LAMA: Location-Aware Medium Access for Wireless Sensor Networks

Kiran Yedavalli and Bhaskar Krishnamachari

Department of Electrical Engineering - Systems
University of Southern California, Los Angeles, CA 90089, USA

Abstract. We present a novel location-aware medium access protocol for wireless sensor networks. In this protocol, the contending nodes make use of their location information to reduce collisions and improve the overall performance. We study the application of this protocol to the problem of medium access in one-hop data-gathering applications. We evaluate it in terms of delay and energy consumption and compare it with location-unaware medium access protocols using simulations. Results show that the location-aware medium access protocol can take advantage of the location distribution of nodes to provide significantly lower delay and energy consumption compared to location-unaware protocols.

1 Introduction

Location awareness of sensor nodes is increasingly common in many wireless sensor network applications. For example, protocols such as GPSR [1] have used it to provide efficient routing. In this paper, we propose a novel medium access protocol called the *location-aware medium access* (LAMA) protocol, that makes use of the location awareness of sensor nodes to provide efficient wireless medium access.

The main idea in the LAMA protocol is the separation of collision domains of nodes using spatial partitioning. A tree-based space partitioning procedure is used to adaptively partition the space until each node can transmit its packet successfully, without collisions. The key point here is that spatial partitioning allows us to leverage the location distribution of sensor nodes to provide efficient medium access.

In this paper, we study the application of the LAMA protocol to the one-hop medium access problem that occurs frequently in many applications such as neighbor discovery, data gathering, etc. An important aspect of this problem in such applications is that, each sensor node has a *single* packet to transmit to the sink. For example, in neighbor discovery, which is an essential part of many routing protocols, the sink is the node that discovers its neighboring nodes. A single packet is sufficient for each neighboring node to transmit its ID to the sink.

Another important application in which the above *single-packet medium access* problem arises is sensor-data-gathering. Representative examples of this application include, (i) sensor data collection by the cluster-head in a hierarchical structural health monitoring infrastructure, (ii) mobile inventory of goods on the shelves of ware-houses, (iii) air-borne data gathering from sensor nodes spread over large geographical areas such as forests and farms. The two main metrics of interest in these applications are the delay for the sink to receive a pre-determined number of packets and the resultant energy expenditure of the sensor nodes.

We evaluate the performance of the LAMA protocol in terms of the above two performance metrics and compare it with three location-unaware medium access protocols – HT-split, optimal p-persistent slotted CSMA, and the IEEE 802.15.4 standard. We show through simulations that the LAMA protocol can take advantage of the location distribution of sensor nodes to provide significantly lower delays and energy consumptions compared to location-unaware medium access protocols.

The rest of the paper is organized as follows. In the next section we describe the assumptions and metrics associated with the single-packet medium access problem and in Section 3, we present the location-aware medium access protocol in detail and discuss its implementation aspects. In Section 4, we present the protocol performance evaluation results and discuss its scope in Section 5. We present the related work in Section 6. Finally, we conclude and briefly discuss the future directions of our work in Section 7.

2 Problem Description

In this section, we describe the assumptions and performance metrics associated with the single-packet medium access problem.

The one-hop sensor network has n contending sensor nodes, not including the sink (which does not contend for the channel), each with a single data packet to be transmitted. The locations of all the nodes, including that of the sink, are known, but their deployment density is unknown. Time is divided into slots and each node transmits its packet only at the beginning of a time slot. If more than one node transmits in the same time slot, it results in a collision. Otherwise, if a single node transmits in a time slot, it results in successful transmission of the packet. On successful transmission, the node is no longer in contention of the medium. The sink uses explicit acknowledgement (ACK) and negative-acknowledgement (NACK) packets to indicate successful transmission and collision, respectively, to the sensor nodes. The sink broadcasts the ACK/NACK packets as soon as the data packet transmission is completed by the sensor nodes. We assume that the time slot length includes the transmission of both the data packet and the ACK/NACK packet.

In order to study the intrinsic performance advantages of the LAMA protocol, we isolate the random errors due to noise and wireless channel non-idealities such as multi-path fading and shadowing[1]. We assume that the transmission powers of the sensor nodes and the sink are such that all links in their radio ranges are symmetric and error free.

[1] These non-idealities affect all the medium access protocols in consideration in equal measure. Therefore, it is very reasonable to isolate them to evaluate the protocols' performance.

We consider the following two performance metrics for the single-packet medium access problem:

1. **Delay for the first k packets** ($D(k)$): The number of time slots required for the sink to successfully receive the first k packets from the sensor nodes.
2. **Energy consumption per node for the first k packets** ($E(k)$): The average number of transmissions per sensor node for the sink to successfully receive the first k packets.

3 Location-Aware Medium Access (LAMA)

Now, we describe the LAMA protocol applied to the single-packet medium access problem and illustrate its working through examples.

The main idea in the LAMA protocol is a tree-based splitting of space that adaptively reduces the collision domain of sensor nodes until each node is able to transmit its packet successfully. The protocol starts out by splitting the space in the radio range of the sink into m equal partitions. Each partition is a separate collision domain. At each step, only nodes belonging to the current partition are allowed to transmit their packets. When a partition has more than one nodes, their transmission leads to collision. In the event of a collision, the current partition is further split into m equal partitions. This continues until the current partition has at-most a single node in it. The protocol moves onto the next partition after all nodes in the current partition have successfully transmitted their packets. This process of space splitting builds a tree with m branches at each split, where, each branch is a separate collision domain. The leaves of the tree are collision domains with at-most a single sensor node in them and therefore, successful transmissions can take place only from the leaves of the tree.

We illustrate the LAMA protocol through an example shown in Figure 1, in which the space is a square whose half-diagonal is equal to the radio range of the sink (the sink is located at the center of the square). In this example the space is split into $m = 4$ equal *square* partitions at each level. Figure 1(a) shows the square space splitting and Figure 1(b) shows the corresponding tree. The space contains 14 sensor nodes numbered 1 through 14. The numbers in the tree show the nodes involved in collision at each branch. At time slot 1, the space is split into 4 equal squares and nodes 1,2, and 3 transmit their packets as all of them belong to partition 1 at level 1[2]. Since this results in collision, partition 1 of level 1 is further split into 4 equal partitions.

Now, each partition has a single node. Therefore, node 1 successfully transmits its packet at time slot 2, node 2 at time slot 3 and node 3 at time slot 4. Time slot 5, allotted to partition 4 at level 2 of partition 1 at level 1, goes idle because it does not have any nodes in it. Similarly, nodes 4, 5, and 6 collide at time slot 6 and transmit successfully in time slots 7, 8, and 9. Time slot 10 goes idle as there are no nodes in partition 4 at level 2 of partition 2 at level 1. Nodes 7, 8, and 9 collide at time slot 11 and after time slot 12 goes idle, they successfully transmit

[2] We follow the convention of counting partitions from left to right and bottom to top.

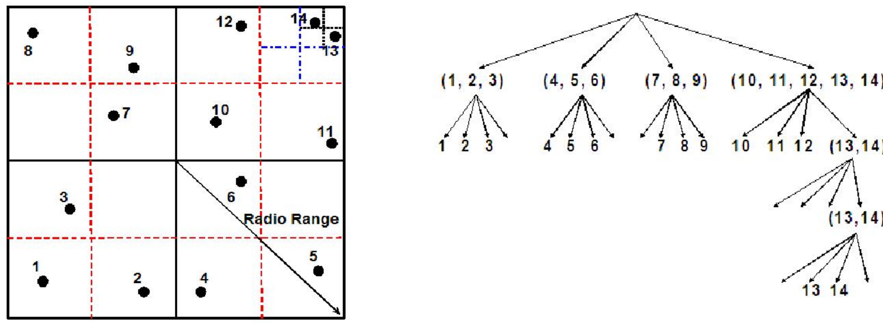

(a) (b)

Fig. 1. Example of the LAMA protocol for $m = 4$. (a) The square space splitting (b) The corresponding tree.

their packets at time slots 13, 14, and 15, respectively. At time slot 16, nodes 10, 11, 12, 13 and 14 belonging to partition 4 at level 1 transmit their packets and collide, leading to further splitting of that partition into 4 partitions at level 2. Since partitions 1, 2, and 3 at level 2 have nodes 10, 11, and 12, respectively, a single node each, all of them transmit their packets successfully at time slots 17, 18, and 19 respectively. At time slot 20, nodes 13 and 14 transmit their packets and collide. This results in further splitting of partition 4 at level 2 of partition 4 at level 1. Due to absence of nodes in partitions 1, 2, and 3 at level 3, time slots 21, 22, and 23 go idle. In time slot 24, nodes 13 and 14 transmit again, collide, and the partition is further split into 4 partitions at level 4. Due to absence of nodes in partition 1 at level 4, time slot 25 goes idle. Finally, nodes 13 and 14 successfully transmit their packets in time slots 26 and 27 respectively.

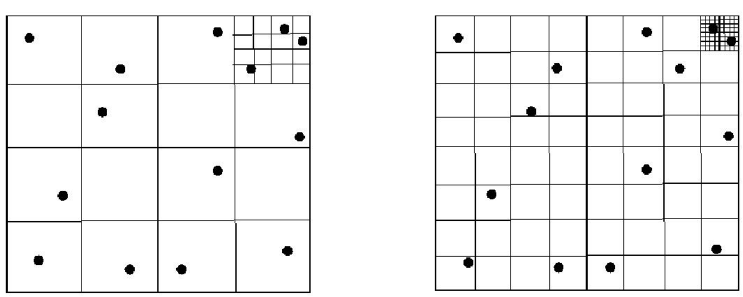

(a) (b)

Fig. 2. (a) 16-split strategy, ($m = 16$), $D(n) = 31$. (b) 64-split strategy, ($m = 64$), $D(n) = 99$.

Thus, in the above example, the delay for the sink to receive packets from all the 14 sensor nodes is $D(n) = 27$ time slots. Also, since the space is split into 4 equal square partitions at each level we call it a 4-split strategy. Similarly, Figure 2 illustrates the space splitting for 16-split and 64-split strategies.

3.1 Implementation Aspects

A key aspect in the implementation of the LAMA protocol is the determination of the nodes that belong to the current partition. This can be achieved by issuing *location tokens*, that contain the boundaries of the current partition, to the nodes at each time step. The location tokens are generated using the *Location Token Generator* (LTG), shown below for the m-split strategy, where m is a power of 4. The LTG uses the current splitting level, the partition numbers of all the levels, the sink location and its radio range to determine the boundaries of the current partition. The equations show that the boundaries are calculated relative to the lower left corner of the square space.

LTG$(L, \{P(l) : 1 \leq l \leq L\}, (s_x, s_y), S)$:

$$\mathbf{x_1} = \left(s_x - \frac{S}{2}\right) + \sum_{l=1}^{L}[(P(l) - 1) \bmod \sqrt{m}] \cdot \frac{S}{(\sqrt{m})^l}; \quad \mathbf{x_2} = x_1 + \frac{S}{(\sqrt{m})^L};$$

$$\mathbf{y_1} = \left(s_y - \frac{S}{2}\right) + \sum_{l=1}^{L}\left\lfloor\frac{P(l) - 1}{\sqrt{m}}\right\rfloor \cdot \frac{S}{(\sqrt{m})^l}; \quad \mathbf{y_2} = y_1 + \frac{S}{(\sqrt{m})^L};$$

Return (x_1, x_2, y_1, y_2);

- L: current level in the space splitting tree.
- $P(l)$: partition number at level l for the current partition.
- (s_x, s_y): location coordinates of the sink.
- S: side length of the square whose half-diagonal is equal to the radio range of the sink.
- x_1 is the left vertical boundary, x_2 is the right vertical boundary, y_1 is the lower horizontal boundary, and y_2 is the upper horizonal boundary.

The implementation of the protocol depends on where the location tokens are generated – at the sink or at the sensor nodes. In the former, the sink has to run the LTG and transmit the location token to the sensor nodes. This can be achieved by piggy-backing the location tokens on the ACK/NACK packets. In the latter, the sensor nodes themselves run the LTG and generate the location token at the beginning of each time slot. The advantage of the latter over the former is the small size of the ACK/NACK packets. This advantage is obtained at the cost of shifting the computational load of LTG from the sink to the sensor nodes.

Nevertheless, in either case, each sensor node decides if the location token belongs to it by verifying if its location falls within the boundaries specified by the location token. If the location token belongs to a node it transmits its packet, otherwise, it is ignored. Figure 3 shows the sink and sensor node state diagrams for the LAMA protocol for the implementation in which the sensor nodes determine the location tokens by themselves.

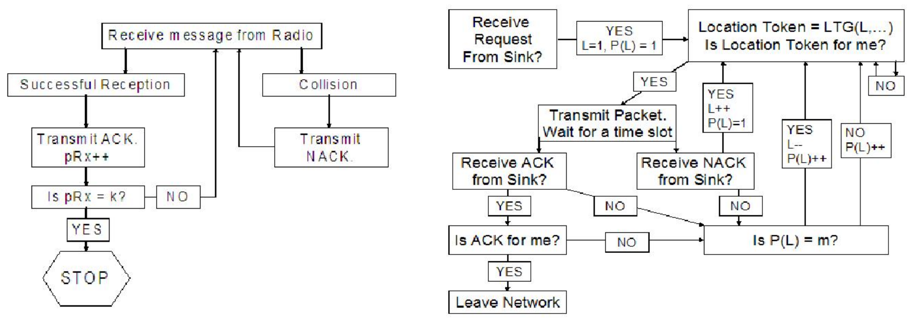

(a) Sink

(b) Sensor Node

Fig. 3. LAMA protocol state diagram at the sink and at the sensor node (pRx: packets received, L: split level, $P(L)$: partition at split level L).

4 Evaluation

In this section, we present results of performance evaluation of the LAMA protocol using simulations. We also present results of a comparative study with location-unaware MAC protocols.

4.1 LAMA Protocol

We consider a square space of $S \times S$ sq. length units with $S = 16$ (simulation results for other values of S did not shown any major difference.), populated by n sensor nodes, and the sink lodged at the center of the square. We evaluate the performance of the LAMA protocol for three different sensor node distributions – grid-random, even-random, and uniform-random. In grid-random distribution, the space is divided into a grid of 256 equal sized squares and nodes are placed such that each grid square is occupied by at most a single node. In even-random distribution, the space is divided into n equal sized partitions and each partition has at-most one sensor node. The procedure to divide the space into n equal partitions is described in the Appendix. In uniform-random deployment, each node is placed uniformly at random within the square space. We consider three symmetric[3] square space splitting strategies – 4-split, 16-split, and 64-split. The

[3] For an m-split symmetric square space splitting strategy, m should be a power of 4.

simulation results are averaged over 1000 random trials with 100 different random seeds. In each random trial the locations of the sensor nodes are different.

Now we discuss the performance of the LAMA protocol in terms of the delay and energy consumption per node as a function of the number of nodes (n) in the radio range of the sink, *i.e.,* for $k = n$. (please refer to Section 2).

Figure 4 shows the results for grid-random deployment of nodes. According to the figure, there exists a delay-energy trade-off for varying space splitting strategies. With increasing resolution of space splitting (4-split to 64-split), while the expected delay increases, the expected energy consumption per node decreases. This is expected because, even though increasing resolution reduces the number of collisions, thus reducing the energy consumption, it increases the number of idle time slots resulting in longer delays. However, it can be observed from the 4-split and 16-split graphs that the delay-energy trade-off vanishes after a certain number of nodes. For node numbers greater than ≈ 130, the 16-split strategy provides lower delay and lower energy consumption simultaneously, compared to 4-split strategy

The reason for this is that, for lower number of nodes, higher resolution space splitting (16-split) renders many time slots idle without any packet transmissions, thus increasing the delay and reducing the energy consumption. However, for higher number of nodes, higher resolution space splitting avoids more idle time slots on an average and thus reduces the delay. The results show that, for 64-split strategy, the reduction in energy consumption is almost negligible for higher number of nodes compared to the increase in delay (an order of magnitude) compared to 16-split. For the rest of the evaluation we consider the performance of only 4-split and 16-split as the delay due to 64-split is an order of magnitude higher.

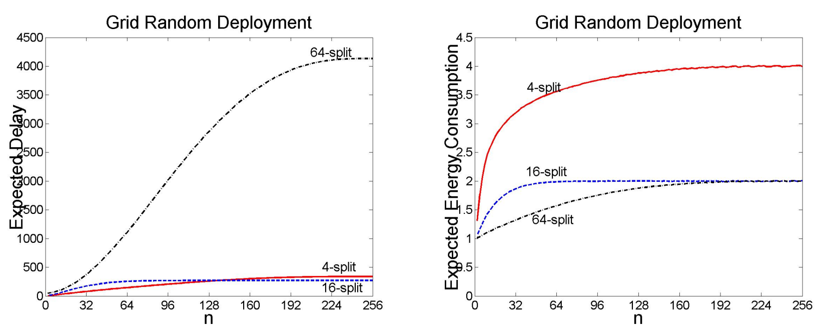

(a) (b)

Fig. 4. Expected delay and expected energy consumption per node due to 4-split, 16-split and 64-split strategies for grid-random placement of nodes.

Figure 5 plots the delay and energy consumption as a function of n for the 16-split strategy for the three random location distributions. According to the figure, the delay and energy consumption in grid-random deployment is lower than uniform-random deployment for all values of n, implying that the LAMA protocol can take advantage of node location distribution to provide better performance. However, the relative performance of 16-split for even-random deployment depends on the number of nodes. The reason for this is that, for lower number of nodes, the distribution of nodes in even-random deployment is more spread out in space, on an average, compared to that of uniform-random or grid-random deployments, resulting in lower number of split levels for even-random deployment compared to the other two deployments. Lower number of split levels implies lower delays and lower energy consumption. The graph for even-random deployment shows dips in delay and energy consumption for $n = 16$, $n = 64$, and $n = 256$. This is because, for these values of n, for even-random deployment, the space partition boundaries match exactly with that of the 16-split strategy, thus reducing the delay and energy consumption compared to that of their neighboring values.

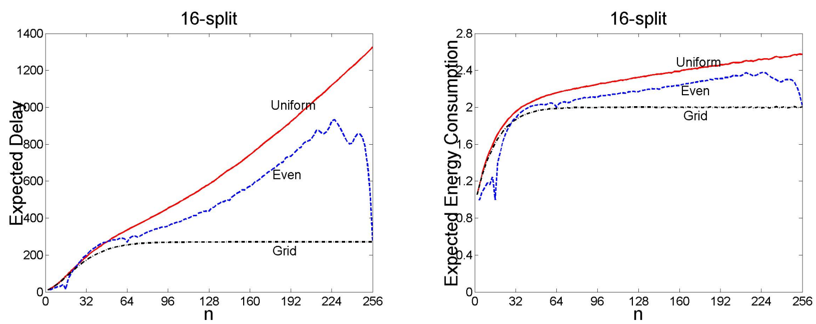

(a) (b)

Fig. 5. Expected delay and expected energy consumption per node due to 16-split strategy for three different location distributions.

For $n = 256$, even-random distribution is identical to grid-random distribution. Therefore, the delay and energy consumption are the same for both for these distributions for this value of n. The Figure also shows that the delay and energy consumption is almost constant for grid-random distribution for higher values of n. The reason for this is that, for grid-random deployment, with increasing number of randomly deployed nodes, the node density becomes more uniform across all split partitions. The corresponding split levels remain constant irrespective of the number of nodes (as long as the number does not cross the number of grid squares, 256), once a certain node density is crossed. This results in an almost constant delay and energy consumption for grid-random distribution of nodes.

4.2 Comparative Study

Next, we compare the performance of the LAMA protocol with the following three location-unaware protocols:

1. **HT-Split**: In this protocol [2], the collision domains of the sensor nodes are isolated probabilistically rather than spatially as done in LAMA protocols. The protocol starts out by all sensor nodes in the radio range of the sink tossing a coin, and the subset of nodes with a heads (H) transmitting their packet. If there is a collision, all nodes with a H in the first level, again toss a coin and the subset of nodes with a H in both the present and the previous levels transmit their packets. This is continued until a single node has H from all the previous levels and the present level. Once this node finishes transmitting its packet, the node with a tails (T) in the present level and a H in all the previous levels transmits its packet. This process of descending and ascending the "tree" of coin tosses continues until all nodes transmit their packets.
   We have chosen to compare the performance of the LAMA protocol to that of the HT-split protocol to show that the LAMA protocol, in addition to taking advantage of collision domain separation like the HT-split protocol, also takes advantage of the nodes' location distribution, to provide better delay and energy efficiency.
2. **Optimal p-persistent Slotted CSMA**: In the p-persistent slotted CSMA protocol ([2], [3]), each contending node senses the channel at the beginning of each time slot and if the channel is free it transmits its packet with probability p. If the channel is not free, the node attempts to transmit its packet in the next available free time slot with probability p. When the packet length is equal to that of a single time slot, this protocol is identical to p-persistent slotted Aloha. Even if the packet length is equal to multiple time slots, the packet-length-normalized delay and energy consumption will be identical to that of p-persistent slotted Aloha [3].
   Intuitively, the delay in p-persistent slotted Aloha can be minimized by dynamically changing the probability of transmission p for each time slot, to be the inverse of the number of nodes yet to successfully transmit their packets at the beginning of that time slot. In order to achieve this, the sink should keep track of the number of nodes in its radio range and the number of nodes that were able to successfully transmit their packets to it. The sink can then determine the probability of transmission that minimizes the delay and piggyback this information on the acknowledgement messages to the sensor nodes. We use this *optimal* p-persistent slotted CSMA protocol to benchmark the performance of the LAMA protocols.
3. **IEEE 802.15.4**: In order to compare the LAMA protocols' performance with a state-of-the-art, off-the-shelf, MAC protocol for sensor networks, we chose the recently standardized IEEE 802.15.4 protocol for low-rate, low-power personal area networks ([4], [5]). We adopt the 2.4 GHz ISM band and the star-topology options provided by the protocol. We consider the beacon-enabled mode in which the nodes are time synchronized with each other and use a variant of the non-persistent slotted CSMA-CA as the MAC protocol. The time period between two beacons is called a super-frame and it is divided into an active period and an optional inactive period. All communications take place in the active period and the inactive period can be used to power-down

the nodes to conserve energy. However, for fairness in comparison to other protocols, we ignore the inactive period and consider only the active period. The active period of the super-frame in turn consists of the contention access period (CAP) and the contention free period (CFP). In the CAP, channel access is through slotted CSMA-CA and in the CFP the channel access is through guaranteed time slots (GTS) which are mainly used for low-latency applications. In our implementation we ignore the CFP and assume that the the entire super-frame is made up of the CAP. We use default values for all the parameters of the standard; readers can refer to [4] or [5] for more details.

Figure 6 shows the expected delay and energy consumption for the above three location-unaware protocols as a function of n. Clearly, the optimal p-persistent slotted CSMA protocol is the best among the three location-unaware protocols. Also, the IEEE 802.15.4 standard, with the default parameters, performs the worst for higher number of nodes (after about $n = 40$). The main reason for this is that, for high number of nodes, due to multiple back-offs, most nodes quickly reach the highest back-off stage which has the lowest probability of transmission. The advantage of low probability of transmission is off-set by high number of nodes, leading to higher probability of collision. For the rest of the evaluation, we do not consider the IEEE 802.15.4 protocol as it gives delay which is orders of magnitude worse than the other two protocols, for high number of nodes. Thus, we compare the 4-split and 16-split strategies of the LAMA protocol with the HT-split and optimal p-persistent slotted CSMA protocols.

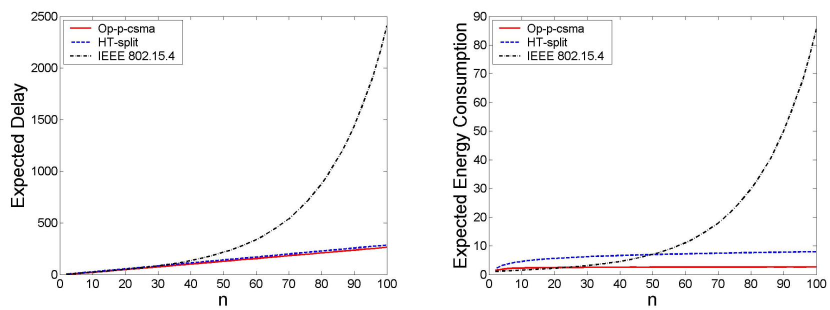

(a) (b)

Fig. 6. Expected delay and expected energy consumption for HT-split, optimal p-persistent slotted CSMA and IEEE 802.15.4 standrad.

Figure 7 compares the simulation results of the LAMA protocols with that of the two location-unaware protocols as a function of the number of nodes (n) and Figure 8 compares them as a function of k ($< n$) for $n = 200$. The main observations can be summarized as follows:

– The LAMA protocols clearly take advantage of the location distribution of the sensor nodes. This can be seen from Figures 7(a), (b) and (c). With in-

creasing order in the deployment of nodes, the delay is lower for the LAMA protocols compared to location-unaware protocols. This is because, with increasing order the number of split levels decrease, on average.
– As a result of the above advantage, for grid-random deployment of nodes, for high number of nodes, the delay due to 16-split is 60% lower and, simultaneously, the energy consumption is 30% lower compared to that due to optimal p-persistent slotted CSMA.
– Even for uniform-random deployment, the 4-split strategy performs close to or better than the optimal p-persistent slotted CSMA protocol.
– The gains for the LAMA protocols, in terms of lower delay, are higher for higher number of nodes. This is because the advantage due to location distribution becomes more significant when the number of nodes is higher. Nevertheless, the 4-split strategy gives lower delay compared to location-unaware protocols even for low number of nodes, albeit for higher energy consumption than for optimal p-persistent slotted CSMA. Even in this case, the energy consumption due to the 4-split strategy is lower than for HT-split protocol.
– Similar trends can be observed in Figure 8 for delay and energy consumption for the first k successful packet receptions.
– Interestingly the energy consumption graphs in the above figure show a periodicity as a function of k. This is because, in the tree of space splitting, there is a surge in energy consumption just before branching happens, when all nodes in a partition transmit their packets leading to collision and further splitting of the partition. This surge is the greatest for level 1 partitions and progressively reduces for higher levels. This is clear from the 4-split graph in which the biggest energy surges occur at multiples of 50 of k, which is the average number of nodes per partition at level 1.

5 Discussion

In this section, we discuss various issues inherent to the LAMA protocol and elaborate on its scope.

1. We have illustrated and evaluated the performance of the LAMA protocol for the case in which the space is symmetrically split into m equal squares at each level. However, intuition suggests that, the shape of space splitting does not affect the performance of the LAMA protocol as long as the its tree structure remains the same. For the same number of nodes and the same node location distribution, if the space is considered to be circular and if it is split into m equal sectors at each level (as shown in Figure 9(a) for $m = 4$), then, on an average, the delay and energy consumption of nodes would remain the same as that for square splitting. This intuition is verified by the simulation results shown in Figures 9(b) and (c).
2. The location of the sink is a crucial part in the implementation of the LAMA protocol. However, for some one-hop data-gathering applications such as localization [6], the location of the sink is not available. In fact, the application

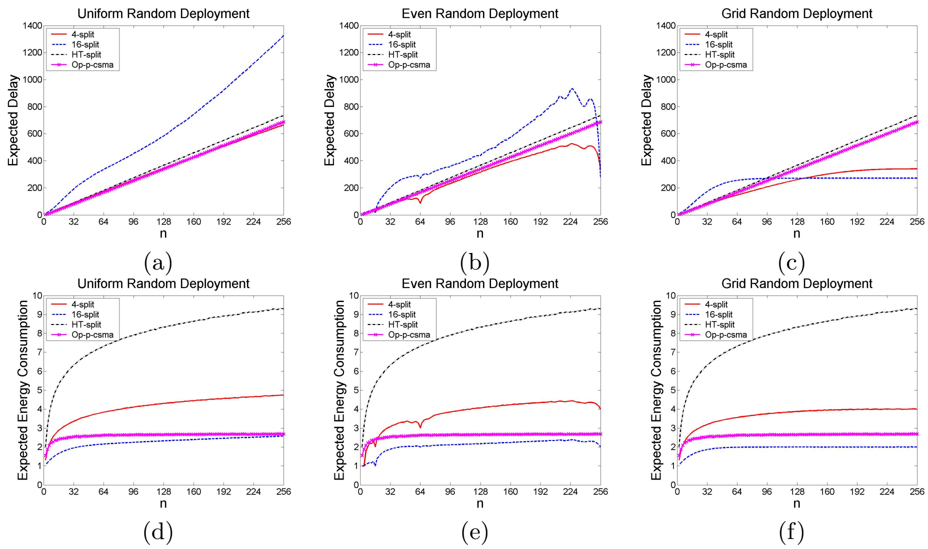

Fig. 7. Comparison of location-aware – 4-split, 16-split – and location-unaware – HT-split and optimal p-persistent slotted CSMA – medium access protocols as a function of n.

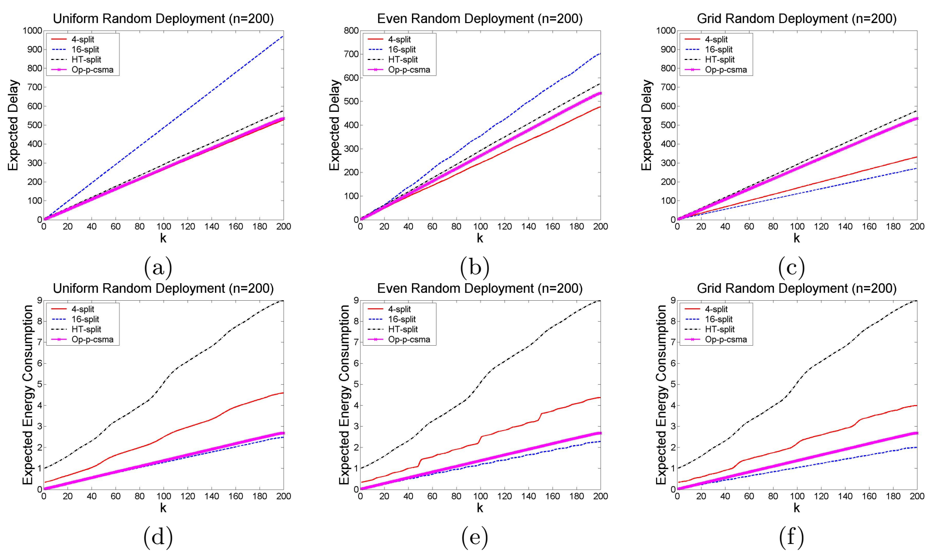

Fig. 8. Comparison of location-aware – 4-split, 16-split – and location-unaware – HT-split and optimal p-persistent slotted CSMA – medium access protocols as a function of k.

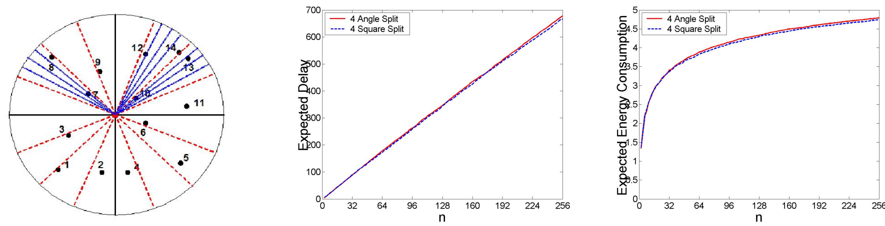

(a) (b) (c)

Fig. 9. (a) 4–Angle–Split Strategy. Comparison of (b) expected delay and (c) expected energy consumption, for angular splitting and square splitting strategies for uniform-random deployment of nodes.

has to determine the location of the sink. This problem can be solved by first assuming an approximate location for the sink and then using the LAMA protocol to obtain data-packets from the sensor nodes. We propose to use transmission power control for this purpose.

The main idea here is that the sink assumes the location of the nearest sensor node and uses this location to obtain packets from other nodes in its radio range. The sink can obtain the location of the nearest sensor node by using power control, in which, its transmission power is incremented by small steps starting from the lowest power until it is able to reach a sensor node and receive a packet from it. In the possibility of the existence of more than one nodes in the lowest connected radio range of the sink, the nodes can contend for the channel using a random medium access scheme such as p-persistent slotted CSMA. The key observation here being that, for typical node densities, the number of nodes in the lowest connected radio range of the sink is very low compared to the typical operational radio range.

For example, Tmote-sky [7] devices have a radio range of above 100 m for the highest transmission power of 0 dBm. For the lowest transmission power of -25 dBm the radio range is less than 6 m. Thus, for a uniform node density, the number of nodes in the lowest-power connected radio range is at-least two orders of magnitude lower than the number in the highest-power radio range.

For the case in which more than one node exists in the lowest connected radio range of the sink, the delay in determining its nearest sensor node is the delay for the first packet to reach the sink. This delay depends on the density of node distribution, the topology of the network and the reliability of the wireless channel. If the number of nodes in the lowest-power connected radio range is a and if the nodes use p-persistent slotted CSMA [3] with T back-off slots (since each node chooses uniformly at random to transmit, the probability of transmission at each time slot is $p = \frac{1}{T}$), then the expected number of time slots for the first successful transmission is given by:

$$ap(1-p)^{a-1} + 2 \cdot (1 - ap(1-p)^{a-1}) \cdot ap(1-p)^{a-1}$$
$$+3 \cdot (1 - ap(1-p)^{a-1})^2 \cdot ap(1-p)^{a-1} + \cdots \quad (1)$$
$$= \frac{1}{ap(1-p)^{a-1}} \quad (2)$$

Figure 10 shows the behavior of the above equation for different values of a and T. As the figure shows, the value of T can be chosen such that the delay due to determination of the sink's location is very low, usually much lower than 10 time slots.

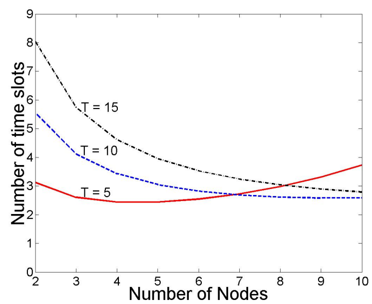

Fig. 10. The expected number of time slots for the first successful transmission.

3. In the LAMA protocol we have illustrated and evaluated, the space is split into m equal parts at each level. However, the value of m can be changed, adaptively, at each level, depending on the sensor node deployment density. For example, for a given sensor node density, the number of nodes in partitions of higher split levels is lower than that in lower split levels. This fact can be taken advantage to adaptively reduce the value m for higher split levels, thus, reducing the number of idle time slots and consequently reducing the delay. Or alternatively, since lower number of nodes contend for the channel at higher split levels, random medium access techniques such as CSMA-CA could be used in conjunction with the LAMA protocol; and potentially reduce the delay and energy consumption.

6 Related Work

The contributions of this paper are two fold - the LAMA protocol and the introduction of the single-packet medium access problem. In this section, we will discuss the prior related work in both these directions.

To the best of our knowledge, very little research has been done for the single-packet medium access problem. Tay *et.al.* in [8] have proposed a collision minimizing CSMA protocol applicable to event-driven applications in wireless sensor networks. The main focus of their work is in minimizing the delay for the first successful packet reception. However, their work differs from ours in that the authors assume that each node in the network has multiple backlogged packets to transmit, whereas, the single-packet medium access problem is defined by the existence of a single packet for transmission at each node.

There have been some prior efforts to incorporate location information in medium access control protocols. Corbett *et. al.* in [9] propose a hybrid TDMA – Contention based protocol for multi-hop sensor networks that uses the locations of nodes for spatial reuse and time slot allocation to avoid collisions and interference. The space is divided into hexagonal cells, similar to cellular networks, and nodes within each cell use contention based medium access. In contrast to this, in our work, we use the locations of nodes to solve the problem of medium access within a cell. Liu *et. al.* in [10] use the location information of nodes within one-hop to provide energy efficiency and fault tolerance, even though the medium access is through contention-based random-access schemes. In our work, the medium access itself is based on the locations of nodes. Nadeem *et. al.* in [11] use the location information in tandem with the capture effect to increase throughput in IEEE 802.11 DCF networks. In this, the location information is used to increase the spatial reuse efficiency and better manage interference leading to additional concurrent transmissions, thereby increasing the overall throughput of the protocol. Again, this work differs from ours in that, we solve the one-hop medium access problem using the location information of nodes in contrast to the multi-hop one.

7 Conclusions & Future Work

In this paper, we have presented a novel location-aware medium access (LAMA) protocol for wireless sensor networks. We studied in detail the application of this protocol for the single-packet medium access problem that appears frequently in many sensor network applications such as neighbor discovery, sensor data gathering, etc. The defining aspect of the single-packet medium access problem is that each sensor node in the one-hop radio range of the sink has a single packet to transmit. We illustrated the working of the LAMA protocol for this problem using examples and discussed its implementation aspects. The main idea in the LAMA protocol is a tree-based hierarchical partitioning of space to progressively reduce the collision domains of nodes until there are no collisions.

We then presented results from a thorough performance evaluation of the LAMA protocol in comparison to three location-unaware MAC protocols – HT-split, optimal p-persistent slotted CSMA, and the IEEE 802.15.4 standard protocol – using simulations. We evaluated the protocol for three different location distributions of nodes – uniform-random, even-random, grid-random. Results showed that the LAMA protocol takes advantage of the location distribution of nodes to provide significant gains – up to 60% lower delay and 30% lower energy consumption, simultaneously – compared to the chosen location-unaware MAC protocols.

In the future, we would like to analyze the LAMA protocol and derive closed form expressions for delay and energy consumption. The LAMA protocol presented in this paper is open to many possible improvements and enhancements as discussed in Section 5. In the future, we wish to explore such possible performance enhancers for the protocol. We are also interested in studying the performance of the protocol in real systems implementation. We wish to address real implementation concerns such as the effect of errors in the locations of nodes on the performance of the protocol. Another aspect of interest to us is the application of this protocol for 3-dimensional sensor node deployments.

Appendix

Procedure to obtain even-random distribution of nodes: Below, we illustrate the steps to divide a square region into n equal area partitions:

1. Divide the square into $x = \lfloor \sqrt{n} + 0.5 \rfloor$ vertical partitions.
2. Each vertical partition will be divided into a minimum of $y_{min} = \lfloor \frac{n}{x} \rfloor$ horizontal partitions.
3. Let $r = n - x \cdot y_{min}$. Determine the number of horizontal partitions y_i in each vertical partition i $(1 \leq i \leq x)$ using,

$$y_i = \begin{cases} y_{min} + 1, & 1 \leq i \leq r \\ y_{min}, & r < i \leq x \end{cases}$$

4. The width of vertical partition i is $W_i = S \cdot \frac{y_i}{n}$, where S is the side length of the square.
5. Finally, divide vertical partition i into y_i equal parts.

References

1. Karp, B., H.T.Kung: Greedy Perimeter Stateless Routing for Wireless Networks. In: International Conference on Mobile Computing and Networking. (2000)
2. Bertsekas, D., Gallager, R.: Data Networks (Second Edition). Prentice Hall (1991)
3. Bruno, R., Conti, M., Gregori, E.: Optimization of Efficiency and Energy Consumption in p-Persistent CSMA-Based Wireless LANs. IEEE Transactions on Mobile Computing **1**(1) (2002) 10–31
4. Callaway, E., Gorday, P., Hester, L., Gutierrez, J.A., Naeve, M., Heile, B., Bahl, V.: Home Networking with IEEE 802.15.4: A Developing Standard for Low-Rate Wireless Personal Area Networks. IEEE Communications Magazine (2002) 70–77
5. Ramachandran, I., Das, A.K., Roy, S.: Analysis of the Contention Access Period of IEEE 802.15.4 MAC. Accepted for publication in ACM Transactions on Sensor Networks. (2005)
6. Yedavalli, K., Krishnamachari, B., Ravula, S., Srinivasan, B.: Ecolocation: A sequence based technique for rf localization in wireless sensor networks. In: The Fourth International Conference on Information Processing in Sensor Networks (IPSN 2005), Los Angeles, CA (2005)
7. : (http://www.moteiv.com/products/docs/tmote-sky-datasheet.pdf)
8. Tay, Y.C., Jamieson, K., Balakrishnan, H.: Collision Minimizing CSMA and its Applications to Wireless Sensor Networks. IEEE Journal on Selected Areas in Communications (2004)

9. Derek J. Corbett and David Everitt: A Partitioned Power and Location Aware MAC Protocol for Large Ad Hoc Networks. In: European Wirless 2005, Nicosia, Cyprus (2005)
10. Liu, B.H., Bulusu, N., Pham, H., Jha, S.: A Self-Organizing, Location Aware Media Access Control Protocol for DS-CDMA Sensor Networks. In: IEEE International Conference on Mobile Ad Hoc and Sensor Networks. (2004) 528–530
11. Nadeem, T., Ji, L., Agarwala, A., Agre, J.: Location Enhancement to IEEE 802.11 DCF. In: IEEE INFOCOM, Miami, Florida, USA (2005)