International Journal of Distributed Sensor Networks, 3: 87–104, 2007

ISSN: 1550-1329 print/1550-1477 online
DOI: 10.1080/15501320601069515

Taylor & Francis
Taylor & Francis Group

Sensing with One or with Four? A Comparison of Two IEEE 802.15.x Protocols for Use in Sensor Networks

VOJISLAV B. MIŠIĆ and JELENA MIŠIĆ

University of Manitoba, Winnipeg, Manitoba, Canada

In this paper we compare the pertinent features of the two emerging technologies for wireless sensor networks: IEEE Standards 802.15.1 and 802.15.4. We review the main features of the MAC protocols defined by those standards, describe their operation, and compare them in terms of characteristics such as performance (access and end-to-end packet delays), bandwidth utilization, and scalability for the deployment of large networks. Our findings indicate that there is no clear winner in all categories; the best protocol (and the underlying technology) to use, are heavily dependent upon the requirements for a particular sensing application. While the main focus of our analysis is the MAC layer, certain important parameters of the Physical (PHY) layer are considered as well, together with some other networking aspects. The results of this analysis should be of interest to the designer and operators of wireless sensor networks.

Keywords IEEE Std 802.15.1; IEEE Std 802.15.4; Medium Access Control (MAC) Protocols; Bluetooth; Time-Division Multiplex (TDM); Carrier-Sense Multiple Access with Collision Avoidance (CSMA-CA)

1. Introduction

The wireless sensor network must satisfy a number of stringent and, quite often, conflicting requirements. Sensors must be small and cheap to produce and operate. They must operate on battery power, often for prolonged periods of time. They are often deployed or scattered in a given area without the knowledge of their precise physical position. At the same time, sensors must reliably monitor the physical phenomena of interest and gather data about it at a predefined rate. The values of the data rate span a wide range from a few bytes per second, for applications in forestry and agriculture, to perhaps tens of kilobytes per second, for surveillance and military applications; the traffic burstiness may also vary considerably among different applications.

Reliable event detection using minimal energy resources requires simultaneous achievement of several sub-goals. First, packet loss along the path from source to the sink has to be minimized; at the Physical (PHY) layer, packets can be lost due to noise and interference, while at the Medium Access Control layer (MAC) layer, losses may be incurred by collisions. (Since sensors are continuously monitoring the environment and sending data, retransmissions of lost packets are not necessary.) Packets can also be lost because of the limited capacity of device buffers: when the buffer is full when the new

Address correspondence to Vojislav B. Mišić, University of Manitoba, 75A Chancellors Circle, Winnipeg, Manitoba, R3T 2N2, Canada. Tel.: 1-204-474-8688; Fax.: 1-204-474-6791. E-mail: vmisic@cs.umanitoba.ca

packet arrives, the new packet may be discarded, or the packet at the head of the buffer can be dropped so that the new packet is accommodated.

Second, packet waiting has to be minimized, including queueing delays experienced in various devices along the data path, but also delays due to congestion in the network. (Queueing delays are the responsibility of the MAC layer, while congestion detection and control are performed at the transport layer.) Finally, packet propagation should take place along the shortest paths, while avoiding congested nodes and paths; this is the responsibility of the network layer.

Given that the protocol stack on sensor nodes—which are battery-operated and have limited computational capabilities—has to be as simple as possible, we conclude that simultaneous minimization of packet losses and improvement in efficiency (with the goal of maximizing the lifetime of the network) necessitate that some of the aforementioned functions of different layers are performed together. In other words, cross-layer optimization of network protocol operation is needed; the feasibility of this optimization is determined by the communication technology used to implement the network. All of the aforementioned optimizations are severely constrained by the limited computational and storage resources of the sensing nodes.

Obviously, the performance of a wireless sensor network is determined by a number of factors, not the least important of which is the choice of the communication and network technology, and the associated protocols [6]. In particular, performance parameters such as throughput, delay, and losses, are critically dependent on the network technology. As stated by [6], "the success of wireless sensor networks as a technology rests on the success of the standardization efforts to unify the market and avoiding the proliferation of proprietary, incompatible protocols that, although, perhaps optimal in their individual market niches, will limit the size of overall wireless sensor market."

Although many network technologies have been proposed in recent years, two among them are gaining widespread popularity as major candidates for implementation and deployment of sensor networks. One is Bluetooth [4, 5], which was originally envisaged as simple cable replacement technology, but subsequently developed into a generic network standard with many applications in wireless personal area networks; it has recently been standardized as IEEE Std 802.15.1 [11]. The other is IEEE Std 802.15.4 standard for low rate wireless personal area networks [13], which is the foundation for the recently adopted ZigBee standard [31]. Although both technologies use the same frequency band—the Industrial, Scientific, and Medical band at 2.4GHz[1]—their underlying protocols differ in many aspects. It is of interest, then, to compare these technologies from the viewpoint of their potential use in wireless sensor networks, and thus provide useful insights for the designers and operators of such networks.

Such comparison is the focus of the present paper, which will look at the number of characteristics of the Medium Access Control (MAC) layer, including network parameters (size) and device parameters (buffer size), as well as traffic parameters such as packet arrival rates, duty cycle, and duration of inactivity periods. Specifically, we will try to address the following questions:

1. What is the access delay inserted by the MAC protocol, which ultimately determines the freshness of the sensed information received by the sink?
2. What is the end-to-end packet delay achievable with the chosen MAC protocol?
3. Can the network reach the regime in which delays are unacceptably high?

[1]Note that 802.15.4. networks may operate in other frequency bands at 868 and 915MHz, but with much reduced data rates.

4. What is the effective bandwidth offered to the sensing application, i.e., what is the maximum possible event detection reliability for a particular MAC?
5. What amount of buffering at the source nodes is sufficient to guarantee acceptably low packet losses?
6. How easy is to integrate power/congestion control with a particular MAC?
7. How scalable is the design of large sensor networks using the given technology?

We will also investigate some important characteristics at the Physical (PHY) layer, most notably the immunity of transmissions to errors caused by noise and interference.

The paper is organized as follows. In Sections 2 and 3 we present the pertinent properties of 802.15.1 and 802.15.4 MAC protocols, respectively. The important characteristics of their respective PHY layers are mentioned as well. Section 4 is devoted to the comparison of these two technologies from the viewpoint of their suitability for deployment in sensor networks. Finally, Section 6 presents a brief summary and outlines the main implications of our research.

2. Basic Properties of Bluetooth Technology

We will first discuss the properties of Bluetooth technology, which is slightly more mature and certainly more widespread than the 802.15.4 technology. Bluetooth was originally intended as a simple communication technology for cable replacement [5]. However, its use has been steadily growing in a diverse set of applications [1]. Bluetooth operates in the Industrial, Scientific, and Medical band at 2.4GHz using Frequency Hopping Spread Spectrum technique, which makes it highly resilient to the noise and interference from other networks operating in the same band, such as IEEE 802.11b and IEEE 802.15.4 [9]. Each piconet hops through the available RF frequencies in a pseudo-random manner; the hopping sequence, which is determined from the Bluetooth device address of the piconet master, is known as the channel [3]. Each channel is divided into time slots of $T = 625\mu s$, which are synchronized to the clock of the piconet master.

2.1 Bluetooth Piconet Operation

Bluetooth devices are organized into piconets, small networks with up to eight active nodes and up to 255 inactive ones [5]. One of the active nodes is the master of the piconet; it invites other nodes to join the piconet as slaves. Within the piconet, Bluetooth uses a variant of the Time Division Multiple Access (TDMA) polling protocol, in which all communications are performed under the control of the piconet master. In Bluetooth, the master polls the slaves by sending them packets with appropriate identification and (when available) data. Slaves can talk back to the master only when addressed, and only immediately after being addressed by the master. As master (downlink) and slave (uplink) transmissions occur in alternative slots, this scheme is known as Time-Division Duplex (TDD); it is shown schematically in Fig. 1(b).

This approach is collision free and, consequently, more energy efficient than the collision-based MACs used in some other protocols, such as 802.11 and 802.15.4 [13].

2.2 Bluetooth Packet Types

The raw data rate of up to 1Mbps [4] in the original standard, or up to 3Mbps with the Enhanced Data Rate (EDR) option in version 2.0 of the standard [5], and the default transmission range of 10 to 100 meters make Bluetooth networks suitable for medium-rate

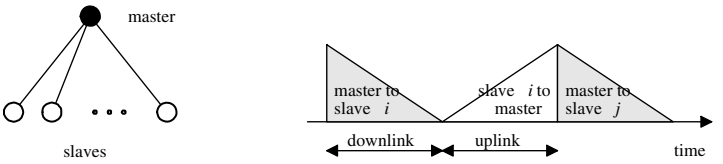

(a) A single piconet. (b) The Time Division Multiplex (TDD) mechanism.

FIGURE 1 Bluetooth operation.

TABLE 1 ACL packet types in Bluetooth. Packets of 2- and 3- types are available only under the ER option [5]

Type	Slot(s)	Payload (bytes)	FEC	Asymmetric data rate (kbps total)
DM1	1	17	2/3	217.6
DH1	1	27	none	341.6
DM3	3	121	2/3	516.2
DH3	3	183	none	692.0
DM5	5	224	2/3	514.1
DH5	5	339	none	780.8
2-DH1	1	54	none	791.2
2-DH3	3	369	none	1347.2
2-DH5	5	681	none	1563.7
3-DH1	1	85	none	1062.4
3-Dh3	3	554	none	2001.0
3-DH5	5	1023	none	2355.2

Wireless Personal Area Networks (WPANs). The same qualities mean that Bluetooth is suitable for the construction of low cost sensor networks, offering coverage of sensing areas with a diameter of several tens to several hundred meters [1].

While several types of logical links may be established in a Bluetooth piconet, it is likely that the sensor network will profit the most from the ACL logical links [5]. The data rate that can be achieved with different packet types is shown in Table 1.

2.3 Intra-piconet Polling Schemes

The polling scheme is obviously the main determinant of performance of Bluetooth piconets, and one of the main determinants of performance of Bluetooth scatternets. As usual, the main performance indicator is the end-to-end packet delay, with lower delays being considered as better performance. There are, however, at least two other requirements to satisfy. First, the piconet master should try to maintain fairness among the slaves, so that all slaves in the piconet receive equal attention in some shorter or longer time frame. (Of course, their traffic load should be taken into account.) Second, Bluetooth devices are, by default, low power devices, and the polling scheme should be sufficiently simple in terms of computational and memory requirements.

As noted above, the current Bluetooth specification does not specifically require or prescribe any specific polling scheme [5]. This may not seem to be too big a problem, since optimal polling schemes for a number of similar single-server, multiple-input queueing system are well known [16]. However, the communication mechanisms used in Bluetooth are rather specific, because

- all communications are bidirectional (i.e., there cannot exist a downlink packet without an uplink packet, or vice versa),
- the master polls the slaves using regular packets, possibly without data payload (i.e., all polls and responses thereto take at least one slot each),
- all slave-slave communications have to be routed through the master (i.e., there can be no direct slave-to-slave communication), and
- the master does not know the status of queues at the slaves, because there are no provisions for exchange of such information in the Bluetooth packet structure.

As a consequence, the existing results cannot be applied directly, and a number of polling schemes have been proposed and analyzed [7, 20]. These schemes may roughly be classified using the following criteria: the number of frames exchanged during a single visit to the slave, the bandwidth allocated to each slave, and the sequence in which slaves are visited. We note that the first criterion is probably the most important for sensor networks, although the others may also come into play in some sensing applications. Consequently, we mention only the *1-limited service* polling, in which the master visits each slave for exactly one frame, and then moves on to the next slave; *exhaustive service* polling, in which the master stays with the slave as long as there are packets to exchange in either the downlink or the uplink direction; and *E-limited service* polling, where the master stays with a slave until there are no more packets to exchange, or for a fixed number M of frames ($M > 1$), whichever comes first. In fact, 1-limited and exhaustive service polling may be considered as special cases of E-limited service, where the limit M equals 1 and ∞, respectively.

In traditional polling systems, exhaustive service performs better than either 1-limited or E-limited service [16]. In Bluetooth, however, the E-limited service has been found to offer the best performance; it even allows end-to-end delays to be optimized through judicious choice for the value of M [17].

2.4 On Power Saving Operation and Scatternets

A group of independent piconets interconnected through shared devices, or bridges, forms a scatternet. Depending on the role of the bridge, they are designated as Master/Slave or Slave/Slave bridges, as shown in Figs. 2(a) and 2(b) respectively. A master/slave bridge acts as the master in one piconet, and as a slave in another piconet. A slave/slave bridge acts as slave in both piconets. While the Bluetooth standard recognizes those two types of bridges, it poses no restrictions on the number of bridges that may participate in a piconet (provided the number of active slaves does not exceed seven), neither does it pose restrictions on the number of piconets that each bridge can participate in.

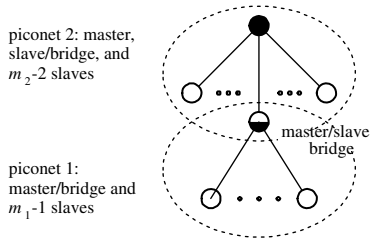

(a) Connection through a Master/Slave, or MS, bridge.

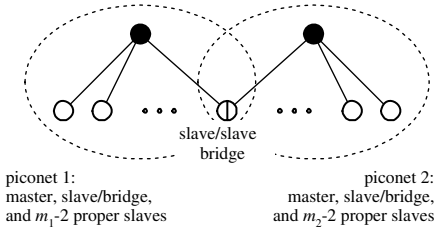

(b) Connection through a Slave/Slave, or SS, bridge.

FIGURE 2 Two Bluetooth piconets form a scatternet.

As most Bluetooth devices—bridges included—have only one radio interface, the bridge has to visit different piconets in different time periods. Consequently, both the intra-piconet polling scheme and the inter-piconet (bridge) scheduling scheme are important factors that determine the performance of a Bluetooth scatternet [20].

Bridge scheduling can use a fixed or adaptive sequence of *rendezvous* points – prearranged time when the bridge should be present in a particular piconet. The rendezvous schedule can be global for the entire scatternet, but the construction of a globally optimal schedule has been shown to be NP-complete, even under very favorable conditions [14]. Local rendezvous schedules can be made on a per-piconet basis, but this may give rise to synchronization problems as soon as there are three or more piconets in a scatternet. Fortunately, a simple yet scalable approach of *walk-in* scheduling has been proposed [21], in which the bridge visits each of its piconets in sequence, while the master of a piconet polls its bridges as if they were ordinary slaves; if they are present, a packet exchange is initiated. This approach can be combined with E-limited intra-piconet polling. The best performance in terms of end-to-end packet delays is obtained when the bridge remains in a given piconet during a single packet exchange, and then moves on to the next piconet [22].

It should be noted that the time during which a device can be absent from the piconet is limited in two ways. First, the Bluetooth standard contains provisions for the so-called *supervision timer*, the duration of which can be set within certain limits. If this timer expires, the piconet master is allowed to assume that the device has permanently left the piconet and to reassign its active address to another device. Second, the use of frequency hopping requires synchronization of slave device clocks to that of the piconet master. If a device is absent and, consequently, does not listen to the channel for a prolonged period of time, its clock may drift away from the master's clock and the device may not be able to synchronize with the piconet again. The latter mechanism is particularly important when designing the power saving schemes in which slave devices may be sent to sleep, as it imposes the limit on the duration of the sleep period.

We note that the implementation of Bluetooth scatternets is subject to another, rather serious constraint. Namely, while the Bluetooth standard allows the formation of scatternets by joining two or moro piconets, it does not support routing (since device addresses have local, intra-piconet significance only) and does not prescribe any particular bridge scheduling mechanism. As a result, all such support must be developed "from scratch," and actual implementations run the risk of being incompatible with one another. Still, bridge scheduling has been addressed by many researchers and significant advances have been made; but routing has received much less attention, and most of the work reported so far relies on the upper protocol layers to provide such support. In practice, most Bluetooth implementations do not attempt to build a scatternet. Instead, the piconet master acts as an access point that links the piconet to another, wired or wireless network, which serves as the backbone upon which more complex networks can be built. Although this approach is convenient for personal area networks, commercial and residential, it is rather unsuitable for sensor networks where that other networking infrastructure simply does not exist in most cases.

2.5 Power Saving Modes

There are three power saving mechanisms in Bluetooth. First, the master can park some slaves; up to 255 such slaves can exist in a given piconet. A parked slave gives up its active device address, which the master can reuse by inviting some other device to join the piconet. A parked slave has to be awakened by an explicit command of the master. Furthermore, the active mode of a Bluetooth device has two sub-modes or states [5]. In the SNIFF state, the device becomes active periodically; it listens for master transmissions for a predetermined period of time; and goes back to the SNIFF mode afterwards. In the HOLD

state, the device simply becomes inactive for a specified period of time and becomes active again when this period expires. In both states, the device retains its active device address within the piconet. Both SNIFF and HOLD states can be initiated by either the master or the slave device, and the duration of each state can be negotiated between the two.

Although all three aforementioned states are labeled as "power saving," the slave device in question may use the time of absence to join another piconet as a bridge, or do whatever else it chooses to.

Regarding the bridging function, the HOLD state is the preferred mechanism to implement it, for the following reasons. First, the duration of each HOLD state can be negotiated anew, whereas in the SNIFF state it is determined once for a sequence of SNIFF states. Second, there is no restriction on the time period in which the device is active upon returning from the HOLD state; in the SNIFF state this time is limited also.

2.6 Building Sensor Networks in Bluetooth

The suitability of Bluetooth as the platform to implement sensor networks has been investigated by building a Bluetooth protocol stack for the TinyOS operating system [15]. The experiments were conducted on actual Bluetooth devices, known as BTnodes, which were developed at ETH Zurich [2]. The BTnodes were equipped with two radios to enable multihop networking. The network was, then, tested for throughput and energy consumption. The results suggest that Bluetooth based sensor networks could be appropriate for event-driven applications that exchange bursts of data for a limited time period.

In another approach, the feasibility of an activity management approach with the goal of maintaining the sensor network reliability at the desired level, has been investigated in a simulated Bluetooth scatternet [24]. Reliability is defined as the number of packets received in a given time period at the network sink [30]. In this setup, each piconet used E-limited intra-piconet polling [17], and walk-in bridge scheduling [21]; both these approaches were shown to offer good performance and excellent scalability. The traffic model is derived from a relatively high-bandwidth, low cost surveillance based sensing application where compressed, still images are taken as a result of event detection and sent to the sink. (This setup may be used in applications such as road traffic control or asset protection.) Simulation results confirm that the adaptive activity management mechanism is capable of maintaining the reliability at the desired level while minimizing the buffer losses in all intermediate bridges.

3. Basic Properties of IEEE Std 802.15.4 MAC

The IEEE Standard 802.15.4 describes the MAC and PHY layers of a low rate WPAN (LR-WPAN) suitable for implementing networks with little or no infrastructure [13]. The other WPAN standards include Std 802.15.1, also known as Bluetooth [11], which is suitable for medium rate WPANs, as well as Std 802.15.3 [12], which is intended for high data rate WPANs.

3.1 802.15.4 Cluster Operation

In an IEEE 802.15.4-compliant WPAN, a central controller device commonly referred to as PAN coordinator builds a network, commonly referred to as a piconet or cluster[2], with

[2]In subsequent discussions we will use the latter term, in order to distinguish it from the Bluetooth piconet.

other devices within a small physical space known as the personal operating space. Two topologies are supported: in the star topology network, all communications, even those between the devices themselves, must go through the PAN coordinator. In the peer-to-peer topology, the devices can communicate with one another directly (as long as they are within the physical range), but the coordinator must be present. The standard also defines possible channel access mechanisms, depending on whether a beacon frame (sent periodically by the coordinator) is used to synchronize communications, or not. Beacon enabled networks use slotted carrier sense multiple access mechanism with collision avoidance (CSMA-CA), while the non-beacon enabled networks use simpler, unslotted CSMA-CA. In this work, we will focus on the performance of beacon enabled networks with slotted CSMA-CA; the unslotted access mechanism, being very similar to the one used in IEEE 802.11 standard, will not be considered here.

In beacon enabled networks, the coordinator divides its channel time into superframes [13]. Each superframe starts with the transmission of a network beacon, followed by an active portion and an optional inactive portion as shown in Fig. 3. The coordinator interacts with the nodes in its cluster during the active portion of the superframe, and may enter a low power mode during the inactive portion. The superframe duration, SD, is equivalent to the duration of the active portion of the superframe, which cannot be longer than the beacon interval BI. The active portion of the superframe is divided into 16 slots of equal size. Each slot consists of $3 \cdot 2^{SO}$ backoff periods which gives the shortest active superframe duration *aBaseSuperframeDuration* of 48 backoff periods.

The access mode in the CAP period is slotted CSMA-CA, similar to 802.11. In this case, the transmission of a packet begins with a backoff countdown, with the initial count chosen at random to avoid contention. If the active portion of the superframe ends while the countdown is in progress, the countdown will be frozen during the inactive portion of the superframe and will resume immediately after the beacon in the next superframe. After the countdown, the device listens to the channel to make sure it is idle; this is referred to as the Clear Channel Assessment (CCA). The standard prescribes two CCAs in two successive backoff periods, both of which must successfully pass for the transmission to begin.

A successful transmission is optionally acknowledged within a predefined time period; the timing constraints dictated by the standard preclude the possibility that the acknowledgment will collide with a packet transmission. The absence of acknowledgment indicates that the transmission has failed (which may be due to a collision with another transmission, or packet blocking because of buffer overflow) and must be repeated. If the remaining time after the countdown does not suffice for the two CCAs, the packet transmission, and subsequent acknowledgment, all of those activities are deferred to the next superframe. A more detailed discussion of the operation of the 802.15.4 MAC layer in the slotted CSMA-CA mode and its performance limitations can be found in [29].

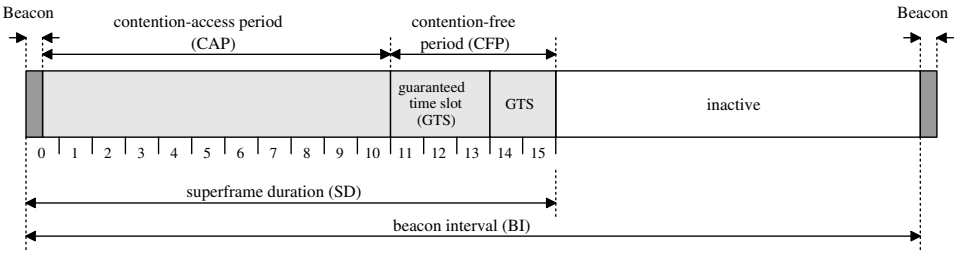

FIGURE 3 The composition of the superframe under IEEE Std 802.15.4, adapted from [IEEE, 2003*b*].

Within the time slots of the active portion of the superframe, the coordinator may reserve some slots to allow dedicated access to some devices. These slots are referred to as guaranteed time slots or GTS, and together they comprise the so-called contention-free period (CFP). Any device can request a GTS for an uplink and/or a downlink transmission, but the actual allocation is ultimately performed by the cluster coordinator. The list of allocated GTS slots is announced in the beacon frame; the device then simply waits until its designated slot and transmits the data without any backoff countdown and without any contention. The presence of GTS slots may be utilized for establishing a dedicated channel for certain devices, for example, the devices that serve as bridges to connect two or more clusters, as explained below. In this manner, the contention in the CAP period will be reduced (as there are fewer devices to use it), although the duration of the CAP period will be shortened.

The basic time unit of the MAC protocol is the duration of the so-called backoff period. Access to the channel can occur only at the boundary of the backoff period. The actual duration of the backoff period depends on the frequency band in which the network is operating. Namely, the standard allows the use of either one of three frequency bands: 868-868.6MHz, 902-928MHz and 2400-2483.5MHz (also known as ISM band). In the two lower frequency bands, BPSK modulation is used, giving the data rate of 20kbps and 40kbps, respectively. Each data bit represents one modulation symbol which is further spread with the chipping sequence. In the third band, the O-QPSK modulation is used before spreading; in this case, four data bits comprise one modulation symbol which is further spread with the 32-bit spreading sequence. The maximum packet payload, in this case, is 112 bytes or 13 backoff periods (nominally it is 127 bytes, but at least 15 bytes are taken up by the header). Both the packet payload and header are protected by a 16-bit CRC sequence, but there is no FEC or repetition.

3.2 An Example of Cluster Operation

The operation of this algorithm is illustrated in Fig. 4, where we show the concurrent activities of two devices *A* and *B*, together with the status and timing of the shared channel. For reasons of clarity we assume that transmissions do not require acknowledgements.

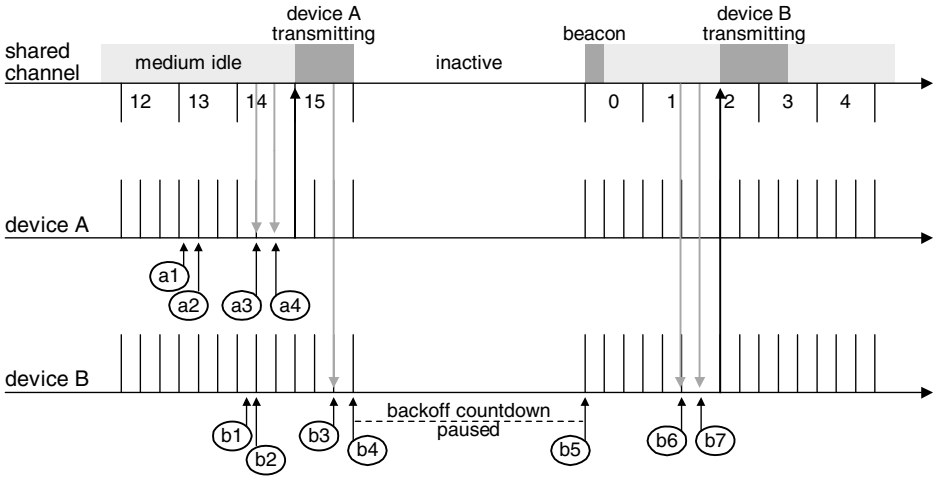

FIGURE 4 Pertaining to the operation of IEEE Std 802.15.4 MAC algorithm.

For device A, the events occur as follows. A new packet arrives (a1), the algorithm initializes the variables NB, *CW*, and *BE* to their initial values, and then locates the boundary of the next backoff period (a2). As *BE* has been set to *macMinBE*=2, the random backoff count is chosen as 3, and the waiting period lasts for three backoff periods. At the end of waiting, the MAC sublayer determines that the remaining time in the superframe is sufficient to accommodate all of the necessary activities. The first CCA (shown as a gray arrow) senses the medium to be idle (a3), and the value of *CW* is decreased to 1. After the next backoff period (a4), the medium is again sensed to be idle, the value of *CW* is decreased to 0, and device A may begin its transmission.

Device B receives a packet some time later (b1), initializes its variables *NB*, *CW*, and *BE*, and waits until the boundary of the next backoff period (b2). *BE* has also been set to 2, but the random backoff count is chosen as 4, and the waiting period lasts for four backoff periods. As the CCA senses a transmission in progress (b3), the algorithm moves on to step (4): it resets *CW* to 2, increases *NB* to 1, and increases *BE* to 3. The random backoff count is chosen as 7 and the backoff countdown begins. Since the superframe ends during the countdown (b4), it is paused until the next superframe, and is resumed at the next beacon frame (b5). When the waiting period is finished, the channel is sensed to be idle (b6), and the value of *CW* is decreased to 1. At the end of the next backoff period, the channel is again sensed to be idle, *CW* is decreased to zero, and device B can begin its transmission.

3.3 Interconnecting Clusters

IEEE 802.15.4 standard does not mention the bridging between clusters in beacon enabled mode. However, the bridging function is made possible by the existence of the inactive part of the superframe and, optionally, the presence of guaranteed time slots. During the inactive portion of the superframe, the nodes may enter a low power mode in order to conserve energy, or perform the bridging (interconnection) function. Let us now describe a possible scenario for a master/slave bridge, in which the shared device would act as the coordinator in one cluster and an ordinary node in another; these two clusters will be referred to as source and sink clusters, respectively. When the active part of the superframe is completed in the source cluster, its cluster coordinator (i.e., the bridge) switches to the sink cluster and waits for the beacon. When the beacon is detected in the sink cluster, the bridge can behave in one of two ways. If the GTS slots were previously reserved in the sink superframe for the purpose of cluster interconnection, bridge will wait until its dedicated GTS slot(s) and then transmit the packet immediately without undertaking the backoff procedure, as there can be no contention. If the bridge uses the CSMA-CA mechanism to deliver its data, it will execute the regular CSMA-CA transmission procedure, just like the other nodes in the sink cluster, until the end of the current superframe; then it can switch back to the source cluster. When the bridge (in CSMA-CA mode) is ready to move to its source cluster it freezes its backoff counter and resumes backoff counting upon return to the sink cluster. In this way, the bridge behaves as an ordinary node in a sink cluster, since all nodes have to freeze their backoff counters during the inactive portion of the superframe. Upon return to the source cluster, the bridge, acting as the coordinator, transmits the beacon and a new superframe may begin. Both the GTS based and the CSMA-CA based bridge operation are presented in Fig. 5.

Regarding synchronization of the source and sink superframes, the simplest case in the master/slave configuration is to have both clusters use the same beacon interval. The duration of the active portion of the source superframe should not exceed the corresponding active portion of the sink superframe. In this case, the synchronization is easy and can be accomplished in a short time. While other parameter values are also possible, they

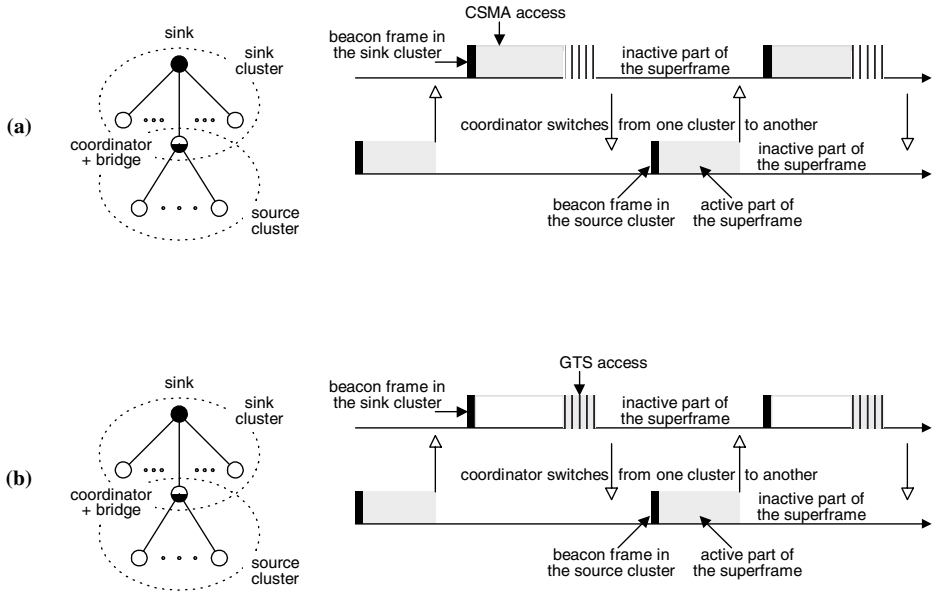

FIGURE 5 Operation of a two-cluster network with the bridge working in CSMA-CA and GTS mode respectively.

would lead to more complex synchronization procedures and longer packet delays. The same observation holds for the slave/slave bridge topology.

Some initial results indicate that such bridging function is indeed feasible [18]. Under small packet arrival rates and small cluster sizes, the use of CSMA-CA access mechanism for the bridge leads to better performance than the use of GTS. However, for moderate to high loads GTS based bridge significantly outperforms CSMA-CA bridge. In the cases of larger traffic in sink cluster than in the source, which can be expected if the sink serves more than one source cluster, the bridge should use GTS access in the sink cluster, as this approach guarantees data delivery to the sink and separates the collision domains which makes the design of larger networks scalable.

4. Comparison of Bluetooth and 802.15.4 Network Standards

After the individual description of IEEE 802.15.1 and IEEE 802.15.4 standards, we will give direct comparison of their properties from the standpoint of their feasibility for implementing a sensor network.

4.1 Access delay and Fairness

Bluetooth has a polling based MAC protocol and its access delay depends on the order in which the master polls the slaves, as well as on the amount of packets which are exchanged between master and slave in one visit. Mathematically speaking, the packet service time directly depends on the piconet cycle time, i.e., the time needed for the master to visit each slave in the piconet according to a given sequence. It has been shown that under low, symmetric traffic load, the exhaustive scheduling offers the lowest access delay, while 1- and E-limited policies perform better under high load [20]. Moreover, both 1- and E-limited policies have inherent fairness, unlike the

exhaustive polling where any given station can keep the master busy for a long period of time.

If the load is not uniform, slaves with lower traffic will experience lower access delays; some adjustment is possible if the E-limited policy is used.

Again, this is not too important an issue for sensor networks where the traffic load is expected to be way below the maximum value that the Bluetooth piconet can sustain.

In 802.15.4 networks operating in beacon enabled, slotted CSMA-CA mode, access delay also depends on the traffic load [25]. Regarding fairness, the built-in CSMA-CA contention resolution mechanism helps establish long-term fairness among nodes with comparable traffic intensity. However, in case of asymmetric traffic, the nodes with higher traffic can take up a disproportionately large share of the bandwidth and thus lead to deteriorating quality of service for nodes with low traffic. This property can be abused by malicious users to launch a Denial of Service attack against an 802.15.4 network [28].

4.2 Can the Network Reach the Regime when Delays are Unacceptably High?

A Bluetooth piconet can reach such a regime only if the duration of the piconet cycle becomes extremely long, which in turn means that exhaustive polling is used. This represents also a security problem, since one malicious node can bring down the entire piconet.

Under 1- or E-limited polling, this can happen only if a device attempts to send an inordinately high amount of traffic. Note, however, that only that device will experience high delays, while the traffic from other nodes will not suffer.

Note that all communications in the piconet are controlled by its master, which can adapt the polling policy dynamically to achieve the best performance under current traffic conditions.

An IEEE 802.15.4 network can reach the saturation regime if the traffic load—in other words, the number of nodes and their packet arrival rates—exceeds certain limits [27]. For example, for the packet size of 30 bytes (including PHY and MAC headers), saturation is reached with 30 nodes each having packet arrival rate of three packets per second (total of 45 bytes per second). Under the packet size of 90 bytes, saturation is reached with only 15 nodes at the same packet arrival rate.

Saturation also can represent a security problem since a couple of malicious nodes can quickly bring the network down, as shown in [28].

4.3 How much Buffering is Required?

Buffering refers to the fact that sensor devices have limited computational and storage abilities, hence the size of input buffers will by necessity be rather limited.

In many sensing applications, a single packet can hold the data for one measurement point, both in Bluetooth and 802.15.4 networks. In such cases, the amount of buffering is not very important. Namely, if a new packet finds that the input buffer is full, it can still be admitted but the packet which is currently awaiting processing at the head of the buffer will be dropped. This buffer management policy (often referred to as push-out) ensures that the sensing application will always receive the most recent data to work with. Alternatively, the new packets may simply be dropped until space becomes available in the buffer (the blocking policy). Note that the choice of the buffer management policy does not affect the throughput; but the end-to-end packet delays will be lower with the push-out policy.

In some applications, such as surveillance and health care, one measurement can require several packets: for example, a low resolution image from a surveillance camera or a set of measurements from an ECG or EEG probe. One measurement thus results in a

packet burst at the PHY and MAC levels. In such cases, the size of the buffer should be sufficient to hold the entire packet burst, as splitting it into several chunks may increase the delays and complicate further processing. Moreover, intra-piconet polling should be done using E-limited policy, so as to make sure that the entire burst is preserved in transmission [19]. On the other hand, an 802.15.4 network is unable to guarantee that the burst will be preserved, regardless of its operating mode. However, its lower maximum data rate makes it rather inappropriate for use in applications where bursty traffic is to be expected.

At the same time, exhaustive scheduling with large buffers increases spatial and temporal correlation of sensed data, which may be undesirable from the security point of view – a malicious node should not be allowed to inject large amounts of useless or bogus data into the network. Therefore, for most applications, the size of device buffers should be dimensioned to hold several packets, possibly five to ten at most.

4.4 What is the Effective Bandwidth Left to the Application, i.e., what is the Maximum Possible Event Detection Reliability for a Particular MAC?

The main concern in sensor network applications of Bluetooth is that the actual polling of slaves must use one slot packets, even if there is no actual data to send. In the worst case, which will limit the maximum throughput of the network to only around 200kbps. However, when the uplink traffic uses longer packets and/or the EDR option, the maximum throughput can reach much higher values. This makes Bluetooth piconet suitable for applications that require higher bandwidth, such as surveillance and military ones.

It should be noted that applications that require full bandwidth available in Bluetooth may pose other problems as well. In particular, the size of bridge buffers in Bluetooth scatternets may become a critical parameter when a bridge carries the traffic from several piconets toward the network sink. Some preliminary results indicate that buffers which hold 20 to 40 baseband packets can enable a Bluetooth scatternet to sustain the traffic load close to its rated capacity [23].

On the other hand, IEEE 802.15.4 has a maximum raw data rate of 250kbps, which is only one-third or less of the maximum capacity in Bluetooth. Unlike Bluetooth, where all communications are managed by the piconet master, in 802.15.4 networks the traffic originating at one node affects the activities of the others. Under large traffic which can originate from many nodes, there will be many collisions and many deferred transmissions. This results in severe congestion and all nodes experience large delays. In such a situation which is termed as saturation, throughput drops to few percent of the raw data rate. Since in IEEE 802.15.4 backoff window cannot exceed the value of 31 and the packet size is limited to 127, network can easily reach the saturation regime. Our results show that the highest throughput of 25% occurs for a packet size of 90 bytes (including PHY and MAC headers), 5 active stations in the network (we did not check for smaller number station) under superframe size of 48 backoff periods. This puts a limit of effective data rate of 62.5 kbps per cluster, or around 12.5 kbps per node. However, under 15 active nodes the total throughput drops to 18% and this drop continues with the increase of the cluster size.

4.5 How Easy is to Integrate Power/Congestion Control with the MAC?

Among the most important requirements for sensor networks is the maximization of their lifetime due to high costs (and sometimes even infeasibility) of maintenance activities. Sensor lifetime can be extended by adjusting the frequency and ratio of active and inactive periods of the network – essentially, by sending all the nodes in the network to sleep for some time [30].

However, many sensor applications (such as surveillance, health care, and structural health monitoring) require continuous monitoring of relevant variables and events. In such cases, letting the network sleep for some, even very short, time is simply out of the question. In such cases, duty cycle management can be applied at the level of individual sensor nodes, provided the number of sensors covering a given physical area is larger than the minimum number based on the required data rate. The desired packet rate received at the sink can then be achieved by adjusting the number of sensors that are active at any given time. However, this activity control must be coupled with the MAC layer and physical layer in order to be able to distinguish congestion from inadequate activity of the nodes. This approach has the additional benefit of extending the network lifetime, since the duty cycle of individual sensors can be adjusted to any desired value. For example, a duty the duty cycle of 0.5% will allow a single AAA battery to power an off-the-shelf radio transceiver at $10mA$ for two years [10].

Cross-layer activity management algorithms have been developed for Bluetooth scatternets [24]. They essentially control the number of slaves which are active at any given time, based on the feedback information about the actual vs. desired reliability. (We have discussed the power saving modes available in Bluetooth in Section 2.5.)

Power management is also supported by the 802.15.4 standard in its beacon enabled mode with slotted CSMA-CA, where the interval between the two beacons is divided into active and inactive parts, and the sensors can switch to low-power mode during the inactive period as shown in Fig. 3. In case redundant sensors are used, scheduling of active and inactive periods along the lines discussed above can also be used. Some initial results obtained by implementing this approach with both centralized and distributed algorithms are encouraging [26]. While both approaches have similar accuracy, the distributed algorithm is simpler to implement and requires less computational resources than the centralized one.

4.6 How Scalable is the Design of Large Sensor Networks using given Technology?

Walk-in bridge scheduling in Bluetooth allows each bridge to move between piconets at its own time points. Therefore, the design of large sensing scatternets is only the issue of event sensing reliability required at the sink and the number of piconets included in the sensing. Under ideal conditions each piconet should contribute the equal portion of the total event sensing reliability, which on the other hand can not be larger than the effective piconet bandwidth for the given packet type and size. Individual event sensing reliability of the piconet should be used to control the number of active slaves. We have to note that total event sensing reliability at the sink would be affected by the losses at intermediate bridges which brings the issues of topology, routing, and the bridge's buffer size.

Regarding the use of 802.15.4 for larger sensor networks, we note that this can be achieved using GTS based bridges. The overall network size should be carefully designed so that event sensing reliability of all clusters does not exceed the capacity of the sink. Also, the packet size, the total number of sensing nodes, and the sleeping technique should be chosen for each sensing cluster. At the network level, the bridge buffer size and routing algorithm are important design parameters.

4.7 Issues at the PHY Layer

Finally, let us discuss an important issue at the PHY layer, namely the immunity to noise and interference. Both the 802.15.1 and 802.15.4 networks operate in the ISM band at 2.4GHz, which is already hosting wireless LAN/PAN standards such as 802.11b as well as

other radiation sources, and a lot of interference may be expected. As mentioned, Bluetooth uses FHSS and is very resilient to interference, as several reports have found [9]. According to the exhaustive simulation results reported in [32], when ten fully loaded piconets with seven slaves each, operate in the room with dimensions 10m × 20m, the packet error rate for DH1 packets was only around 0.03; this rate increases to 0.3 when the same experiment was repeated with 100 co-located piconets.

When operating in the ISM band, IEEE 802.15.4 networks use a 16-ary quasi-orthogonal modulation technique. Four data bits represent one modulation symbol and that symbol is further encoded into a 32 bit chip sequence. There are 16 nearly-orthogonal Pseudo-Noise chip sequences. Each chip sequence is modulated onto the carrier using offset quadrature phase shift keying (O-QPSK). Since the chip rate is 2Mcps and raw data rate is 250kbps, the maximum supported ratio of bit energy to the noise power spectral density of $\frac{E_b}{N_0} = 8$. According to the properties of QPSK, the Bit Error Rate is determined using known expression given for example in [8]. Therefore, without the interference, we should expect BER slightly less than 10^{-4}. This is confirmed in the section 6.1.6 of the standard where the Packet Error Rate (PER) of 1% is expected on packets which have 20 bytes including MAC and physical level headers. However, in the presence of interference in the ISM band, it is more realistic to expect BER around 10^{-3} and Packet Error Rate more than 28% for packets with 27 bytes of payload and 15 bytes of headers. (Packet Error Rate can be calculated as $PER = 1 - (1 - BER)^X$ where X is packet length including MAC and physical layer header expressed in bits).

Although the experiments in [32] cannot be directly translated into BER, the interference due to ten co-located piconets is likely to be way above the maximum value that the PHY layer of an 802.15.4 cluster could handle.

Moreover, other emission sources, such as 802.11 devices, are more likely to damage 802.15.4 networks which use the entire band, than Bluetooth networks which hop in a random fashion and do not stay long enough at any given 1MHz channel within the ISM band.

Therefore, we may conclude that Bluetooth does perform better in terms of immunity to noise and interference.

In terms of transmission range, both standards prescribe a similar range—the default value of which is around 10 meters or 30 feet. The range may be adjusted by tuning the power of the device transmitter, provided the radio subsystem hardware supports power management.

5. Matching the Technology to the Application

Given the considerations presented in the last Section, we can identify the sensing applications for which one or the other communications technology would be well suited.

Bluetooth appears well suited for applications that carry higher traffic load and/or bursty traffic. With the recent addition of the EDR option, even image transmission becomes possible. The range of applications that fit the above description is wide: it includes health care, health monitoring of large civil structures such as bridges, buildings, and the like, surveillance, and military. The last two areas are well served by the fact that all communications in Bluetooth are performed under the control of the piconet masters, as well as the relative insensitivity of Bluetooth communications to noise and interference.

Bluetooth can also perform well in more complex applications where several piconets are connected to form a larger network. However, in order to obtain satisfactory performance,

intra-piconet polling must use the E-limited policy, while bridge scheduling must be performed using the walk-in approach.

The downsides of Bluetooth include, first and foremost, the small number of active devices in the piconet and the comparatively complex procedure through which devices go in and out of the active state. Also, the Time-Division Duplex communications impose a performance penalty on communications that, on their way from the source to the network sink, need to pass through one or more bridges.

Networks operating under the 802.15.4 standard are suitable for applications which do not require high bandwidth, as is the case with many environmental, agricultural, logistics, and other applications. An 802.15.4 network can easily include tens or even hundreds of devices – certainly many more than a Bluetooth network. More complex networks can easily be constructed from 802.15.4 clusters, and adaptive power management can be used to optimize energy consumption while maintaining the desired throughput at the network sink. (Comments about walk-in scheduling apply to 802.15.4 networks as well.)

However, an 802.15.4 network can easily get saturated, due to the use of self-regulating contention-based channel access. In such cases, the network effectively ceases to function. This makes it vulnerable to hostile behavior, even to simple attacks that would pose only a limited threat to a comparable Bluetooth network.

6. Conclusion

In this paper we have compared IEEE 802.15.1 (also known as Bluetooth) and 802.15.4 WPAN technologies from the aspect of their possible deployment in sensor networks. Our discussion shows that both technologies can be used to implement large-scale sensor networks over large physical areas. Bluetooth offers much higher effective throughput per piconet, and is more resistant to DoD attacks, than 802.15.4. On the other hand, 802.15.4 networks support piconets with higher number of devices, and the procedure to join or leave the network is simpler than in Bluetooth. Overall, both network technologies appear suitable for use in sensor networks, provided those observations are taken into account.

About the Author

Vojislav B. Mišić received his PhD in Computer Science from the University of Belgrade, Yugoslavia, in 1993. He is currently Associate Professor of Computer Science, at the University of Manitoba in Winnipeg, Manitoba, Canada. His research interests include systems and software engineering and modeling and performance evaluation of wireless networks. He is a member of ACM, AIS, and IEEE.

Jelena Mišić received her PhD degree in Computer Engineering from the University of Belgrade, Yugoslavia, in 1993. She is currently Associate Professor of Computer Science at the University of Manitoba in Winnipeg, Manitoba, Canada. Her current research interests include wireless networks and security in wireless networks. She is a member of the IEEE Computer Society and SCS society.

References

1. O. B. Akan and I. F. Akyildiz. ESRT: Event-to-Sink Reliable Transport in Wireless Sensor Networks. In *IEEE/ACM Transaction on Networking (to appear)*, October 2005.
2. J. Beutel, O. Kasten, and M. Ringwald. BTnodes – A Distributed Platform for Sensor Nodes. In *Proc. 1st Intl. Conf. on Embedded Networked Sensor Systems (SenSys)*, pp. 292–293, November 2003.

3. Bluetooth SIG. *Specification of the Bluetooth System*. Version 1.1, Feb. 2001.
4. Bluetooth SIG. *Specification of the Bluetooth System*. Version 1.2, Nov. 2003.
5. Bluetooth SIG. *Core Specification of the Bluetooth System*. Version 2.0 + EDR, Nov. 2004.
6. E. H. Callaway, Jr. *Wireless Sensor Networks, Architecture and Protocols*. Auerbach Publications, Boca Raton, FL, 2004.
7. A. Capone, R. Kapoor, and M. Gerla. Efficient polling schemes for Bluetooth picocells. In *Proceedings of IEEE International Conference on Communications ICC 2001*, vol. 7, pp. 1990–1994, Helsinki, Finland, June 2001.
8. V. K. Garg, K. Smolik, and J. E. Wilkes. *Applications of CDMA in Wireless/Personal Communications*. Prentice Hall, Upper Saddle River, NJ, 1998.
9. N. Golmie, R. E. Van Dyck, and A. Soltanian. Interference of Bluetooth and IEEE 802.11: simulation modeling and performance evaluation. In *Proceedings 4th ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems*, pp. 11–18, Rome, Italy, July 2001.
10. J. A. Gutiérrez, E. H. Callaway, Jr., and R. L. Barrett, Jr. *Low-Rate Wireless Personal Area Networks*. IEEE Press, New York, NY, 2004.
11. IEEE. Standard for part 15.1: Wireless medium access control (mac) and physical layer (PHY) specifications for wireless personal area networks (WPAN). IEEE standard 802.15.1, IEEE, New York, NY, 2002.
12. IEEE. Standard for part 15.3: Wireless MAC and PHY specifications for high rate WPAN. IEEE Std 802.15.3, IEEE, New York, NY, Sept. 2003.
13. IEEE. Standard for part 15.4: Wireless MAC and PHY specifications for low rate WPAN. IEEE Std 802.15.4, IEEE, New York, NY, Oct. 2003.
14. N. Johansson, U. Körner, and L. Tassiulas. A distributed scheduling algorithm for a Bluetooth scatternet. In *Proceedings of the International Teletraffic Congress – ITC-17*, pp. 61–72, Salvador de Bahia, Brazil, Sept. 2001.
15. M. Leopold, M. Dydensborg, and P. Bonnet. Bluetooth and Sensor Networks: A Reality Check. In *1st ACM Conference on Sensor Networks*, November 2003.
16. H. Levy, M. Sidi, and O. J. Boxma. Dominance relations in polling systems. *Queueing Systems Theory and Applications*, **6**, 2, 155–171, 1990.
17. J. Mišić, K. L. Chan, and V. B. Mišić. Admission control in Bluetooth piconets. *IEEE Transactions on Vehicular Technology*, **53**, 3, 890–911, May 2004.
18. J. Mišić, J. Fung, and V. B. Mišić. Interconnecting 802.15.4 clusters in master-slave mode: queueing theoretic analysis. In *Proc. I-SPAN 2005*, Las Vegas, N V, Dec. 2005.
19. J. Mišić and V. B. Mišić. Performance analysis of Bluetooth piconets with finite baseband buffers. *Wireless Communications and Mobile Computing*, **5**, 8, 917–925, Dec. 2005.
20. J. Mišić and V. B. Mišić. *Performance Modeling and Analysis of Bluetooth Networks: Network Formation, Polling, Scheduling, and Traffic Control*. CRC Press, Boca Raton, FL, July 2005.
21. J. Mišić, V. B. Mišić, and K. L. Chan. Walk-in bridge scheduling in Bluetooth scatternets. *Springer Science, Cluster Computing*, **8**, 3, 197–210, 2005.
22. J. Mišić, V. B. Mišić, and K. L. Chan. Performance of Bluetooth scatternets under E-limited polling and walk-in bridge scheduling. *Dynamics of Continuous, Discrete and Impulsive Systems, Series B: Applications & Algorithms*, **13**, 1, 49–76, Feb. 2006.
23. J. Mišić, V. B. Mišić, and G. R. Reddy. On the performance of Bluetooth scatternets with finite buffers. In *Proc. WWAN 2005 IBnternational Workshop on Wireless Ad Hoc Networking (ICDCS'05 Workshops)*, pp. 865–870, Columbus, OH, June 2005.
24. J. Mišić, G. R. Reddy, and V. B. Mišić. Activity Scheduling based on cross layer information in Bluetooth sensor networks. *Computer Communications*, to appear, 2006.
25. J. Mišić, S. Shafi, and V. B. Mišić. Performance of 802.15.4 beacon enabled PAN with uplink transmissions in non-saturation mode – access delay for finite buffers. In *Proc. BroadNets 2004*, pp. 416–425, San Jose, CA, Oct. 2004.
26. J. Mišić, S. Shafi, and V. B. Mišić. Activity management through Bernoulli scheduling in 802.15.4 sensor clusters. In *Proc. BroadNets 2005*, Boston, MA, Oct. 2005.
27. J. Mišić, S. Shafi, and V. B. Mišić. The impact of MAC parameters on the performance of 802.15.4 PAN. *Ad hoc Networks*, **3**, 5, 509–528, Sept. 2005.

28. V. B. Mišić, J. Fung, and J. Mišić. MAC layer security of 802.15.4-compliant networks. In *Proc. WSNS'05, held in conjunction with IEEE MASS05 2005*, Washington, DC, Dec. 2005.
29. V. B. Mišić, S. Shafi, and J. Mišić. Avoiding the bottlenecks in the MAC layer in 802.15.4 low rate WPAN. In *Proc. HWISE2005 (ICPADS 2005 Workshops, vol. 2)*, pp. 363–367, Fukuoka, Japan, July 2005.
30. Y. Sankarasubramaniam, Ö. B. Akan, and I. F. Akyildiz. ESRT: event-to-sink reliable transport in wireless sensor networks. In *Proc. 4th ACM MobiHoc*, pp. 177–188, Annapolis, MD, June 2003.
31. ZigBee Alliance. ZigBee specification. ZigBee Document 053474r06, Version 1.0, ZigBee Alliance, San Ramon, CA, June 2005.
32. S. Zrbes. Considerations on link and system throughput of Bluetooth networks. In *Proceedings of the 11th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC 2000*, vol. 2, pp. 1315–1319, London, UK, Sept. 2000.