominance Protocol for Wireless Medium Access. *IEEE Trans. Industrial Informatics*, 3(2):120–130, 2007.

[23] Nuno Pereira, Björn Andersson, Eduardo Tovar, and Paulo Carvalho. Improving the reliability of WiDom in a single broadcast domain. In *SIES*, pages 144–147. IEEE, 2009.

[24] Nuno Pereira, Bjorn Andersson, Eduardo Tovar, and Anthony Rowe. Static-priority scheduling over wireless networks with multiple broadcast domains. In *RTSS '07: Proceedings of the 28th IEEE International Real-Time Systems Symposium*, pages 447–458, Washington, DC, USA, 2007. IEEE Computer Society.

[25] Bhaskaran Raman, Kameswari Chebrolu, Naveen Madabhushi, Dattatraya Y. Gokhale, Phani K. Valiveti, and Dheeraj Jain. Implications of Link Range and (In)stability on Sensor Network Architecture. In *WiNTECH '06: Proceedings of the 1st International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, pages 65–72, New York, NY, USA, 2006. ACM.

[26] Matthias Ringwald and Kay Römer. BitMAC: A Deterministic, Collision-free, and Robust MAC Protocol for Sensor Networks. In *Proceedings of Second European Workshop on Wireless Sensor Networks (EWSN 2005), Istanbul, Turkey*, pages 57–69, january 2005.

[27] Jang-Ping Sheu, Chi-Hsun Liu, Shih-Lin Wu, and Yu-Chee Tseng. A priority MAC protocol to support real-time traffic in ad hoc networks. *Wireless Networks*, 10(1):61–69, 2004.

[28] J. L. Sobrinho and A. S. Krishnakumar. Quality-of-Service in Ad Hoc Carrier Sense Multiple Access Wireless Networks. *Selected Areas in Communications, IEEE Journal on*, 17(8):1353–1368, august 1999.

[29] J.L. Sobrinho and A.S. Krishnakumar. Real-time traffic over the ieee 802.11 medium access layer. *Bell Labs Technical Journal*, pages 172–187, autumn 1996.

[30] Kannan Srinivasan and Philip Levis. RSSI is Under Appreciated. In *Proceedings of the Third Workshop on Embedded Networked Sensors (EmNets)*, 2006.

[31] Hongyi Wu, Anant Utgikar, and Nian-Feng Tzeng. SYN-MAC: A Distributed Medium Access Control Protocol for Synchronized Wireless Networks. *MONET*, 10(5):627–637, 2005.

[32] Xue Yang and Nitin H. Vaidya. Priority scheduling in wireless ad hoc networks. In *MobiHoc '02: Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing*, pages 71–79, New York, NY, USA, 2002. ACM.

[33] Tiantong You, Chi-Hsiang Yeh, and Hossam Hassanein. CSMA/IC: A New Class of Collision-free MAC Protocols for Ad Hoc Wireless Networks. In *ISCC '03: Proceedings of the Eighth IEEE International Symposium on Computers and Communications*, pages 843–848, Washington, DC, USA, 2003. IEEE Computer Society.