

A wireless sensor networks MAC protocol for real-time applications

E. Egea-López · J. Vales-Alonso · A. S. Martínez-Sala · J. García-Haro · P. Pavón-Mariño · M. V. Bueno Delgado

Received: 16 May 2006 / Accepted: 29 June 2006 / Published online: 3 November 2006
© Springer-Verlag London Limited 2006

Abstract Wireless sensor networks (WSN) are designed for data gathering and processing, with particular requirements: low hardware complexity, low energy consumption, special traffic pattern support, scalability, and in some cases, real-time operation. In this paper we present the virtual TDMA for sensors (VTS) MAC protocol, which intends to support the previous features, focusing particularly on real-time operation. VTS adaptively creates a TDMA arrangement with a number of timeslots equal to the actual number of nodes in range. Thus, VTS achieves an optimal throughput performance compared to TDMA protocols with fixed size of frame. The frame is set up and maintained by a distributed procedure, which allows sensors to asynchronously join and leave the frame. In addition, duty cycle is increased or decreased

in order to keep latency constant below a given deadline. Therefore, a major advantage of VTS is that it guarantees a bounded latency, which allows soft real-time applications.

Keywords Wireless sensor networks · MAC protocols · Real-time applications

1 Introduction

Wireless sensor networks (WSNs) are a new paradigm of telecommunication networks. WSNs are designed to perform efficient data collection and environment monitoring, among other applications. WSNs share key properties with Mobile Ad-hoc NETWORKS (MANETs): decentralized control, wireless broadcast nature, self-configuring capabilities, multi-hop routing and ephemeral topologies. However, unlike MANETs, WSNs must support: (a) specific traffic patterns, characterized by very long idle periods and sudden peak transmissions, (b) long run battery-powered deployment, which yields to tight energy constraints, and (c) device (hardware and software) simplicity. Therefore, two fundamental goals of WSN protocols are energy saving and traffic/environment adaptivity.

In addition, there are new incoming proposals of combined sensor and actors (devices that act upon events) networks [1], yielding the so-called Wireless Sensor and Actor Network (WSAN) model. WSANs are by nature alarm-driven systems, where the reaction time (from sensor detection to actor action) must usually be bound. Thus, WSAN proposals must add real-time operation as a requirement for their associated protocols.

E. Egea-López (✉) · J. Vales-Alonso · A. S. Martínez-Sala · J. García-Haro · P. Pavón-Mariño · M. V. Bueno Delgado
Department of Information Technologies and Communications, Polytechnic University of Cartagena, Cartagena, Spain
e-mail: esteban.egea@upct.es

J. Vales-Alonso
e-mail: javier.vales@upct.es

A. S. Martínez-Sala
e-mail: alejandros.martinez@upct.es

J. García-Haro
e-mail: joang.haro@upct.es

P. Pavón-Mariño
e-mail: pablo.pavon@upct.es

M. V. Bueno Delgado
e-mail: mvictoria.bueno@upct.es

WSNs major sources of energy waste are related to radio communication issues [2]. Namely, *collisions*, *idle listening*, *overhearing* packets addressed to other nodes, and *packet overhead* (sending and receiving too many control packets).

Since nodes do not know when they will receive packets from their neighbors, they are always listening to the channel (idle listening) and the radio is kept in receiving mode, consuming energy. Stemm and Katz [3] state that idle listening is the dominant factor. Thus, radios must be *turned off* during periods of inactivity to save energy.

Besides, sudden traffic peaks are likely to happen in WSNs. High loads may collapse the network, degrading its performance (throughput and latency) and raising power consumption. Consequently, adaptation to extreme situations is mandatory for WSN protocols. Device limitations (both hardware and software), additionally impose that algorithms and protocols be simple.

In this paper we propose the virtual TDMA for sensors(VTS) Medium Access Control (MAC) protocol. VTS provides a Time Division Multiple Access (TDMA) access scheme, in which the number of available slots *dynamically* adjusts to the number of nodes present in a cell (cluster) of nodes. Such a mechanism, after a transient adjustment phase, leads to a *scalable* and collision-free MAC protocol that consumes *considerably less energy* than contention-based protocols and *has a bound packet latency* (providing support for soft real-time services). VTS also addresses network setup and synchronization issues. The trade-off is the average latency, which is slightly worse than contention protocols under low/medium loads.

As most of the sensor network proposals [4] VTS periodically puts nodes to sleep to reduce power consumption, which results in *listen/sleep* cycles. Our protocol employs a synchronization procedure similar to S-MAC [2] to establish the listen/sleep schedule. However, unlike S-MAC, only one node can transmit in every listen/sleep cycle. Thus, every cycle becomes what in a TDMA context is called a *timeslot*.¹ By following an extremely simple procedure, the nodes in a cluster will transmit in different timeslots. Therefore, when each node is finally transmitting in a different timeslot, a *frame* of timeslots has been built in a *distributed way*. VTS allows frame adjustment, that is, to increase or reduce the number of timeslots, which

improves throughput compared to a TDMA frame with a fixed number of timeslots. With this TDMA-like access there is no contention for data transmission and latency is guaranteed. Finally, in order to meet real-time deadlines, VTS nodes adapt the duration of the sleep interval (*dynamic duty cycle*). This way, when new nodes join a TDMA frame, the sleep interval is reduced in order to keep constant maximum latency.

The rest of the paper is organized as follows: Sect. 2 contains related work on MAC protocols for WSNs. Section 3 introduces a suitable application scenario for the VTS protocol. In Sect. 4 the basis of the S-MAC protocol is reviewed to introduce VTS synchronization procedure. Section 12 thoroughly describes the VTS protocol. A performance analysis of VTS is presented in Sect. 12. Finally, Sect. 19 concludes and suggests future work.

2 Related work

A considerable research effort has been devoted to WSNs in the last few years. Many new protocols and applications are currently being proposed and tested. WSN MAC protocols focus mainly on energy efficiency. Latency in message delivery is not usually a metric to be optimized. Most of the proposals can be classified in one of the classical categories: contention or TDMA based.

MAC contention protocols are simple, scalable and flexible. Their major drawback is a high idle listening time (the dominant factor of energy waste). WSN contention-based proposals presently extend the Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) mechanism, applying additional schemes to reduce overhearing and idle listening:

- (1) Out-of-band signaling requires additional radio channels [5], and hardware is more complex and expensive.
- (2) Coordinated scheduling of listen time, which was first proposed by the sensor-MAC (S-MAC) protocol [2].

S-MAC introduces a procedure to synchronize nodes on a common structure that yields a shared listen/sleep cycle among neighbor nodes. This schedule reduces idle listening and, therefore, energy consumption.

The timeout-MAC (T-MAC) [6] protocol improves S-MAC by using an adaptive cycle length. The listen/sleep interval duration adapts to traffic fluctuations and obtains a better energy profile. This family of MAC protocols is relatively simple but does not guarantee

¹ In this paper, we refer to a listen/sleep cycle as timeslot, cycle or frame, depending on the context. A set of listen/sleep cycles is called a superframe. The *duty cycle* is the ratio of listen interval to total time $DC = \frac{T_{listen}}{T_{slot}}$.

latency. In contrast, with a similar complexity, our protocol keeps latency bounded (see Sect. 12).

TDMA protocols assign timeslots to nodes, avoiding collisions and idle listening. However, in ad-hoc and sensor networks, establishing and maintaining a superframe of timeslots is a complex task. In addition, if the number of nodes dynamically changes, which is likely to occur in WSNs, scheduling must be readapted. All TDMA proposals for WSNs (and MANETs) utilize contention stages to setup and maintain a properly organized TDMA. Our protocol also belongs to this category.

There are a number of these proposals for MANETs and WSNs:

- (1) The Five Phase Reservation Protocol (FPRF) [7], which provides a distributed algorithm to solve the problem of slot allocation in multi-hop networks. FPRF allocation procedure performs well at the expense of a great complexity, and does not implement any energy saving mechanism.
- (2) In Eyes MAC (EMAC) [8] a node can be active or passive. Active nodes own a timeslot and form a network backbone that performs routing tasks. Passive nodes use contention periods to send data. EMAC is focused on the increase of network lifetime, whereas latency or throughput is not addressed.
- (3) The Lightweight Medium Access Protocol (LMAC) protocol [9] is a modification of EMAC in which each node selects a timeslot using slot occupancy information from its one-hop neighbors. Its main limitations are that the number of available slots is fixed and the nodes listen to unused slots.

Therefore, LMAC latency and throughput degrades at low loads [4]. In comparison with these proposals,

VTS is simpler and does not fix the number of timeslots; therefore, it achieves a better performance.

3 Motivating scenario

In this section, we introduce an example of a suitable application for VTS. Let us assume a storage area for dangerous chemicals. This storage area may be a warehouse or a freight dock. Chemicals are carried in container or barrels. Every one of these containers carries a wireless sensor, which is able to periodically measure critical safety parameters of its load. These critical parameters must arrive to a control center before a given deadline (T_d). Let us also assume that

containers enter or leave the storage area in random way.

In this scenario it is clear that sensors must be able to guarantee a bound delivery time for packets in order to fulfill the deadline. Since nodes (containers) leave and join randomly the area, the network must be capable of self-organization. The VTS protocol achieves both goals, as we will see in the following sections.

4 S-MAC overview

S-MAC [2] is a contention-based protocol that reduces energy consumption by means of several mechanisms.


Periodic listen and sleep forces nodes to activate periodically for a small time interval (the listen period); the rest of the time the nodes turn off their radio and sleep (the sleep period). A listen/sleep cycle is also called a frame (see Fig. 1). The ratio of the listen interval to the sleep interval is the duty cycle. Neighbors achieve and maintain a coordinated sleeping time, synchronizing their listen/sleep schedules by means of the short SYNChronization (SYNC) packet.

SYNC packets correct clock drifts and are used to discover new neighbors. In a stationary situation, each node broadcasts a SYNC packet after a fixed number of frames (N_c) to maintain synchronization. Within a frame, the listen interval is subdivided into SYNC period (for SYNC packets) and data period (for data packets), as shown in Fig. 1. Nodes perform carrier sense during a random number of slots (contention) before transmitting SYNC. If two nodes contend for transmitting a SYNC in the same cycle, it may happen that:

- (1) Nodes choose a different number of carrier sense slots. As a result, the node with a higher number defers its SYNC transmission (losing contention), and makes another attempt in the next synchronization cycle.
- (2) Both nodes select the same slot. A collision occurs which is not detected by any of them. After N_c cycles, both nodes contend again for SYNC transmission. Transmission of an information packet occurs in the data period inside the frame (see Fig. 1). Nodes can make use of the data period in any frame: synchronization and data periods operate independently.

Collision avoidance is based on CSMA/CA. It uses a request to send(RTS)/clear to send(CTS)/data/ACK sequence, with a fixed back off contention window. Notice that S-MAC uses two independent periods of

Fig. 1 S-MAC listen-sleep frame



contention in every cycle, one for SYNC and one for data transmission. To avoid *overhearing*, all the nodes sleep either at the beginning of the sleep period or immediately after receiving a RTS or a CTS not addressed to them, and they wake up when the next frame starts. This scheme (periodic listen and sleep) significantly reduces idle listening. However, there is an undesirable effect on packet latency, because nodes must wait for the next listen period to send their data. To overcome this issue, S-MAC proposes a technique called *adaptive listening*: nodes which overhear a RTS or CTS packet wake up at the end of the transmission, instead of waiting for their next scheduled listen time. Thereby, if a node is the next-hop destination, its neighbor is able to immediately pass the data to it.

Summarizing, S-MAC reduces idle listening and provides an efficient mechanism to synchronize nodes. Nevertheless, it does not act upon the other major sources of energy waste: collisions and packet overhead. It even increases them: there are two contention intervals (SYNC and data) every cycle, and a new control packet (SYNC). Moreover, S-MAC cannot guarantee packet latency.

4.1 S-MAC synchronization

In S-MAC, when a node initializes, it keeps listening for a certain amount of time. If it receives a SYNC packet, it adopts its listen/sleep schedule and tries to send its own SYNC in the next available chance. Otherwise, the node chooses its own schedule and broadcasts it using a SYNC packet. After N_C listen/sleep cycles, nodes broadcast a new SYNC packet to maintain synchronization.

Figure 2 illustrates this effect for a network with M neighbor nodes, being node A_1 the first node sending a SYNC. Any other node will follow A_1 's schedule. The rest of the nodes compete to send a SYNC packet in the next scheduled SYNC time. Nodes that lose contention compete again every cycle until they send a SYNC. Let the i th listen/sleep cycle be called *timeslot* i (t_i), the evolution of the network is as follows:

- At t_1 , node A_1 sends the first SYNC. A cycle counter is set to N_C , which decreases by one every

cycle. When the cycle counter reaches 0 a new SYNC is delivered.

- At t_2 , the rest of the nodes try to send a SYNC packet, but only the contention winner sends it. Let us assume that an arbitrary node A_2 wins contention and so it sets its cycle counter to N_C . We say that node A_2 has *captured* this timeslot.
- At t_3 , the remaining nodes try to transmit their SYNC packet, but, again, only the contention winner actually sends it. Node A_3 wins the contention (captures timeslot) and sets its cycle counter to N_C .
- At t_M , the last present node sends its SYNC packet and sets its cycle counter to N_C .
- From t_{M+1} to t_{N_C} , there are timeslots without SYNC transmissions.
- At t_{N_C} , A_1 sends a SYNC again.
- At t_{N_C+1} , A_2 sends a SYNC again.
- And so on.

Let us notice that S-MAC implicitly defines a TDMA-like arrangement of N_C timeslots, even though, in fact, it is not used, because S-MAC allows all nodes to contend for sending data every listen period. On the contrary, *VTS takes advantage of this property to setup and maintain an adaptive TDMA frame*.

5 VTS description

As stated in Sect. 2, TDMA access schemes are the natural way to keep latency bound and to reduce energy consumption, since there is neither contention nor collisions. VTS constructs a TDMA structure with the exact number of timeslots needed, that is, VTS dynamically adjusts the number of timeslots (N_C) in the TDMA frame to the number of nodes in the cell. In a stationary state, the protocol ensures that each node owns a single timeslot, not shared with any other node. In this situation, a *virtual superframe of timeslots* is created. The word *virtual* means that nodes do not know the superframe arrangement: neither its limits, nor their relative position in the superframe. They just *independently* keep a counter with the superframe length (N_C). VTS synchronization procedure works as

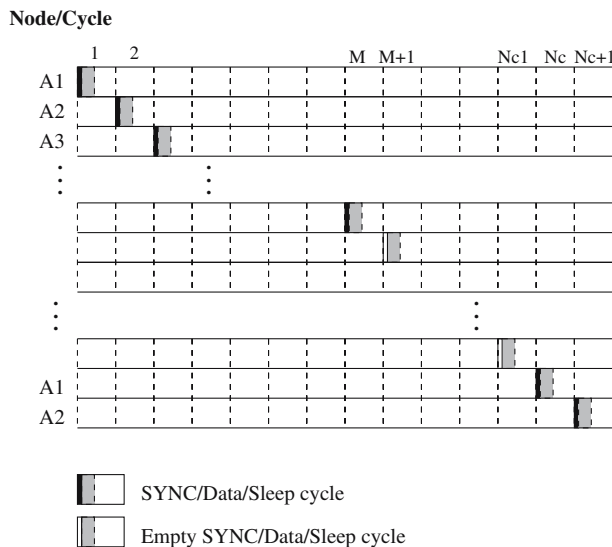


Fig. 2 S-MAC and VTS setup evolution

S-MAC (Sect. 3), but, unlike S-MAC, VTS nodes are only allowed to send data in their captured cycle, i.e., *nodes only send packets, any kind of packet, every N_C cycles after their firstly sent SYNC packet.* It can be seen in Fig. 2 that a superframe of length N_C is virtually established.

Briefly, a VTS node contends every cycle until it captures a timeslot (wins the contention). From then on, the node only sends packets every N_C cycle. After a number of network setup (initialization) cycles, the nodes adjust their superframe length counter to their number of known neighbors. If nodes leave the cell, the superframe length is distributedly reduced. In order to allow new nodes to join the superframe, there is always a short contention period before packet transmission, where new nodes can contend with the owner of the timeslot. Finally, in order to keep a constant latency, nodes change the duration of their sleep period when new nodes join the superframe.

In the following sections a detailed description of VTS operation is provided.

5.1 Network setup

The network setup phase lasts from nodes activation, to the definition of a (still of fixed size) virtual superframe. In VTS we prefer the name Control (CTL) packet to SYNC packet, because the packet is used also for other purposes, which are discussed in next sections. VTS setup stage behaves exactly as the S-MAC synchronization mechanism presented in Sect. 2. That is, a node contends to send a CTL packet every listen cycle until it actually sends it. Then, the node only can send a new

CTL packet after N_C cycles, and it *must* send it as a keep-alive beacon. When all the nodes have sent their first CTL packet, the virtual superframe of N_C timeslots is formed. Let us notice that nodes are not aware of the implicit timeslot allocation. They just know that they are only allowed to transmit packets every N_C cycle. Thus, the slot allocation procedure is simple and fully distributed. Let us also notice that nodes know who their neighbors are when they receive CTL packets. In the example of Fig. 2 it is assumed that all the nodes are initialized simultaneously and that they always capture timeslots consecutively. Let us see what would happen if these assumptions did not hold:

1. *If a node is initialized after the superframe has been formed.* It will wait for a CTL packet to join the listen/sleep schedule. Once it has been received, the new node tries to send its own CTL packet in the next scheduled timeslot. If this timeslot is owned by another node, both of them contend for the timeslot. The contention winner becomes the owner of the timeslot. The loser retries to send its CTL picketing the next timeslot. If a node also owns this timeslot, the contention winner will own the timeslot and the loser will keep trying. Eventually, an empty timeslot will be reached, which is captured by the only remaining node trying to access the medium at that moment. If more than one new node initialized, the only difference would be that the number of contending nodes for a timeslot would be higher. Eventually, every node would be assigned an empty timeslot. We call this process of multiple deallocation and reassignment “allocation loop”.
2. *If more than one node selects the same access instant during contention there is a collision.* In this case, two or more nodes send their CTL packets simultaneously. However, the contenders are not aware of the collision and each one considers that its own CTL has been correctly sent. After N_C cycles (the superframe length) they will try to send their CTL packet again, so there will be a new contention. The contention winner will own the timeslot. The loser will keep trying to send the packet, as discussed in the previous case. In the unlikely event of a new collision this sequence is repeated.
3. *A node wins contention but there is an error in CTL packet reception.* Since the node has won contention, it is the owner of the timeslot, even though other nodes cannot decode the packet due to channel errors. After N_C timeslots it will send its CTL and, unless there is another channel error, the rest of the nodes will get noticed of its presence.

5.2 Adjustment of virtual superframe length

With the previous setup procedure, a virtual superframe of fixed length N_C is created. If N_C were kept fix, protocol performance would be poor.

On one hand, if N_C is less than the maximum number of neighbors, nodes cannot exclusively own a timeslot. In this case, there would always be contention between at least two nodes in every slot. On the other hand, if N_C is greater than the actual number of neighbors, VTS latency and throughput are negatively affected, since they are proportional to superframe length.

To overcome these situations, the number of timeslots should be adapted to the actual number of nodes. Therefore a node adjusts the initial cycle counter (set to N_C) to the real number of neighbors in the cell, that is, *to the number of received CTL packets from distinct nodes up to that moment*. This is done a number of timeslots (let it be N_S , a protocol parameter) after node initialization, From then on, the node dynamically adapts to the possibility of nodes joining and leaving the cell:

1. *New nodes join the cell:* anytime a node receives a CTL packet from an unknown neighbor, the superframe length is updated ($N_C = N_C + 1$).
2. *Nodes leave the cell:* within a superframe the mandatory CTL packet from the timeslot owner must be received. Therefore, CTL serves as a keep live beacon, which allows to signal missing neighbors. However, a single CTL packet missed does not mean that its corresponding node has actually left the network: its CTL packet may have been corrupted or it may have been a collision with a new node joining the cell. Hence, a node is considered missed only after a certain number of inactivity timeslots (let it be N_I). On such event, the frame length and number of known nodes is updated ($N_C = N_C - 1$).

Let us note that in both cases there is a transient period before stability: in the first case the incoming node “steals” a slot, and it causes a new allocation loop. In the second case, nodes are not aware of the position of the lost node in the superframe, so they cannot properly adjust their cycle counter. To overcome it nodes randomly select a value within zero and the number of known neighbors. This solution requires a full reallocation of positions in the superframe. Since these events are supposed to be unlikely, this scheme was preferred because it keeps the protocol extremely simple. It should be remarked here that sensor nodes do not usually move, that is, networks are assumed to

be static, and a node leaves the cell only when it has depleted its battery.


5.3 Data exchange and control packets

At the beginning of each timeslot, all the nodes wake up and listen. The owner of the timeslot performs a carrier sense (choosing a random slot from a fixed contention window), and broadcasts a single and short control packet (CTL); see Fig. 3, which is used as:

- Synchronization and schedule discovery (as S-MAC SYNC packet).
- Keep-alive beacon. It is mandatory for a node to send a CTL packet in its owned slot, since its neighbors must know that the node is active.
- New node discovery. CTL packets include source address. Thus, new nodes are added to the list of known neighbors as CTL packets arrive.
- Channel reservation: RTS information is included in CTL packets. This way, non-addressed nodes may go to sleep just after CTL packet reception.

VTS uses the CSMA/CA mechanism for data delivery. The following types of transmissions exist:

1. Unicast packet transmission. A CTL_{RTS} packet is sent by the owner of the timeslot. Non-addressed nodes change to the sleep state immediately, avoiding overhearing. Destination node replies using a CTS transmission is finished after a Data/ACK sequence and both nodes go to sleep.
2. Broadcast packet transmission. A CTL_{BCAST} packet is sent by the owner of the timeslot. Destination is a broadcast address. All the nodes keep listening. Immediately, sender sends the broadcast packet, that is, without waiting for any CTS reply. After receiving the packet nodes go to sleep. No ACK is sent.
3. No data transmission. CTL_{SYNC} packet is sent. Nodes adjust the clock reference, clear sender inactivity counter and go to sleep.



T: Packet type {RTS,BCAST,SYNC} **Sq#:** Sequence number
DEST: Destination address **SLEEP:** Time to sleep period (ms)
SRC: Source address **CRC:** Cyclic Redundancy Code
S: State (used as schedule identity)

Fig. 3 CTL packet format

Control packet overhearing is reduced this way. A single CTL packet performs synchronization and discovery, reservation and keeps-alive functions.

5.4 Single-hop cluster latency

Let T_C be timeslot duration. In a single-hop cluster in steady state (i.e., all the nodes are the owners of a timeslot) any data transmission between two nodes has a maximum latency (L) given by:

$$L \leq N_C T_C$$

This is the maximum expected latency considering that MAC layer does not enqueue packets. Figure 4 illustrates this expression. Let us assume that the node A_1 generates a packet for any other node in cluster A. In the worst case, this packet arrives just at the end of the activity period (label T_0 in Fig. 4), so a superframe (of length $N_C T_C$) must pass before packet transmission ends (label T_1 in the figure).

5.5 Duty cycle adjustment

In our motivating scenario, packets are due before a given deadline (T_d).

Since the maximum expected latency increase as nodes join the superframe, it is necessary to decrease the duration of a superframe to keep latency below the deadline.

$$L \leq N_C T_C \leq T_d$$

And, since T_C depends on the duty cycle, it is necessary to increase it when a new node joins a superframe.

$$L \leq N_C T_C \leq \frac{N_C T_{listen}}{DC} \leq T_d$$

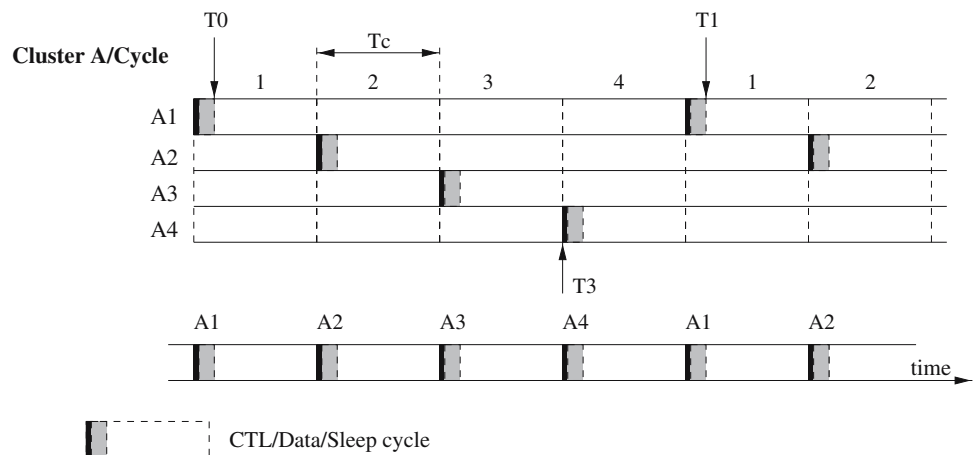
Therefore, we propose to *dynamically adjust the duty cycle*. In order to achieve it, an additional degree of synchronization is necessary. The solution we propose is to use the control center as synchronizer. That is, there is a node, which is directly connected to the control center. This node is obviously the *sink* of the sensor network.

The algorithm we propose is as follows:

1. A node only starts the network setup process (Sect. 7) after receiving a CTL packet from the sink.
2. The CTL packet from the sink includes a field with a suitable DC for a given application. That is, the value of DC that makes latency keep below T_d . In general, the DC given will result in a latency below T_d plus a safety margin. The value of DC depends also on the number of nodes in the superframe. The initial value will be equal to the initial superframe length (before adjustment).
3. The receiving nodes adjust their timeslot duration and set up the superframe arrangement.
4. When the sink adjust the superframe length (Sect. 7), it recomputes the value of DC by taking into account the real number of nodes belonging to the superframe, and informs the rest of node with its CTL packet.
5. Finally, when a new node joins or leaves the superframe, the sink recomputes DC and informs the nodes again. This way, all the nodes (including the sink) adjust the duty cycle at the same time, after receiving the CTL packet from the sink.

With this refinement, latency keeps bounded below a given threshold. The djustment of duty cycle is a trade off between latency and energy consumption, but it acts in both directions: when there is a high number of nodes, latency keeps bounded and consumption is

Fig. 4 VTS TDMA frame



higher, but when there are only a few nodes, energy consumption decreases.

6 Simulation results

In this section, we evaluate VTS through comparative simulations with S-MAC (with and without adaptive listening). S-MAC is chosen as reference because it is a general-purpose protocol, it is well documented and previous results can be found in the open literature [2]. All single-hop experiments are evaluated in a 20-node cell for VTS, S-MAC and S-MAC with adaptive listening.

All figures show the measured parameter versus the packet Inter Arrival time (IA_t) where $IA_t = 0$ means that all the packets are generated at the same time). Unicast packet destination is randomly chosen with equal probability among all the neighbors.

6.1 Simulation configuration

OMNET++² is used as simulation platform. Simulation parameters are selected from reference [2], using the Mica motes³ as underlying hardware. Table 1 shows the main simulation parameters. Additionally, the following options are set for all the simulations:

- A simulation finishes when all the nodes have sent 1,000 data packets (70% unicast and 30% broadcast). Data packets are 100 bytes long.
- Network packet generation starts after a transient time (100 s) plus a random number of cycles uniformly distributed between 0 and 50. Packet generation is then deterministic: packets arrive after a selected IA_t .

6.2 Maximum and average latency

Figure 5a, b show the maximum and the average latency, respectively. A bounded latency is expected for VTS as discussed in Sect. 7. Experiments confirm that in all the cases latency never exceeds the superframe length (as obtained from Eq. 1). VTS effectively adapts the frame length to the actual number of nodes present. In comparison, S-MAC maximum packet latency is clearly not bounded, even with adaptive listening.

Table 1 VTS simulation parameters

Parameter	Value
Radio bandwidth	20 kbps
CTL packet	11 bytes
Listen period	130 ms
Duty cycle	10%
Contention window	31 slots, 1 ms/slot
Consumption in reception state	14.4 mW
Consumption in transmission state	36 mW
Consumption in sleep state	15 μ W
Initial N_C counter	20
Inactivity counter (N_I)	5 superframes
Setup cycles (N_S)	20 cycles
T_C	1.3 s

Under high load conditions (low IA_t), VTS keeps average latency equal to the superframe length (20 $T_C = 26$ s). For low and moderate loads (medium and high IA_t), the average latency depends on the packet generation time, which is uniformly distributed in the timeslot, yielding a latency reduction of one half. In S-MAC, latency depends on the number of nodes contending for the medium. For moderate loads, only a few nodes contend and latency reduces below that of VTS, because VTS nodes must wait for their timeslots to transmit. Adaptive listening is an improvement of S-MAC to reduce latency. Consequently, it outperforms both VTS and S-MAC. Although this is a trade-off between average latency and energy consumption, as it will be shown later.

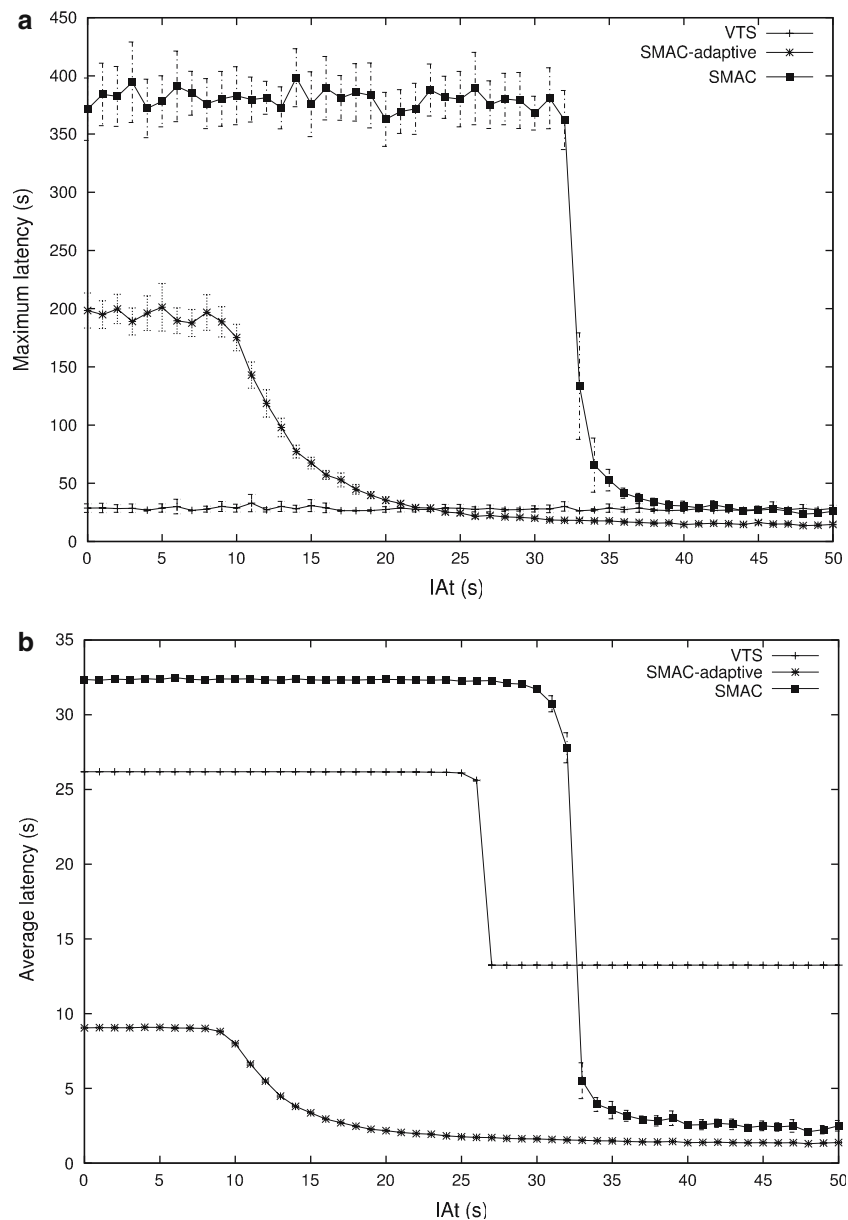
6.3 Power consumption

Figure 6a shows the average network power consumption. S-MAC consumes an 18% more than VTS at high loads (due to the double contention interval per slot and the collisions). S-MAC with adaptive listening consumes two times more power than VTS. This is the cost of reducing latency. In this case, nodes wake up many more times and try to send packets during the scheduled sleep time, consuming the energy of an additional listen interval. As load decreases, the adaptive listening mechanism is not necessary and the behavior is similar to normal S-MAC. In this case, nodes sleep early in VTS, which increases power saving up to 75%. Reduction of VTS power consumption is higher than S-MAC one as load decreases. Under very low loads (high IA_t), VTS will significantly increase the network lifetime compared to S-MAC (the peak at $IA_t = 30$ is due to a steeped decrease in latency 5b, since time to deliver all the packets is reduced and power is measured as total energy consumed divided by total time). In conclusion, in VTS there is a trade-off

² <http://www.omnetpp.org>

³ <http://www.xbow.com>

Fig. 5 a Maximum latency.
b Average latency. Latency for single-hop configuration (average \pm 99% confidence interval)



between latency and energy consumption at low loads, while it guarantees latency at high loads.

6.4 Throughput

Figure 6b shows that at high loads VTS performs slightly better than S-MAC. VTS can handle traffic peaks as properly as a contention protocol. S-MAC with adaptive listening outperforms VTS but at the cost of a higher consumption.

6.5 Transient time

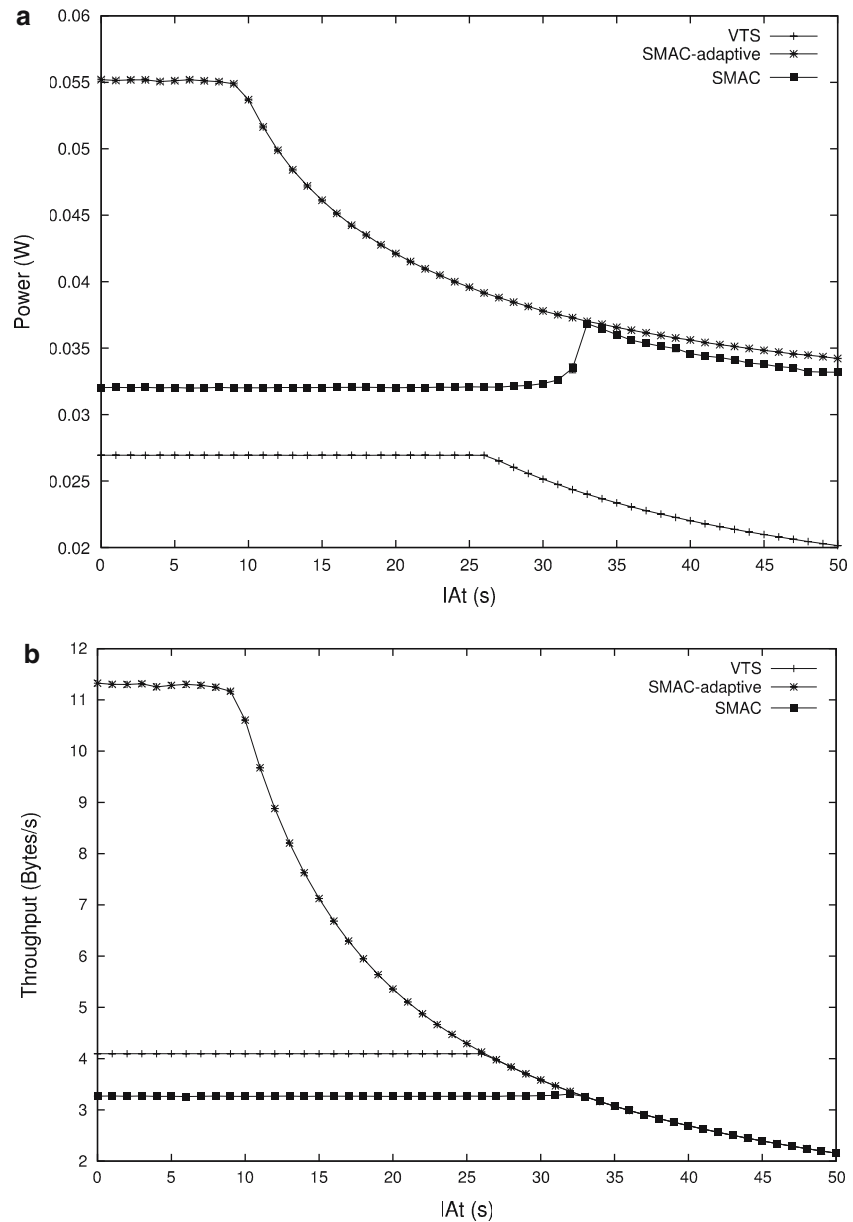
An experiment was conducted to evaluate the time needed to reconfigure the timeslot arrangement when

nodes appear in the network. In this experiment, a 10 nodes network is set and additional nodes join the cell simultaneously (from 1 to 7 nodes), causing allocations loops (see Sect. 7). The results (see Fig. 7a) show that the average transient time until superframe is established again increases linearly approximately with the number of nodes joining simultaneously the area. When nodes leave the cell (again simultaneously, from 1 to 7 nodes), the transient time also increases almost linearly.

6.6 Duty cycle adjustment

The scenario of Sect. 3 is set for the duty cycle adjustment experiments. That is, a single-hop network,

Fig. 6 a Average power consumption. **b** Average throughput per node power consumption and throughput for single-hop configuration (average \pm 99% confidence interval)



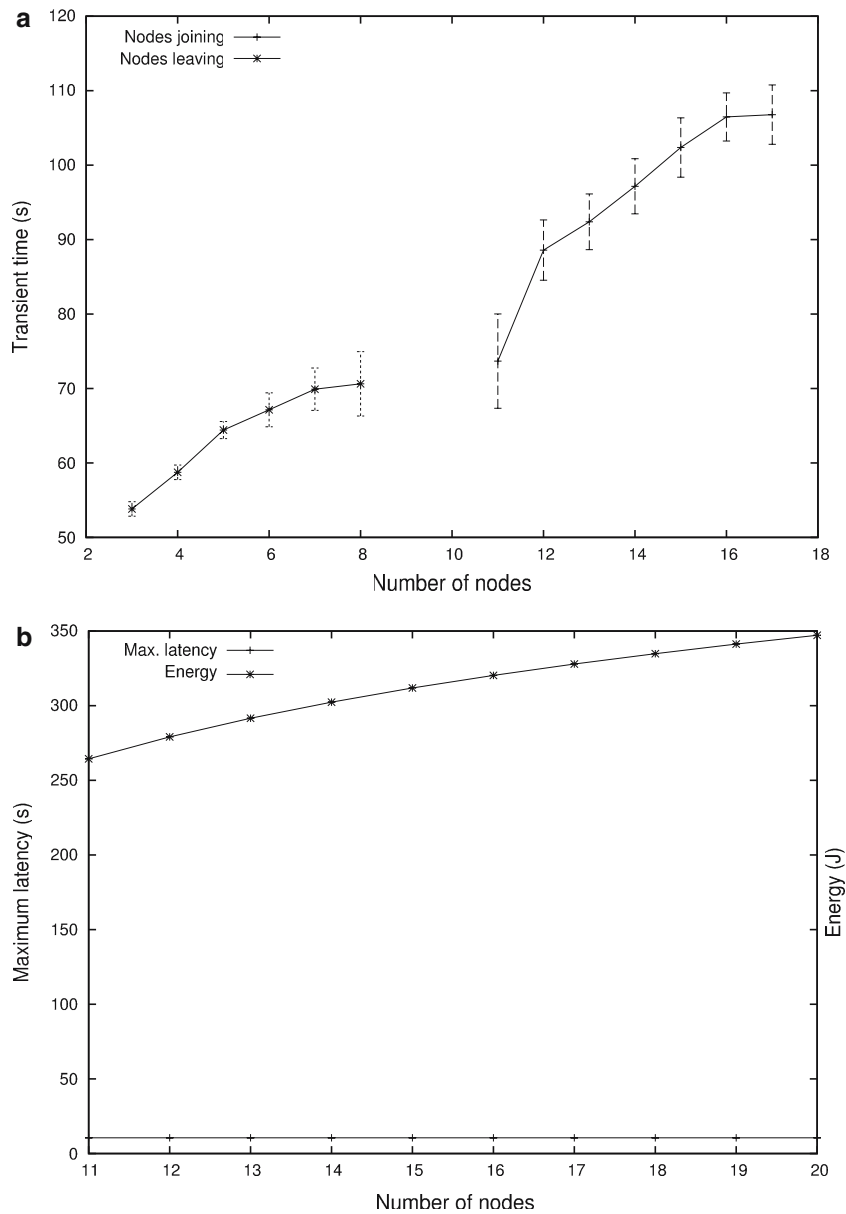
with a sink node and 10 sensor nodes. A number of nodes (from 1 to 10) enter the area at a given time. Nodes adjust superframe length and duty cycle correspondingly.

The deadline is set to $T_d = 15$ s, and the duty cycle (DC) is adjusted to get a margin of 70% of T_d , that is to get a maximum latency of 10.5 s. Nodes send packets continuously after the transient period. As expected, all the experiments exhibit a maximum latency under the upper bound of 10.5 s (see Fig. 7b). The energy (total energy consumed by all nodes, in Joules) increases as the number of nodes on duty cycle increases.

7 Conclusions

In this paper, we proposed VTS, a protocol for WSNs with bounded latency. VTS dynamically creates a superframe of timeslots and adapts its length to the number of actual nodes in a cell for optimum performance. VTS implements a very simple mechanism to adjust and assign timeslots to the nodes (Sects. 7 and 8). VTS further proposes to use a short single control packet to announce any node intentions during its timeslot. Thereby, VTS saves energy by reducing the amount of time a node needs to listen to the channel. Finally, VTS uses dynamic duty cycle adjustment in order to keep

Fig. 7 a Transient time versus final number of nodes.
b Maximum latency and energy consumption for duty cycle adjustment transient time and maximum latency with adaptive duty cycle (average \pm 99% confidence interval)



latency below a given threshold when nodes dynamically join and leave a cell. Simulations reveal that VTS has an excellent power consumption profile, which is crucial in WSN. Under low loads VTS compromises latency and energy consumption, while it guarantees latency at high loads. Additionally, dynamic duty cycle, while trading off latency and energy consumption, keeps latency constant under a given deadline.

Our future work includes refinements to decrease the frame transient time and the development of a generalized mechanism that ensures proper multi-hop operation and an evaluation of scalability.

We also plan to implement and test it with real devices.

Acknowledgments The Economy, Industry and Innovation Council have cofunded this work, with the SOLIDMOVIL project (2I04SU044), supported by Fundacion Seneca, from the Region of Murcia with the ARENA Project (00546/PI/04), with the ARPaq project (TEC2004-05622-C04-02/TCM) by the Spanish Research Council and the CSI-RHET project (TEC2005-08068-C04-01/TCM).

References

1. Akyildiz IF, Kasimoglu IH (1995) Wireless sensor and actor networks: research challenges. Elsevier Ad Hoc Netw 2(4):351–367
2. Ye W, Heidemann J, Estrin D (2004) Medium access control with coordinated, adaptive sleeping for wireless sensor networks. ACM/IEEE Trans Netw 12:493–506

3. Stemm M, Katz RH (1997) Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Trans Commun E80-B*:1125–1131
4. Singh S, Raghavendra C (1998) Power efficient MAC protocol for multihop radio networks *IEEE International symposium on personal, indoor and mobile radio communications*, vol 1, pp 153–157
5. van Dam T, Langendoen K (2003) An adaptive energy-efficient MAC protocol for wireless sensor networks. *SenSys 03-ACM conference on embedded networked sensor systems*, pp 171–180
6. Zhu C, Corson MS (2001) A five-phase reservation protocol (FPRP) for mobile ad hoc networks. *Kluwer Wirel Netw* 7(4):371–384
7. van Hoesel L, Nieberg T, Kip H, Havinga P (2004) Advantages of a TDMA based, energy-efficient, self-organizing MAC protocol for WSNs. *IEEE VTC*
8. van Hoesel L, Havinga P (2004) A lightweight medium access protocol (LMAC) for wireless sensor networks. *INNS 04-International workshop on networked sensing systems*
9. Langendoen K, Halkes G (2005) Energy-efficient medium access control. In: Zurawski R (ed) *Embedded systems handbook*, CRC, Boca Raton