Admission control issues in 802.15.4 sensor clusters

Jelena Mišić,* Shairmina Shafi, and Vojislav B. Mišić†
Department of Computer Science, University of Manitoba
Winnipeg, Manitoba, Canada

PACS numbers: Valid PACS appear here

Contents

I. INTRODUCTION: ADMISSION CONTROL IN WIRELESS SENSOR AND AD HOC NETWORKS

Wireless sensor networks are expected to provide reliable event sensing at the network sink for the longest possible time. These requirements are somewhat conflicting, esp. in the case of networks that use medium access protocols based on the ubiquitous CSMA-CA mechanism, where collisions may substantially affect transmission efficiency and thus impair network lifetime. This is the case with beacon enabled IEEE 802.15.4 network clusters [2], despite the fact that the low rate WPAN (LR-WPAN) standard is considered as an important enabling technology for widespread, standardized deployment of sensor networks [1]. In 802.15.4 clusters, energy is wasted in packet collisions and active measures must be taken to counter this effect. A general solution to the above problem may be found in the mechanism of admission control [12], in which a central authority—most likely, a network coordinator or cluster head—admits or rejects requests by individual nodes to join the network or cluster.

Most common form of admission control is a contention-based mechanism which strives to maintain high utilization in the network under consideration, whilst simultaneously minimizing contention. This goal may be achieved through simple means such as limiting the number of nodes or devices in the cluster, but more complicated schemes that use a combination of traffic characteristics and Quality-of-Service (QoS) indicators, such as delay and throughput, have been proposed as well. The fact that fairness is another desirable goal further complicates the design of suitable admission control techniques.

It is worth noting that admission control techniques have long been studied as a promising approach to solving the generic problem of improving bandwidth utilization whilst maintaining the desired QoS indicators within prescribed limits. The problem was first identified in cellular networks (i.e., mobile telephony), where traffic is predominantly composed of voice calls with well known and standardized traffic characteristics, and issues related to admission control can be treated in isolation from issues related to PHY and other layers. As a result, a number of proposals for admission control in this environment have been described.

More recently, ad hoc and sensor network applications have received much less attention on account of their higher complexity caused by users' mobility and energy considerations. Moreover, admission control in wireless ad hoc and sensor networks is usually treated together with other aspects of the Medium Access Control (MAC) layer functionality, often under the generic label of 'resource allocation' or 'resource management'.

Time-Division Multiple Access (TDMA) techniques, bearing strong resemblance to cellular networks, probably offer the simplest way in which admission control can be enforced. The main challenge, in this case, is to calculate the impact of the traffic generated by a node which requests admission, while actual bandwidth allocation is conducted by the central controller which allocates a certain number of time units to individual nodes. For example, this approach can be used in Bluetooth piconets, where all communications are synchronized to, and controlled by, the piconet master [3]. In other types of networks synchronization may be much more difficult to achieve, which means that TDMA-based approach may not be a viable option for implementing admission control.

Admission control is much more complex in CSMA-based ad hoc networks where all nodes, once admitted, have to contend for medium access. Effective admission control necessitates the availability of a suitable traffic monitoring and shaping (policing) mechanism, which is used in two ways. First, the node which requests admis-

*Electronic address: `jmisic@cs.umanitoba.ca`; This research is partly supported the NSERC Discovery Grant.
†Electronic address: `vmisic@cs.umanitoba.ca`

sion must advertise its traffic parameters, hence the need for monitoring. Second, the node which gets admitted to the network must be able to maintain its traffic rate within the allocated portion of the total available bandwidth; hence the need for traffic shaping. This is a common theme in most, if not all, related proposals: the Admission Control and Dynamic Bandwidth Management scheme [11], the Contention-Aware Admission Control (CACP) [12], both of which are described in the context of IEEE 802.11 networks, as well as in the schemes for peer to peer networks [9, 10].

Sensor networks pose a slightly different set of challenges. In most cases, sensor nodes are not mobile; but on the other hand, sensor nodes are often required to operate on battery power for prolonged periods of time. What this means is, that the network topology is again nonstationary – not because of mobility, but because of the fact that sensor nodes will eventually exhaust their power source and cease functioning. Energy efficiency becomes one of the main factors that affect the design of the network and dictate the choice of communication technology and operating regime, if not *the* main such factor. It is worth noting that energy efficiency is important both for individual nodes and for the network as a whole. In this case, bandwidth utilization may simply be regulated through a sleeping mechanism that is applied to the large number of nodes present in the cluster after being admitted by the cluster coordinator. In this case, each node is free to undertake the transmission of its packets (provided there are some) whenever it wakes up from sleep; this sleep may be coordinated in a centralized manner (i.e., through explicit commands sent by the cluster coordinator) or in a distributed fashion, where each node goes to sleep of its own volition, based on some aggregate information broadcast by the cluster coordinator. An example of the latter solution was originally proposed in [6].

It is also possible to combine the two approaches described above into a single scheme. In this cases, nodes go to sleep, but are not automatically allowed to transmit when they wake up. Instead, they first check with the cluster coordinator whether they are permitted to transmit; if not, they go to sleep again. The cluster coordinator, on the other hand, will only admit (and allow transmission by) the number of nodes that allows the network to perform the sensing task without excessive collisions. A suitable admission scheme can be developed based on the activity management model from [6]. However, this model is rather complex and requires numerical solution of a large number of non-linear equations. While this solution is acceptable for off-line performance analysis, it is virtually intractable when admission decisions have to be made in real time.

In order to find a viable solution for the sensor network environment, we have developed a simplified cluster model that can be repeatedly solved using limited computational resources available to the cluster coordinator. The model uses number of nodes, average packet arrival rate, and packet length as independent variables; from these, it calculates packet service time, which is then used as the major admission condition. This simplified cluster model is developed using the results of simulation experiments conducted on clusters with a variable number of nodes, assuming that the packet arrival rates are not symmetric but follow a uniform distribution instead, with small to moderate deviations from the average value.

The remaining part of the chapter is organized as follows. In Section II we briefly discuss the operation of 802.15.4 MAC layer. In Section III, we discuss the performance of the cluster and identify the main performance indicators that can be used for admission control. The admission control scheme is derived in Section IV; it is computationally lightweight and thus suitable to be applied in resource-constrained sensor cluster. The performance of the admission control scheme is evaluated in Section V. Section VI concludes the chapter.

II. OPERATION OF THE 802.15.4 MAC LAYER

In beacon-enabled 802.15.4 networks, the channel time is divided into superframes delimited by the transmission of network beacons by the network (or PAN) coordinator. The superframe consists of an active portion, during which the coordinator interacts with the nodes, and an optional inactive portion, during which all devices may enter a low power mode to reduce power consumption. The active portion of each superframe is divided into equally sized slots. Each slot consists of $3 \cdot 2^{SO}$ backoff periods which gives the shortest active superframe duration *aBaseSuperframeDuration* of 48 backoff periods. When the cluster operates in the ISM band at 2.4GHz, the duration of the backoff period is 10 bytes, the duration of one slot is 30 bytes, and the maximum data rate is 250kbps.

The active portion of the superframe is divided into a Contention Access Period (CAP) and an optional Contention-Free Period (CFP), as shown in Fig. 1. 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, both of which must successfully pass for the transmission to begin.

A successful transmission is optionally acknowledged within a predefined time period. The absence of acknowledgment (ACK) indicates that the transmission has failed (which may be due to a collision with another transmission, or packet blocking because of buffer overflow) and

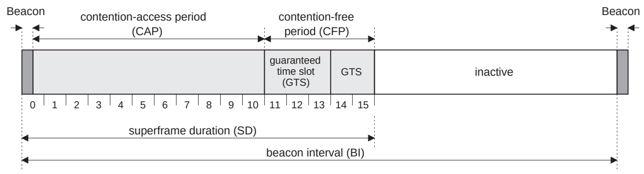

FIG. 1: The composition of the superframe in a 802.15.4 cluster (adapted from [2]).

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 [5].)

In this work, we assume that the 802.15.4 cluster operates in beacon enabled, slotted CSMA-CA mode under the control of a cluster (PAN) coordinator. We assume that all transmissions are sent from the sensor nodes to the cluster coordinator, i.e., there is uplink traffic only, which is common in simple sensor networks. We also assume that all transmissions are acknowledged.

III. PERFORMANCE UNDER ASYMMETRIC PACKET ARRIVAL RATES

In order to develop the admission control algorithm, we have simulated an IEEE 802.15.4 cluster using the object-oriented Petri net simulation engine Artifex by RSoft design, Inc. [8]. The goal of this analysis is to investigate the impact of different parameter values on cluster performance, in particular the effect of the asymmetry of the load originating from individual sensors, and to identify parameters that might be suitable candidates for the implementation of an admission control scheme.

For these simulations, we assumed that the network operates in the ISM band at 2.4 GHz, with the maximum data rate of 250kbps. The cluster had the coordinator and $n = 5, 15$, and 25 ordinary sensor nodes, respectively. Packet arrival rates at each of the ordinary nodes were uniformly distributed in the range of $(1-\delta)\lambda$ to $(1+\delta)\lambda$, where the variation span δ took values from 0.1 to 0.8, and λ was the arrival rate averaged over all nodes. (For example, the load variation of $\delta = 0.5$ corresponds to uniformly distributed arrival rates in the range between 0.5λ, and 1.5λ.) We assume the packet size is fixed at 90 bytes, which includes all PHY and MAC layer headers. Ordinary nodes had buffers that can hold $L = 3$ packets. All other parameters were set to default values prescribed in the 802.15.4 standard [2].

In this setup, we have measured the throughput and the service time. The corresponding values are shown in top and middle rows of Fig. 2 as functions of the average packet arrival rate for the entire cluster λ and the load variation span δ. Note that the mean packet service time does include the effect of packet retransmissions: namely, if the sending node does not receive an acknowledgment from the cluster coordinator, it will assume that the packet is lost due to a collision. The sending node will then repeat the transmission until the proper acknowledgment is received.

From these diagrams, two important observations can be made. First, cluster performance is predominantly dependent on the total traffic load in the cluster, but virtually independent of the load asymmetry, i.e., the load variation among individual cluster nodes.

This observation holds for both throughput and mean packet service time, under a rather wide range of individual node traffic. Note that a variation of $\delta = 0.8$ means that the values of traffic load for individual nodes is uniformly distributed in the range from $0.2\lambda_i$ to $1.8\lambda_i$, where λ_i is the mean packet arrival rate per node.

The second observation is that the cluster effectively goes into saturation beyond a certain traffic load. Saturation is caused by the CSMA-CA algorithm which allows collision of packets sent by different nodes. As long as the number of nodes is small ($n = 5, 15$), the overall range of traffic loads in the diagrams is small and the cluster does not saturate.

However, in the cluster with $n = 25$ nodes, the probability of collisions is much higher. When the traffic load exceeds a certain level, the collisions become prevalent. When this happens, the majority of the bandwidth is taken up by retransmissions and the overall cluster efficiency decreases. At the same time, the sensor nodes are forced to reject or drop newly arrived packets, since the input buffers are occupied by the packets awaiting retransmission.

The cumulative effect of those phenomena is that the throughput experiences a sharp decrease, Fig. 2(c), while the packet service time shows a large increase due to packet retransmission, Fig. 2(f).

Note that collisions, effectively, waste bandwidth; the presence of a large number of collisions means that the nodes are using up their energy resources without actually managing to send useful data. Therefore, the number of collisions must be kept as low as possible in order to conserve energy. A detailed analysis of energy consumption is beyond the scope of this chapter; a suitable activity management algorithm for 802.15.4 networks can be found in [7].

In order to verify that the aforementioned effects are indeed caused by overwhelming number of collisions, we have also measured the probability of successful transmission γ, which is shown in the bottom row of Fig. 2. As can be seen, the success probability is over 0.9 in the cluster with five nodes (note that seemingly large variations are due to the reduced vertical range for γ), and drops to about 0.55 in the cluster with 15 nodes. However, when saturation occurs in the cluster with 25 nodes, the success probability virtually drops to zero, as shown

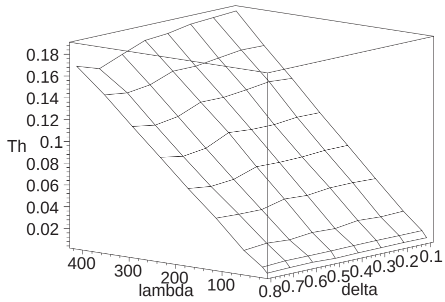

(a)Throughput with $n = 5$ nodes.

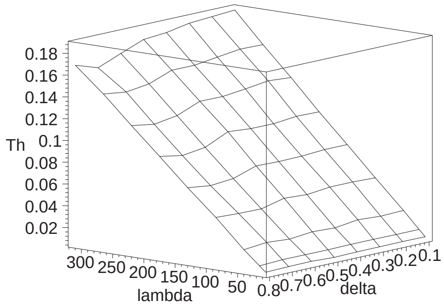

(b)Throughput with $n = 15$ nodes.

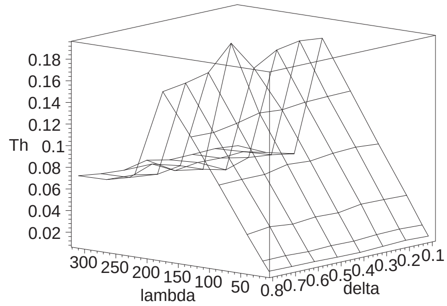

(c)Throughput with $n = 25$ nodes.

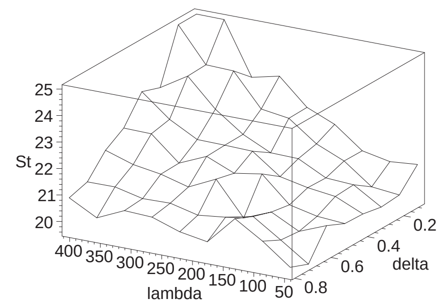

(d)Mean packet service time with $n = 5$ nodes.

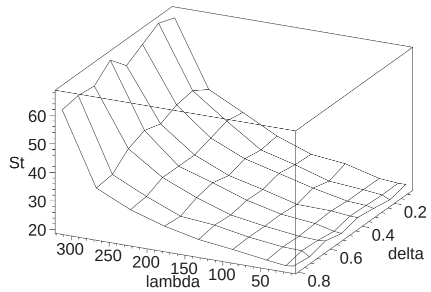

(e)Mean packet service time with $n = 15$ nodes.

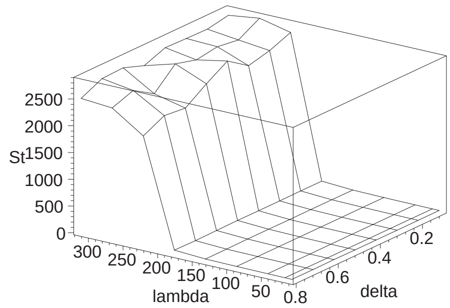

(f)Mean packet service time with $n = 25$ nodes.

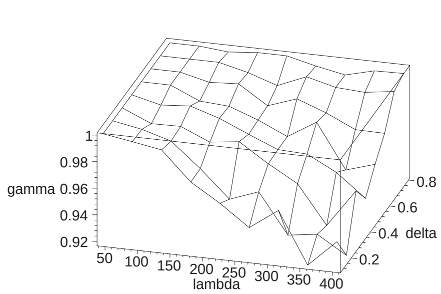

(g)Success probability with $n = 5$ nodes.

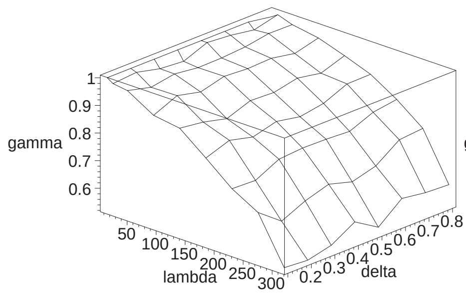

(h)Success probability with $n = 15$ nodes.

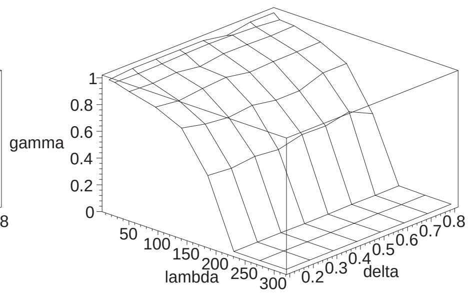

(i)Success probability with $n = 25$ nodes.

FIG. 2: Cluster performance under asymmetric traffic: cluster throughput (top row), mean packet service time (middle row), and probability of successful transmission (bottom row).

in Fig. 2(i), which confirms our analysis.

IV. CALCULATING THE ADMISSION CONDITION

The analysis presented above can be summarized as follows. First, the main factor affecting performance is the *total cluster load*, rather than the degree of asymmetry among the nodes in the cluster. Consequently, the analysis results, and admission schemes derived thereof, for the case of symmetric cluster traffic can safely be utilized for clusters with asymmetric traffic.

Interestingly enough, similar conclusion were made in the context of Bluetooth piconets [3]; however, the underlying mechanism that leads to this independence is quite different since Bluetooth uses TDMA.

Second, the main objective of the admission control in an 802.15.4 cluster is to prevent the onset of the saturation condition.

In view of this, the actual admission algorithm is quite simple. The cluster coordinator monitors the operation of the cluster, in particular the traffic load. When the new node requests admission, the cluster coordinator calculates whether the added traffic will cause the cluster to move into saturation. If the answer is negative, the new node is admitted to the cluster.

The next step is to find a suitable indicator variable, and to devise the actual algorithm for the necessary calculation. In an earlier work, the performance of the network under symmetric load was considered [4], and it was found that the forthcoming saturation regime is reliably predicted by any of the following conditions:

- packet service times which are longer than a certain value (approx. 50 backoff periods);
- transmission success probability lower than about 70%;
- finally, access probability at the node larger than a certain value (around 0.005).

While any of these descriptors can be used for the purpose of admission control, the packet service time appears to be the most accurate indicator of the overall cluster load. Therefore, we will focus on calculating or approximating the packet service time as the main indicator of the cluster operating regime. A promising starting point is the analytical model for an 802.15.4 cluster with uplink traffic described in [6]. But before we proceed, let us first introduce the necessary notation.

Let the PGF of the data packet length be $G_p(z) = \sum_{k=2}^{12} p_k z^k$, where p_k denotes the probability of the packet size being equal to k backoff periods or $10 \cdot k$ bytes. Then, the mean data packet size is $G'_p(1)$ backoff periods.

Let the PGF of the time interval between packet transmission and subsequent acknowledgement (ACK) be $t_{ack}(z) = z^2$; actually its value is between *aTurnaroundTime* and *aTurnaroundTime + aUnitBackoffPeriod* [2], but we round the exponent to the next higher integer for simplicity. Also, let $G_a(z) = z$ stand for the PGF of the ACK packet duration. We note that the timing prescribed by the standard precludes the possibility that an ACK will collide with the transmission of a packet from another device.

Then, the PGF for the total transmission time of the data packet will be denoted with $D_d(z) = z^2 G_p(z) t_{ack}(z) G_a(z)$, while its mean value is $\overline{D_d} = 2 + G'_p(1) + t'_{ack}(1) + G'_a(z)$.

Let us denote the probabilities that the medium is idle on first and second CCA with α and β, respectively, and the probability that the transmission is successful with γ. Note that the first CCA may fail because of a packet transmission in progress (originating from another device), and this particular backoff period may be at any position with respect to that packet. The second CCA, however, fails only if some other device has just started its transmission – i.e., this must be the *first* backoff period of that packet. Since the corresponding probabilities differ, we need two different variables.

The simplified probability that the medium is busy at the first CCA is

$$1 - \alpha = \frac{(1 - (1 - \tau)^{n-1}(G'_p(1) + G'_a(1)))}{\overline{D_d}} \tag{1}$$

where τ denotes the access probability – i.e., the probability that the packet will be transmitted. (γ denotes the probability that the transmission will be successful.)

If the cluster operates in non-saturated regime, we are able to ignore the potential congestion effects at the beginning of superframe due to delayed transmissions from previous superframe, as described in Section II. In this case, the access probability τ is uniform throughout the superframe and the above expression holds.

The probability that the medium is idle on the second CCA for a given node is, in fact, equal to the probability that neither one of the remaining $n-1$ nodes has started a transmission in that backoff period. Then the probability of success of second CCA becomes:

$$\beta = (1 - \tau)^{n-1} \tag{2}$$

.

By the same token, the probability of success of the whole transmission attempt becomes:

$$\gamma = \beta^{\overline{D_d}} \tag{3}$$

The PGF of the time needed to conduct one transmission attempt (potentially unsuccessful) is

$$A(z) = \frac{\sum_{i=0}^{m} \prod_{j=0}^{i} (B_j(z) R_{ud})\, z^{2(i+1)}\, (\alpha\beta T_d(z))}{\sum_{i=0}^{m} \prod_{j=0}^{i} R_{ud}\alpha\beta} \tag{4}$$

where

- $T_d(z) = G_p(z) t_{ack}(z) G_a(z)$ represents the transmission and ACK time without the backoff procedure;
- $R_{ud} = 1 - \alpha\beta$ is the probability that one CCA will not be successful;
- $B_j(z) = \sum_{k=0}^{W_j - 1} \frac{1}{W_j} z^k = \frac{z^{W_i} - 1}{W_j(z-1)}$ represents the PGF for the duration of j-th transmission attempt; and
- $0 \,..\, W_i - 1$ represents the backoff counter value in i-th transmission attempt. If the battery saving mode is not turned on then $W_0 = 7$, $W_1 = 15$, $W_2 = W_3 = W_4 = 31$.

If we assume that transmission will eventually succeed in five attempts, which is reasonable under light to moderate load, the probability distribution of the packet service time follows the geometric distribution and its PGF is

$$\begin{aligned} T(z) &= \sum_{k=0}^{\infty} (A(z)(1-\gamma))^k A(z)\gamma \\ &= \frac{\gamma A(z)}{1 - A(z) + \gamma A(z)} \end{aligned} \tag{5}$$

In this case, mean packet service time can simply be written as

$$\overline{T} = T'(1) = \frac{A'(1)}{\gamma} \tag{6}$$

The period between two consecutive packet service times also depends on the state of the input buffer of the device in question. If the buffer is not empty after packet departure (when positive acknowledgement is received), then the next service time will start immediately; if the buffer is found to be empty after a successful packet transmission, the next service period starts when the next packet arrives. If the probability that the buffer is empty after a packet transmission is denoted with π_0, the mean time period between two packet departures is $\overline{P} = \overline{T} + \frac{\pi_0}{\lambda_i}$, and the period between two transmission attempts by the same device is

$$p = \overline{A} + \frac{\pi_0 \gamma}{\lambda_i} \tag{7}$$

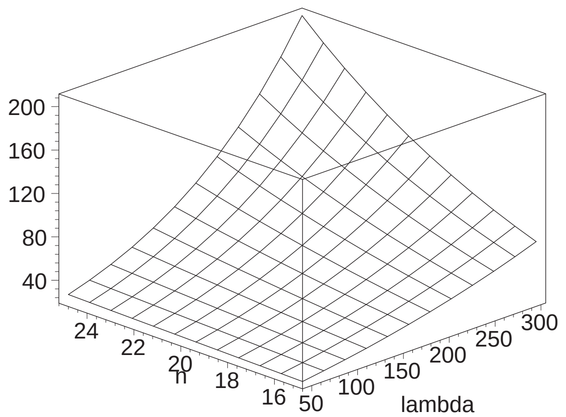

(a)Packet service time.

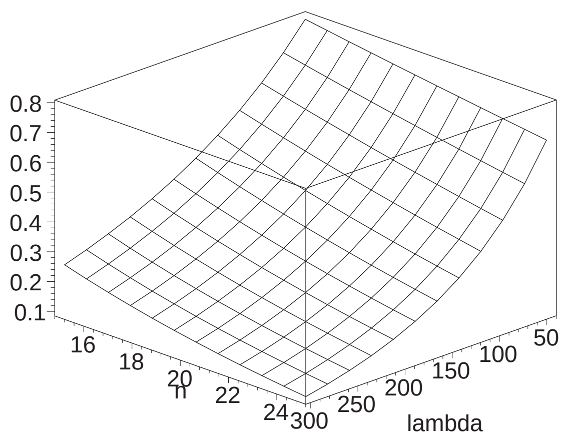

(b)Success probability.

FIG. 3: Approximation of performance indicators.

In a previous work [6], the probability distribution of the length of device queue were derived using the M/G/1/K queueing model coupled with the Markov chain model of the 802.15.4 MAC algorithm. Due to the large computational complexity of the model, the probability π_0 could be obtained only through elaborate numerical calculations. While this approach is appropriate for offline performance evaluation, it is definitely unsuitable for real-time admission control. The limited computational resources of sensor nodes necessitates a much simpler solution.

In order to arrive at a suitable algorithm, we have made two simplifying approximations. First, the expression for the mean time for a transmission attempt was simplified by retaining just the first two members of the Taylor series expansion around point $\tau = 0$. The accuracy of the approximate solution was found to be quite sufficient. Second, we have assumed that the product $\pi_0 \gamma$ is approximately constant and equal to 0.2.

Then, the approximate period between transmission attempts is

$$\widetilde{p} = \widetilde{A} + \frac{0.2}{\lambda_i} \tag{8}$$

In this case the access probability for node becomes:

$$\widetilde{\tau} = \frac{1}{\widetilde{p}}. \tag{9}$$

and we can solve equation 9 algebrically for $\widetilde{\tau}$ using n, λ_i and $\overline{D_d}$ as parameters. Thus, we can obtain approximate values for $\widetilde{A}$, γ, and $\widetilde{St}$ as functions of n, λ_i and $\overline{D_d}$. The latter two are shown in Fig. 3, for $\overline{D_d} = 9$.

V. PERFORMANCE OF ADMISSION CONTROL

In order to verify the performance of the admission control, we have augmented the simulator with the calculation of approximate expressions $\widetilde{\tau}$, $\widetilde{\gamma}$, and $\widetilde{St}$. The cluster was set with variable number of nodes, starting with n=15 nodes and adding one node at a time. The packet arrival rate of each node is featured both in symmetric and in asymmetric approach. In symmetric approach, we assumed the packet arrival rate $\lambda = 2$ packets per second for each and every node. In asymmetric case, each of the nodes featured a variable packet arrival rate λ_i for i=1..n where the load variation span δ took value of 0.5 and the average packet arrival rate $\lambda = 2$ packets per second. In both cases, the packet size was set to nine backoff periods.

After every 120 seconds, a new node applied for admission. The cluster coordinator then performed the approximate calculations outlined above to obtain the approximate packet service time. The node is admitted based on a desirable admission condition. The admission condition is set as the smallest duration of an 802.15.4 superframe – 48 backoff periods in terms of packet service

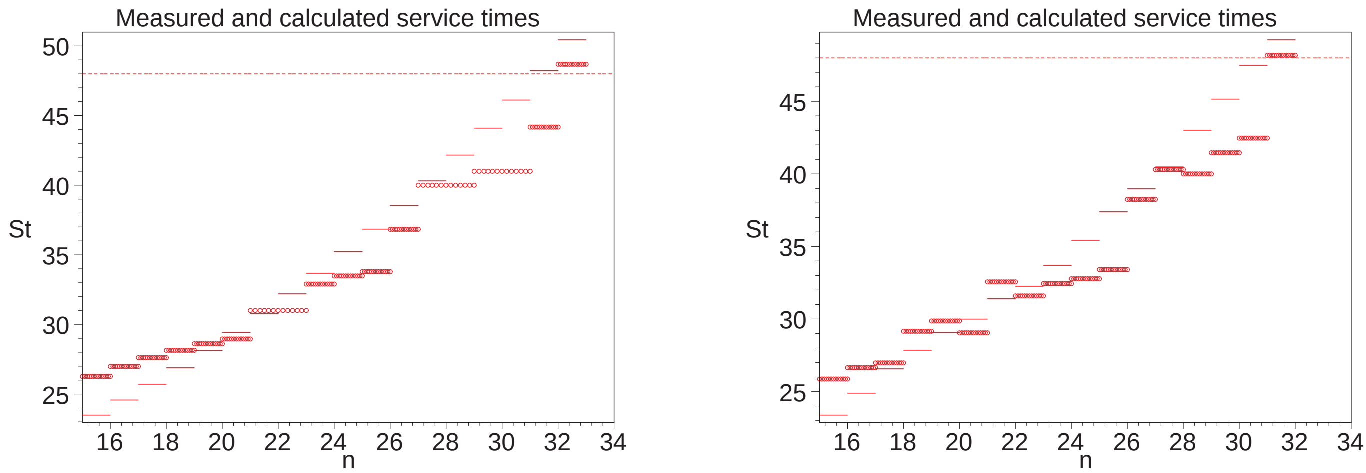

(a)Symmetric traffic. (b)Asymmetric traffic.

FIG. 4: Measured and calculated service times of individual nodes.

time. The mean packet service time is below 48 backoff periods means most of the packets will be processed within a single superframe, which is certainly desirable from the standpoint of energy conservation. The algorithm that the coordinator executes to admit a new node is as follows:

Data: number of existing nodes n ;
packet arrival rates λ_i ;
new node packet arrival rate λ'
Result: admission decision

calculate the new number of nodes $n' = n + 1$;
if *traffic load is asymmetric* **then**
  calculate the new average arrival rate
  $\lambda_{avg} = (\lambda_1 + \lambda_2 + ...\lambda_n + \lambda')/n'$;
else
  calculate $\lambda_{avg} = \lambda$;
end
obtain approximate packet service time $\widetilde{St}$ using n' and λ_{avg} ;
if $\widetilde{St} \leq 48$ **then**
  allow the new device to join ;
  update $n = n'$, $\lambda_{n+1} = \lambda'$;
else
  reject admission ;
  retain existing value of n ;
end

Algorithm 1: Admission policy.

The calculated and measured values for packet service time under symmetric and asymmetric traffic are shown in Fig. 4(a) and Fig. 4(b), respectively. Short horizontal lines show approximate values of service times as calculated by the algorithm, while circle points show measured values (averaged over a five second time window) after admission. The superimposed line at $T = 48$ is set as a desirable admission condition.

We observe very close match between calculated and simulated values for both symmetric and asymmetric traffic. We also observe that the admission condition is same in both cases, allowing 31 nodess to get admitted with an average packet arrival rate $\lambda = 2$ packets per second. Interestingly, measured values of packet service time do not increase all the time with increasing number of nodes in asymmetric case – sometimes simulation results with higher number of nodes fall slightly behind the simulation results with lower number of nodes. The reason for such behaviour will be clear if we observe the diagram shown in Fig. 6, which shows the distribution of packet arrival rates among the nodes in the case of asymmetric traffic. Although we assume that the distribution of load variation is uniform in the range of 50% around the mean, the number of samples (which is between 15 to 31) is not sufficiently high enough to generate a steady average arrival rate λ. As a result, when a new node is admitted, the new average arrival rate may vary within a few percent of the average arrival rate λ. Thus, after admitting a new node, the new measured service time is affected by the current number of nodes as well as the new average arrival rate. For example, the avearge service time generated by a cluster of size 20 nodes with $\lambda = 2.03$ packets per second may be slightly greater than the service time generated by 21 nodes with $\lambda = 1.98$ packets per second.

VI. CONCLUSION

We have developed simplified analytical expressions for the packet service time that can be used for admission control at the 802.15.4 cluster coordinator. Simulation results for symmetric traffic show that simplified admission condition is slightly more conservative than the accurate value.

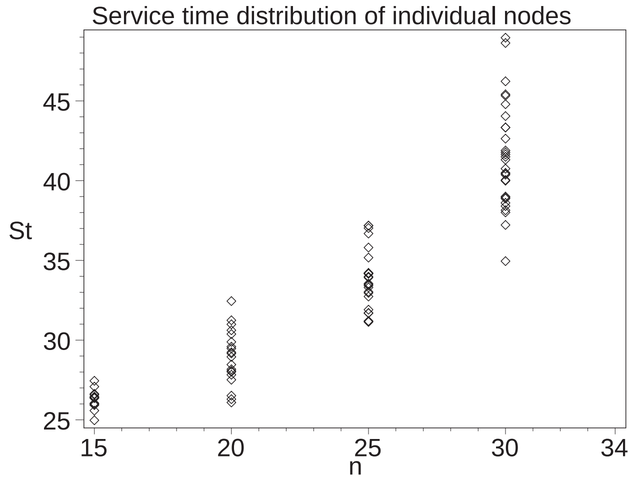

(a)Symmetric traffic.

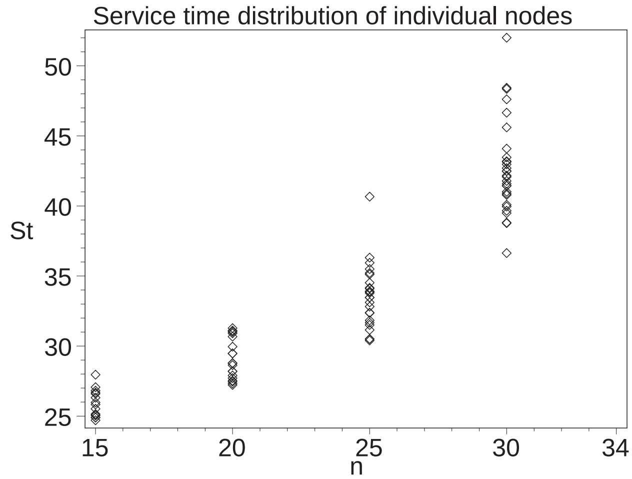

(b)Asymmetric traffic.

FIG. 5: Packet service time distribution of individual nodes.

[1] E. H. Callaway, Jr. *Wireless Sensor Networks, Architecture and Protocols.* Auerbach Publications, Boca Raton, FL, 2004.

[2] 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.

[3] 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.

[4] 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*, pages 416–425, San Jose, CA, Oct. 2004.

[5] J. Mišić, S. Shafi, and V. B. Mišić. Avoiding the bottlenecks in the MAC layer in 802.15.4 low rate WPANs. In *Proc. HWISE2005 (ICPADS 2005 Workshops)*, pages 363–367, Fukuoka, Japan, July 2005.

[6] J. Mišić, S. Shafi, and V. B. Mišić. Maintaining reliability through activity management in 802.15.4 sensor networks. In *Proc. Qshine 2005*, Orlando, FL, Aug. 2005.

[7] V. B. Mišić, J. Fung, and J. Mišić. Cross-layer activity management in a 802.15.4 sensor network. *IEEE Communications Magazine*, 44(1):131–136, Jan. 2006.

[8] RSoft Design, Inc. *Artifex v.4.4.2.* San Jose, CA, 2003.

[9] N. Saxena, G. Tsudik, and J. H. Yi. Admission control in Peer-to-Peer: design and performance evaluation. In *Proc. SASN'03*, pages 104–113, Fairfax, VA, 2003.

[10] N. Saxena, G. Tsudik, and J. H. Yi. Access control in ad hoc groups. In *International Workshop on Hot Topics in Peer-to-Peer Systems*, pages 2–7, Volendam, The Nederlands, Oct 2004.

[11] S. H. Shah, K. Chen, and K. Nahrstedt. Dynamic Bandwidth Management in Single-Hop Ad Hoc Wireless Networks. *Mobile Networks and Applications*, 10:199–217, 2005.

[12] Y. Yang and R. Kravets. Contention-aware admission control for ad hoc networks. *IEEE Transactions on Mobile Computing*, 4(4):363–377, July-August 2005.

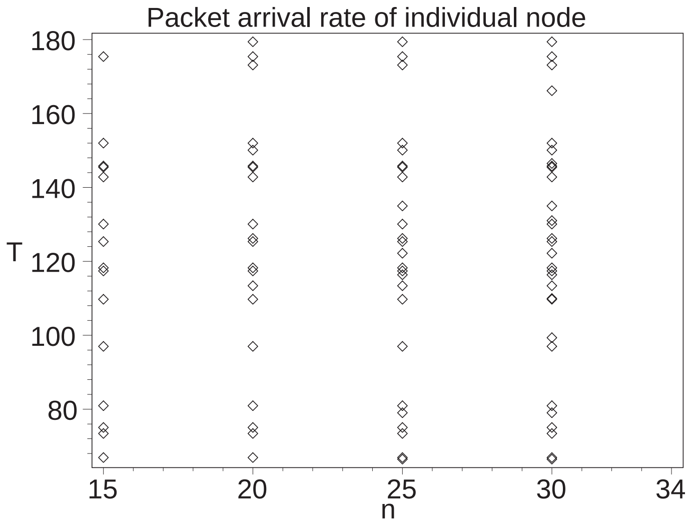

FIG. 6: Packet arrival rate of individual nodes under asymmetric traffic.