

UC San Diego

Technical Reports

Title

Improved Distributed Simulation of Sensor Networks based on Sensor Node Sleep Time

Permalink

<https://escholarship.org/uc/item/73f2501s>

Authors

Jin, Zhong-Yi
Gupta, Rajesh

Publication Date

2008-02-06

Peer reviewed

Improved Distributed Simulation of Sensor Networks based on Sensor Node Sleep Time

Zhong-Yi Jin

Dept. of Computer Science & Engineering
University of California, San Diego
Email: zhjin@cs.ucsd.edu

Rajesh Gupta

Dept. of Computer Science & Engineering
University of California, San Diego
Email: rgupta@cs.ucsd.edu

Abstract

Sensor network simulators are important tools for the design, implementation and evaluation of wireless sensor networks. Due to the large computational requirements necessary for simulating wireless sensor networks with high fidelity, many wireless sensor network simulators, especially the cycle accurate ones, employ distributed simulation techniques to leverage the combined resources of multiple processors or computers. However, the large overheads in synchronizing sensor nodes during distributed simulations of sensor networks result in a significant increase in simulation time. In this paper, we present a novel technique that could significantly reduce such overheads by minimizing the number of sensor node synchronizations during simulations. We implement this technique in *Avrora*, a widely used parallel sensor network simulator, and achieve a speedup of up to 11 times in terms of average simulation speed in our test cases. For applications that have lower duty cycles, the speedups are even greater since the performance gains are proportional to the sleep times of the sensor nodes.

I. INTRODUCTION

Simulations of wireless sensor networks (WSNs) are important to provide controlled and accessible environments for developing, debugging and evaluating WSN applications. In simulations, sensor network programs run on top of simulated sensor nodes inside simulated environments. Since the simulated entities (simulation models) are transparent to simulation users, the states and interactions of sensor network programs can be inspected and studied easily and repeatedly. In addition, the properties of the simulated entities such as the locations of sensor nodes and the inputs to the sensor nodes can be changed conveniently before or during simulations. One of the key requirements for simulating WSNs is high simulation fidelity in terms of temporal accuracy of events and actions. High fidelity simulations usually come at the cost of increased simulation time and poor scalability because significant computational resources are required for running high fidelity simulation models. Parallel and distributed simulators [1] are developed to address these issues by leveraging the combined resources of multiple processors/cores on a same computer and on a network of computers, respectively. They can significantly improve simulation speed and scalability because sensor nodes may be simulated in parallel on different processors or computers. Since different nodes may get simulated at different speeds in this fashion, simulated nodes often need to synchronize with each other to preserve causalities and ensure correct simulation results.

Synchronizations of sensor nodes are illustrated in Figure 1 which shows the progress of simulating in parallel two sensor nodes that are within direct communication range of each other. There are two notions of time in the figure, wallclock time T_W and simulation time T_S . Wallclock time corresponds to the actual physical time while simulation time is the virtual clock time that represents the physical clock time of real sensor nodes in simulations [1]. At wallclock time T_{W1} , the simulation time of Node A (T_{SA}) is T_{S2} and the simulation time of Node B (T_{SB}) is T_{S1} . Node B is simulated slower than Node A in this case ($T_{SB} < T_{SA}$) because the thread or process that is used to simulate Node B either runs on a slower processor or receives fewer CPU cycles from an operating system (OS) task scheduler. At T_{S2} , Node A is supposed to read the wireless channel and continue its execution along different execution paths based on whether there are active wireless transmissions or not. However, despite the fact that the simulation time of Node A already reaches T_{S2} at T_{W1} , Node A probably can not advance any further because at T_{W1} Node A does not know whether Node B is going to transmit at T_{S2} or not (since $T_{SB} < T_{S2}$). There are two general approaches to handle cases like this, a conservative one, e.g., [2] and an optimistic one, e.g., [3]. With the conservative approach, Node A has to wait at T_{S2} until the simulation time of Node B reaches T_{S2} . The optimistic

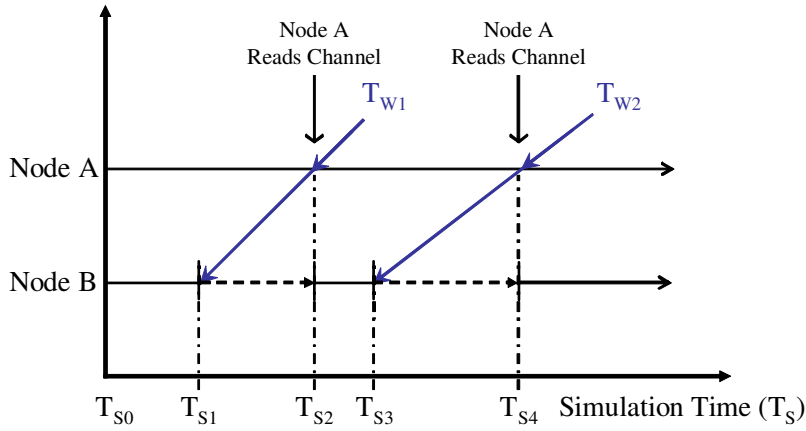


Fig. 1. The progress of simulating in parallel a wireless sensor network with two nodes that are in direct communication range of each other

approach, on the other hand, would allow Node A to advance assuming there will be no transmissions from Node B at T_{S2} . However, the entire simulation state of Node A at T_{S2} has to be saved. If Node A later detects that Node B actually transmits to Node A at T_{S2} , it can correct the mistake by rolling back to the saved state and start again.

Almost all distributed WSN simulators are based on the conservative approach as it is simpler to implement and has a lower memory footprint. The conservative approach mainly involves three time-consuming steps. In the first step, the threads or processes that simulate the waiting nodes (Node A in Figure 1), have to be suspended. This would usually involve context switches to swap in other threads or processes that simulate other nodes. In the second step, the nodes that some other nodes are waiting for (Node B in Figure 1) have to notify the waiting nodes about their progresses. For example, in Figure 1, Node B must notify Node A after it advances past T_{S1} so that Node A can continue. In the last step, the waiting nodes need to be swapped back into execution once they are notified to continue. Step 1 and 3 are time-consuming as they involve context switches. Large numbers of context switches would significantly increase simulation time. Step 2 involves expensive inter-thread or inter-process communications and is generally the slowest step in distributed simulations as the notifications may have to be sent through slow networks to different computers.

Therefore, the performance gains in existing distributed WSN simulators are often compromised by the rising overheads due to inter-node synchronizations. This is reported in both Avrora [4], a cycle accurate simulator that runs over SMP (shared memory multiprocessor) computers, and DiSenS [5], a cycle accurate simulator for simulations over both SMP computers and a network of computers. In the case of Avrora, the reported time of simulating 32 nodes with 8 processors is only about 15% less than using 4 processors because of a large number of thread context switches introduced by synchronizations. In DiSenS, it is actually faster to simulate 4 nodes using 1 computer (dual-processor) than using 2 computers and simulating 2 nodes on each in most of the testing cases. This sub-linear performance in DiSenS is due to the large communication overheads in synchronizing nodes that are simulated on different computers.

In this paper, we describe a novel technique that could significantly reduce the number of synchronizations in distributed simulations of WSNs. Our technique exploits the sleep-often property of sensor network applications and uses sleep times to reduce sensor node synchronizations. It works without any prior knowledge of the sensor network applications under simulations and the number of reductions scale with both network sizes and sleep times. We demonstrate the value of our approach by its implementation in Avrora [4].

We describe our speedup technique in Section II. Its implementation is presented in Section III. In Section IV we describe the results of our experiments followed by a brief overview of the related work and conclusions in Section V and Section VI.

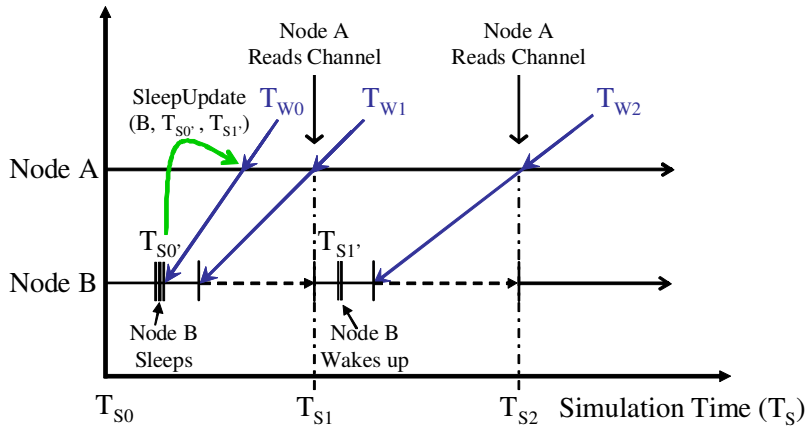


Fig. 2. The progress of simulating in parallel a wireless sensor network with two duty cycled nodes that are in direct communication range of each other

II. A NOVEL TECHNIQUE TO REDUCE SYNCHRONIZATIONS IN DISTRIBUTED WSN SIMULATIONS

Sensor node synchronizations are required for enforcing dependencies between sensor nodes in simulations. Since the dependencies come from the interactions of sensor nodes over wireless channels, the number of required synchronizations in a distributed simulation is inversely proportional to the degree of parallelism in the WSN application under simulation [1].

To reduce the number of synchronizations, we seek to exploit the parallelism available in WSN applications. In particular, we seek to use the information regarding duty cycling of nodes in sensor networks to speed up simulations. Node duty cycling is common in WSN applications for power management purposes. Most WSN applications have very low duty cycles and need to carefully manage their power consumptions in order to function for an extended period of time under the constraint of limited energy supply. In other words, sensor nodes sleep most of the time and do not interact with each other frequently until certain events are detected. For example, in the Great Duck Island sensor network deployment [6], each sensor node reads and transmits data from its sensors every 70 seconds. After successfully transmitting the data, the node is put into a sleep state for the next 70 seconds. Very little power is consumed by a sensor node in the sleep state as its wireless radio is turned off and its processor is put into a low power sleep mode.

Our speedup technique is illustrated in Figure 2 which shows the progress of simulating two duty cycled sensor nodes that are within direct communication range of each other. In the simulation, Node B enters into the sleep state at $T_{S0'}$ and wakes up at $T_{S1'}$. With existing distributed WSN simulators, Node A needs to wait for Node B at T_{S1} although Node B does not transmit anything during its sleep period. To eliminate this type of unnecessary synchronization, our technique keeps track of the time that a node enters into the sleep state and the time it wakes up. When we detect during a simulation that a node is entering into the sleep state, we immediately send both the entering time and exiting time (simulation time) in a *SleepUpdate* message to the neighboring nodes that are within direct communication range. We can detect the entering time because some special instructions have to be executed to put the node processor into the sleep mode. We can track the exiting time because the wakeup time has to be passed to a physical timer before the node is able to enter into the sleep state for a node processor can not execute any instructions during the sleep mode. As a result, neighboring nodes no longer need to synchronize with the sleeping node during the sleep period. For example, when we detect that Node B is entering into the sleep state at $T_{S0'}$, we immediately notify Node A that Node B will be in the sleep state from $T_{S0'}$ to $T_{S1'}$. Once Node A knows that Node B will not transmit between $T_{S0'}$ and $T_{S1'}$, and $T_{S0'} \leq T_{S1} < T_{S1'}$, it no longer needs to wait for Node B at T_{S1} and its lookahead time increases. Lookahead time is defined as the amount of simulation time that a simulated sensor node can advance freely without waiting for inputs from other simulated sensor nodes [1]. The speedup of our technique increases with the durations of sleep periods because the longer the sleep periods, the larger the lookahead time.

We keep the overhead due to tracking of sensor sleep times low by ensuring that the code to detect the sleep state and track wakeup times is only executed when certain sleep events are triggered. Sleep events are integral parts of any WSN simulators and are used in simulations to signal the transitions of nodes from active states to the sleep state. In the case of cycle accurate sensor network simulators, sleep events are associated with the execution of specific machine instructions of the sensor node processors under simulation. Also, the frequency of sleep events is typically not very high and the use of the speedup technique can be controlled based on sleep times. Sleep durations thus can be used to reduce sensor node synchronizations in simulations.

III. SYNCHRONIZATION ALGORITHM AND IMPLEMENTATION

We use Avrora [4] to evaluate the effectiveness of our approach. Avrora is a widely used cycle accurate sensor network simulator. Among all types of sensor network simulators, cycle accurate sensor network simulators [7], [4], [5] offer the highest level of fidelity. They provide simulation models that emulate the functions of major hardware components of a sensor node, mainly the processor. Therefore, one can run on top of them, clock cycle by clock cycle, instruction by instruction, the same binary code (images) that are executed by real sensor nodes. In other words, a sensor network can be simulated by running in simulation the same number of copies of the same images that are on the real sensor nodes of the real sensor network. As a result, accurate timing and interactive behaviors of sensor network applications can be studied in details.

Avrora is written in Java and supports parallel simulations of sensor networks comprised of Mica2 Motes [8]. It allocates one thread for each simulated node and relies on the Java virtual machine (VM) to assign runnable threads to any available processors on an SMP computer. Implementing the speedup technique in Avrora mainly involves developing new code in two areas: synchronization algorithm and channel modeling. Our implementation is based on the Beta 1.6.0 code release of Avrora. Its well documented source code is publicly available.

A. Synchronization Algorithm

The lock-step style synchronization algorithm of Avrora is optimized for parallel simulations on SMP computers but lacks necessary features to support our speedup technique. Our distributed synchronization algorithm is shown in Algorithm 1. It is a generic distributed synchronization algorithm similar to the one in DiSenS [5] and is suitable for both parallel and distributed simulations of WSNs.

Synchronizations are only necessary between neighboring nodes that are within direct communication range of each other. The first step before applying our algorithm is to build a neighbor node list for each node according to the locations of the sensor nodes and the transmission range of their wireless radios. Then, a time stamp is assigned to every node (node id) in the lists to keep the last reported simulation time of that node. This list, named nl in Algorithm 1 is the first required input to the synchronization algorithm. There are two more inputs to the algorithm. The second input id is used to identify the node under simulation. The nodes in nl are neighbors of this node. The third input $bytetime$ is the amount of time to transmit one byte with a wireless radio. It is the maximum lookahead time without synchronizations. Every node starts with that lookahead time because it takes that amount of time for one byte of data to travel from the sender to the receiver. For example, if a node starts at simulation time 0 and wants to read the wireless channel at that time, it can do so because the earliest time that a byte of data can arrive is $0 + bytetime$. Similarly, after synchronizing at time T_S , all synchronized nodes can advance freely up to $T_S + bytetime$ without any additional synchronizations. However, this approach only works if the processor and radio on a real sensor node communicate by exchanging data one byte at a time (*byte-level*).

The variable *intervalclock* in Algorithm 1 is used to ensure that *ClockUpdate* messages are sent by every simulated node once every *bytetime* if the node is not already in the sleep state. These messages update neighboring nodes about the latest simulation time of the sender and ensure neighboring nodes have the right time information to make synchronization decisions according to Condition 1. The interval chosen to send the messages will affect the performance of the algorithm as nodes may have to wait if *ClockUpdate* messages are delayed. The smallest interval one can use with *byte-level* radios is *bytetime* because the actual waiting time must fall on byte boundaries. This can be seen in Algorithm 1 where the *floor* function is used to calculate the lookahead time.

Condition 1. If a node N_i reads data sent by a node N_s over a wireless channel C_k at simulation time T_{SN_i} , then the simulation time of node N_s , T_{SN_s} , must be greater than or equal to T_{SN_i} .

Algorithm 1 Distributed Synchronization Algorithm with Speedup Technique

Require: $nl := \{ \langle nid, nclock \rangle \}$ /*a list of neighboring node ids and their reported simulation time*/
Require: id /*current node ID*/, $bytetime$ /*the amount of time to transmit one byte with a wireless radio*/

- 1: $clock \leftarrow 0$ /*current sim clock to 0*/, $lookahead \leftarrow bytetime$, $intervalclock \leftarrow 0$
- 2: **for** every tuple $\langle nid, nclock \rangle$ in nl **do**
- 3: $nclock \leftarrow 0$ /*initialize simulation time of neighboring nodes to zero before starting the simulation*/
- 4: **end for**
- 5: **while** $clock \leq$ user inputed simulation time **do**
- 6: $waitchannel \leftarrow false$
- 7: execute *next instruction*
- 8: **if** the *instruction* puts a node into the sleep state **then**
- 9: $exitclock = wakeuptime$, $intervalclock \leftarrow exitclock$
- 10: send a $ClockUpdate(id, exitclock)$ message to every nid node in the tuple $\langle nid, nclock \rangle$ of nl
- 11: **else if** the *instruction* reads from the wireless radio **then**
- 12: **if** $lookahead \geq 0$ **then**
- 13: read the wireless radio
- 14: **else**
- 15: **if** $intervalclock \neq clock$ **then**
- 16: $intervalclock \leftarrow clock$
- 17: send a $ClockUpdate(id, clock)$ message to every nid node in the tuple $\langle nid, nclock \rangle$ of nl
- 18: **end if**
- 19: $waitchannel \leftarrow true$
- 20: **end if**
- 21: **end if**
- 22: **if** $waitchannel$ is *false* **then**
- 23: $clock \leftarrow clock + cyclesconsumed$ /*advance clock by the clock cycles of the executed instruction*/
- 24: **end if**
- 25: $updated \leftarrow false$ /*check incoming $ClockUpdate$ messages at least once per instruction*/
- 26: **repeat**
- 27: **for** each received $ClockUpdate(cid, cclock)$ message **do**
- 28: **for** every tuple $\langle nid, nclock \rangle$ in nl **do**
- 29: **if** uid equals to cid **then**
- 30: $nclock \leftarrow cclock$
- 31: **end if**
- 32: **end for**
- 33: **end for**
- 34: $updated \leftarrow true$
- 35: $minclock = \min(nclock)$ in the tuple $\langle nid, nclock \rangle$ of nl /*find the neighbor with the smallest clock*/
- 36: $lookahead \leftarrow (minclock - \text{floor}(clock/bytetime) * bytetime)$ /*on byte boundaries with byte-radio*/
- 37: **if** $lookahead \geq 0$ and $waitchannel$ is *true* **then**
- 38: read the wireless radio /*all neighbors have advanced past the byte boundary this node is waiting on*/
- 39: $clock \leftarrow clock + cyclesconsumed$, $waitchannel \leftarrow false$
- 40: **end if**
- 41: **until** $waitchannel$ is *false* and $update$ is *true*
- 42: **if** $(clock - intervalclock) \geq bytetime$ **then**
- 43: $intervalclock \leftarrow clock$
- 44: send a $ClockUpdate(id, clock)$ message to every nid node in the tuple $\langle nid, nclock \rangle$ of nl
- 45: **end if**
- 46: **end while**

The *SleepUpdate* message described in Section II is replaced in Algorithm 1 with the *ClockUpdate* message because receiving nodes only need to use the wakeup time of the sender to calculate lookahead time. However, to take advantage of a special optimization described in the future work part of Section VI, *SleepUpdate* messages must be used so the time that a node enters into the sleep state is sent as well. The time that a node enters into the sleep state is detected when the *Sleep* instruction of the ATmega128L microcontroller [9] is executed. This instruction can put the microcontroller into different sleep modes based on the values set in the *MCU Control Register*. It is critical to read that register before sending any *ClockUpdate* messages as the speedup technique only works if a microcontroller is put into *Power-Save Mode*. The reason is that the only way to wake up a microcontroller from *Power-Save Mode* is through a timer triggered interrupt. Because of that, we can find the wakeup time by keeping track of the values written to the *Timer Control Register*. However, the same is not true if the microcontroller is put into other low power modes. In those modes, devices connected to the microcontroller via SPI or *I²C* interface, such as the CC1000 radio on the Mica2, can wake up the microcontroller at any time.

The computational complexity of a synchronization algorithm is determined by the total number of synchronization messages that need to be sent and the overheads in sending and processing each of the messages [10]. Our synchronization algorithm has higher computational complexity than the one in Avrora because we choose to implement it as a generic distributed algorithm. In Avrora, a global data structure is used to keep track of the simulation time of each node and therefore a node only needs to update the global data structure once for each clock update. This centralized approach is optimized for parallel simulation over SMP computers but does not support distributed simulations over a network of computers. We implement our synchronization algorithm as a truly distributed algorithm by distributing parts of the global data structure to each node. The penalty is that a node with N nodes in direct communication range has to send a total of N messages for each clock update. However, the penalty is not significant when the number of nodes within direct communication range is not big. In fact, because the synchronization algorithm of Avora works by synchronizing a node with all other nodes regardless of whether they are within communication range or not, our distributed algorithm may even perform better when nodes are sparsely distributed. If performance is an issue in the future, we could choose to optimize our implementation for parallel simulations using a centralized approach.

B. Channel modeling

With our synchronization algorithm, a node can write to a wireless channel long before the packets are read by other nodes (the transmitting code is omitted in Algorithm 1 for simplicity). Because of that, we develop a new wireless channel model that uses a circular buffer to store unread wireless packets. Our channel model uses a similar method as the original channel model in Avrora to map transmitting and receiving time into time slots that are *bytetime* apart. The slot number is used to index into the circular buffer. Note that with the original channel model, a write to a channel right after synchronizations could be dropped as the write may happen before the time slots are carried forward. Our channel implementation does not drop data.

IV. EVALUATION

We conduct a series of experiments to evaluate the performance of our speedup technique. The experiments are conducted on an SMP server running Linux 2.6.9. It has 4 processors (Intel Xeon 2.8GHz) and 2GByte of RAM. For comparison, the same test cases are simulated using both our modified Avrora and the original Avrora. Sun's Java 1.6.0 is used to run both simulators.

To demonstrate the effectiveness of our speedup technique, we choose two programs from the CountSleepRadio example which is a part of the TinyOS 1.1 distribution [11], [12]. These two programs behave exactly like the CntToRfm and RfmToLeds programs used in the experiments of the Avrora paper and serve similar purposes. The only difference is that the new programs can put sensor nodes into sleep states to save power. (The latest TinyOS 2.0 release [13] is not used for our experiments because its radio stack is not fully compatible with the radio model in the version of the Avrora that our code is based on.)

The first program we use for our experiments is CountSleepRadio. It wakes up a node periodically from the sleep state to increase the value of a counter by one and broadcast that value in a packet. Once the packet is sent, it puts the node back into the sleep state to save power. We have to modify this program for some of our




Fig. 3. Average number of synchronizations per node in one-hop networks during 60 seconds of simulation time

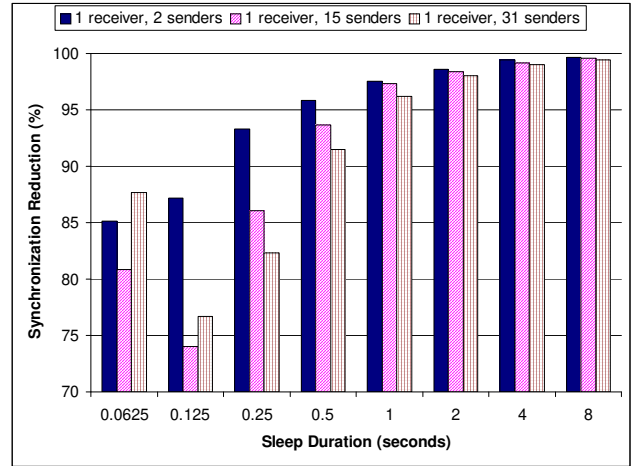


Fig. 4. Percentage reductions of the average number of synchronizations per node in one-hop networks during 60 seconds of simulation time

experiments because the original program has an upper bound on how long a sensor node can stay in the sleep state¹. This is unnecessary and we work around this limitation by using the *Clock* interface of the TinyOS 1.1 directly. The problem has reportedly been fixed in TinyOS 2.0. The counterpart of our CountSleepRadio program is CountReceive. It receives packets sent by CountSleepRadio and flashes different LEDs based on the values in the packets. For simplicity, we identify nodes running CountSleepRadio and CountReceive as senders and receivers respectively in the following sections.

Both CountSleepRadio and CountReceive use the default TinyOS 1.1 CC1000 CSMA (carry sense multiple access) MAC (media access control) which is based on B-MAC [14]. Before sending a packet, the CC1000 MAC first backs off for a random amount of time and then reads its transmitting channel for ongoing transmissions. It only sends the packet if the channel is clear. Otherwise, it backs off for a random amount of time before checking the channel again. As a result, a sender in our experiments reads the wireless channel at least once before each transmission.

A. Performance in one-hop networks

In this section, the performance of our speedup technique is evaluated under various sleep times and network sizes using one-hop sensor networks. One-hop sensor networks are sensor networks set up in such a way that all sensor nodes are within direct communication range of each other. It is a common form of sensor network used in actual deployments. For example, the patch networks deployed on the Great Duck Island [6] for habitat monitoring are one-hop networks. We choose to use one-hop sensor networks to evaluate the speedup technique because the number of nodes that need to be synchronized in one-hop networks can be controlled easily by changing the network size.

In the one-hop sensor network experiments, nodes are laid on a 10 by 10 grid 1 meter apart and their maximum transmission ranges are set to 20 meters. A fixed node is selected as a receiver and the rest as senders. The receiver listens continuously like a gateway node [6], [15] and does not enter into the sleep state. The senders are duty cycled and their sleep durations are varied for different experiments. Sleep duration is how long a node stays in the sleep state before waking up. All results in this section are averages of three runs.

Figure 3 shows the average number of synchronizations per node in one-hop networks during 60 seconds of simulation time. Since all nodes are simulated for the same number of clock cycles ($60 \times$ clock frequency of ATmega128L) in all test cases, the average number of synchronizations per node is a good indicator to the performance of the speedup technique. The synchronization numbers are collected by logging code we add specifically

¹The upper bound is imposed by the timer implementation of TinyOS 1.1. The timer code sets *maxTimerInterval* to 230ms and the physical timers of a real sensor node can not be set to anything larger than that using the timer API.

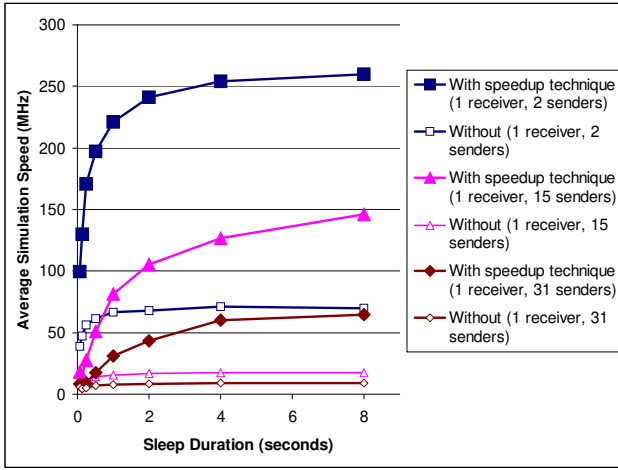


Fig. 5. Average simulation speed in one-hop networks

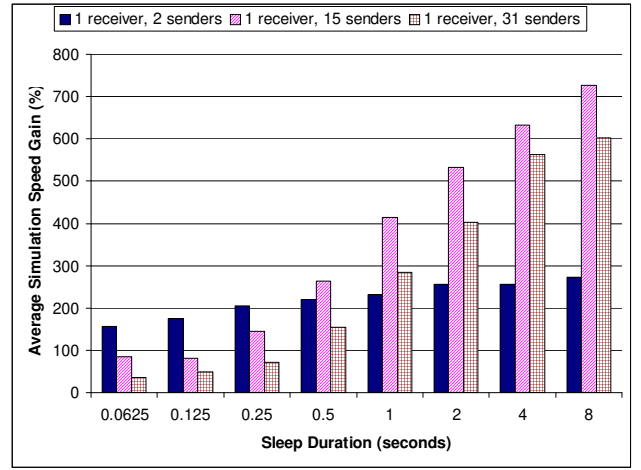


Fig. 6. Percentage increases of average simulation speed in one-hop networks

for evaluation purposes. Figure 4 shows the percentage reductions of the average number of synchronizations per node in one-hop networks during the 60 seconds of simulation time. Figure 5 shows average simulation speed in one-hop networks. The average simulation speed V_{avg} is calculated using Equation 1. The percentage increases of average simulation speed in one-hop networks are shown in Figure 6.

$$V_{avg} = \frac{\text{total number of clock cycles executed by the sensor nodes}}{(\text{execution time of the simulation}) \times (\text{number of sensor nodes in the simulation})} \quad (1)$$

As shown in Figure 3 and Figure 4, the speedup technique significantly reduces synchronizations in all the test cases and the largest percentage reduction is more than 99%. The reduction percentages increase with sleep durations under fixed network sizes except for the 16 (1 receiver, 15 senders) and 32 (1 receiver and 31 sender) node test cases that use $62.5ms$ sleep durations. The unusually high percentage reductions in those cases are results of using the CC1000 CSMA MAC protocol of TinyOS 1.1. When multiple senders in communication range transmit at the same time, the MAC protocol would sequence their transmission time using random backoffs. Since the senders will not return back to the sleep state until packets are successfully transmitted, the sleep times of the senders are sequenced as well. This effectively reduces synchronizations in simulations as the number of nodes that are active at a same time is reduced. The 3-node (1 receiver, 2 senders) test case is not affected by this because we randomly delay the starting time of each node between 0 and 1 second in all our experiments to prevent the nodes from artificially starting at the same time. Therefore, the chance that the two senders in the 3-node test case to transmit at the same time is very low. As the number of nodes in one-hop networks grows, the chance for concurrent transmissions increases and the number of synchronizations decreases in cases like these. The number of synchronizations also decreases with sleep durations in such cases. This is because nodes transmit more frequently and the chance for concurrent transmissions increases. We can see this from Figure 7, a zoomed in view of Figure 3. When sleep duration doubles from $62.5ms$ to $125ms$, the average number of synchronizations per node actually increases for both 16 and 32 node test cases, regardless of whether the speedup technique is used or not. We can also see in the same test cases that the average number of synchronizations per node decreases with network size under fixed sleep durations. The speedup technique can further reduce synchronizations in cases like these because it increases the lookahead time of simulated nodes. With larger lookahead time, running nodes can advance further without synchronizing with nodes in the sleep state. The reductions from applying the speedup technique are greater for larger one-hop networks in those cases as there are more of this type of sequencing in larger one-hop networks under fixed sleep durations.

As shown in Figure 5 and Figure 6, the speedup technique significantly increases average simulation speed in all the test cases and the largest increase is more than 700%. Although the 3-node speedup test case has the highest




Fig. 7. Average number of synchronizations per node in one-hop networks during 60 seconds of simulation time (a zoomed in view of Figure 3)

percentage reduction of the average number synchronizations per node as shown in Figure 4, it does not have the largest average simulation speed increase in Figure 6. This is because the overhead in performing individual synchronizations is very low for 3-node test cases. Context switches are generally not needed for synchronizations in those cases because there are more processors (4) than nodes/threads (3). We can also see in Figure 6 that the growth of the average simulation speed quickly flattens out for large sleep durations in original Avrora but continues after applying the speedup technique. We could expect even better average simulation speed in simulating large one-hop networks after optimizing our generic distributed synchronization algorithm for parallel simulations as discussed in Section III-A.

The speedup technique can not completely eliminate synchronizations caused by having only limited numbers of physical processors available for simulations. We can see this from Figure 3 and Figure 7. When the sleep duration is long enough, the average number of synchronization per node with speedup is similar for all network sizes.

B. Performance in multi-hop networks

In this section, we evaluate the performance of the speedup technique using multi-hop sensor networks. Nodes are laid 20 meters apart on square grids of various sizes. Sender and receivers are positioned on the grids in such a way that nodes of the same types are not adjacent to each other. By setting a maximum transmission range of 20 meters, this setup ensures that only neighboring nodes are within direct communication range of each other. This configuration is very similar to the two dimensional topology in DiSenS [5]. Once again, only senders are duty cycled to keep the experiments simple.

Figure 8 shows the average number of synchronizations per node in multi-hop networks during 20 seconds of simulation time. The percentage reductions of the average number of synchronizations per node in multi-hop networks during the 20 seconds of simulation time is shown in Figure 9. We can see that there are significant reductions in the average number of synchronizations per node in all the test cases using the speedup technique and the reduction percentages scale with sleep durations.

Figure 10 and Figure 11 indicate that the speedup technique significantly increases average simulation speed in all multi-hop test cases. Compared to the one-hop test results in Figure 6, the speed increases scale better with network sizes in multi-hop tests. This is because our distributed synchronization algorithm has less overhead on sensor networks that have smaller numbers of nodes within direct communication range as described in the end of Section III-A.

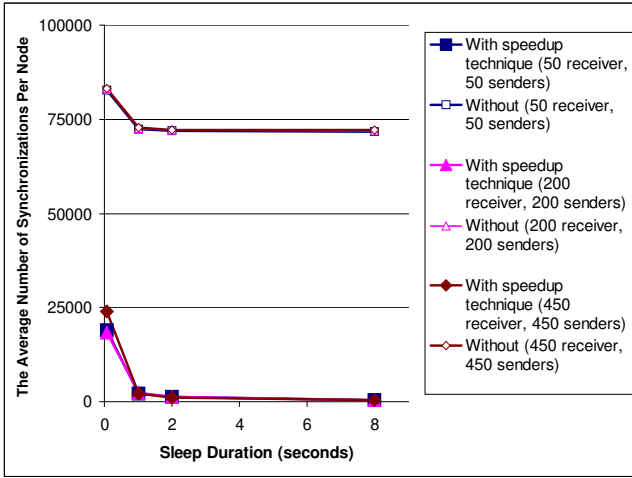


Fig. 8. Average number of synchronizations per node in multi-hop networks during 20 seconds of simulation time

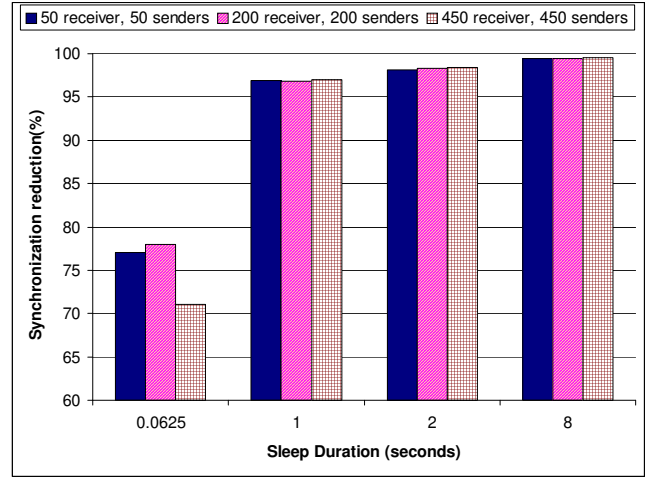


Fig. 9. Percentage reductions of the average number of synchronizations per node in multi-hop networks during 20 seconds of simulation time

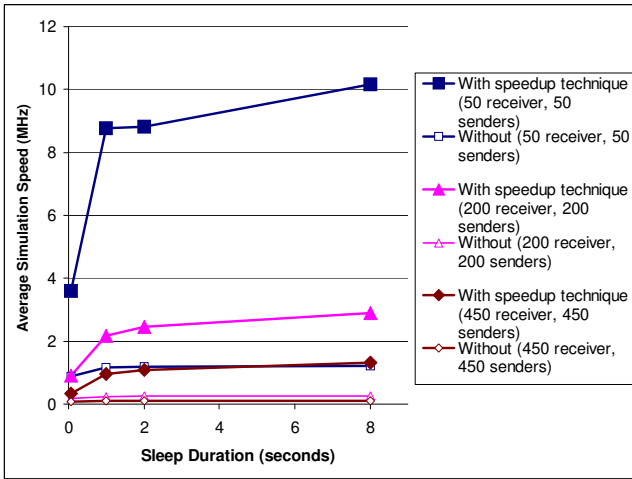


Fig. 10. Average simulation speed in multi-hop networks

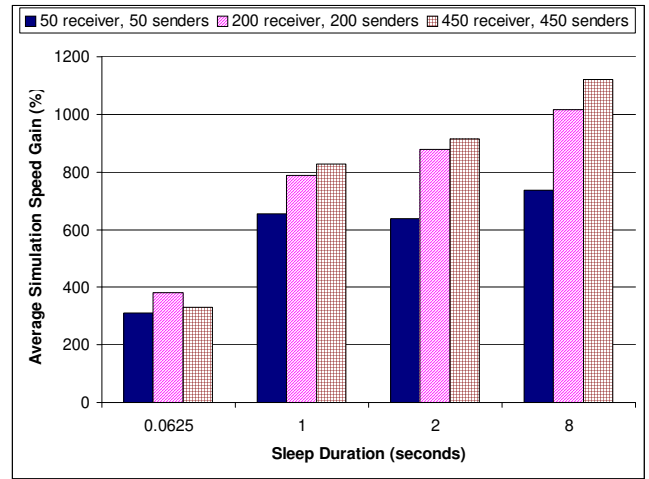


Fig. 11. Percentage increases of average simulation speed in multi-hop networks

V. RELATED WORK

Sensor network simulators can be broadly divided into time-stepped simulators and discrete event driven simulators (DESS) based on the time flow mechanism used to advance simulation time [16]. Most sensor network simulators [17], [18], [19], [4], [20], [5] are DESSs because event driven simulations are usually more efficient than time-stepped ones.

Atemu [7] is the first simulator that supports cycle accurate simulations of sensor networks. It is based on time-stepped simulation technique and has limited scalability. Avrora [4] is designed to be more scalable than Atemu while providing a similar level of simulation fidelity. To increase scalability, Avrora employs distributed event driven simulation techniques and relaxes the radio interface from bit level to byte level. However, the synchronization algorithm of Avrora only supports parallel simulation over SMP computers. DiSenS [5] is the first truly distributed discrete event driven sensor network simulator that supports cycle accurate simulations. It is more scalable than Avrora because it can be used to run a simulation over a network of computers. Both Atemu and Avrora are relatively mature simulators and have their source code publicly available.

There is a large body of work on improving the scalability of distributed discrete event driven simulators in

general. This work can be classified into two groups, those based on conservative synchronization algorithms [2] and those based on optimistic synchronization algorithms [3]. The performance of conservative approaches is bounded by worse case scenarios. The optimistic approaches do not have this limitation but they are usually very complex for implementation and require a large amount of memory to run.

Exploiting lookahead time is a very common conservative approach to improve the scalability of distributed discrete event driven simulators [21], [22]. Our approach is similar to those in the sense that we also improve scalability of distributed discrete event driven simulators by increasing lookahead time. However, our technique is fundamentally different as we use different and application specific characteristics in a different context to increase lookahead time.

VI. CONCLUSION AND FUTURE WORK

We have described a speedup technique that significantly reduces sensor node synchronizations in distributed simulations of sensor networks and consequently improves average simulation speed and scalability of distributed sensor network simulators. We implemented this technique in Avrora, a widely used parallel sensor network simulator and conducted extensive experiments. The significant performance improvements with parallel simulations suggest even greater benefits in applying our technique to distributed simulations over a network of computers because of their large overheads in sending synchronization messages across computers during simulations.

As future work, we plan to merge our implementation into the latest development branch of Avrora. This would make it possible to simulate TinyOS 2.0 based applications with our speedup technique. We also plan to support an optimization that can reduce communication overheads in distributed simulations based on the speedup technique. When a sensor node is in the sleep state, its radio is off and it will not access the wireless channel at all. In other words, when a transmitting node knows that a receiving node is in the sleep state during a simulation, it no longer needs to send packets to the receiver for the entire sleep period. If the sender and receiver are simulated on different computers, the savings in terms of communication time and network bandwidth consumptions could be significant.

REFERENCES

- [1] R. M. Fujimoto, "Parallel and distributed simulation," in *WSC '99: Proceedings of the 31st conference on Winter simulation*. New York, NY, USA: ACM, 1999, pp. 122–131.
- [2] K. M. Chandy and J. Misra, "Asynchronous distributed simulation via a sequence of parallel computations," *Commun. ACM*, vol. 24, no. 4, pp. 198–206, 1981.
- [3] D. R. Jefferson, "Virtual time," *ACM Trans. Program. Lang. Syst.*, vol. 7, no. 3, pp. 404–425, 1985.
- [4] B. L. Titzer, D. K. Lee, and J. Palsberg, "Avrora: scalable sensor network simulation with precise timing," in *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*. Piscataway, NJ, USA: IEEE Press, 2005, p. 67.
- [5] Y. Wen, R. Wolski, and G. Moore, "Disens: scalable distributed sensor network simulation," in *PPoPP '07: Proceedings of the 12th ACM SIGPLAN symposium on Principles and practice of parallel programming*. New York, NY, USA: ACM Press, 2007, pp. 24–34.
- [6] R. Szewczyk, J. Polastre, A. M. Mainwaring, and D. E. Culler, "Lessons from a sensor network expedition," in *EWSN*, 2004, pp. 307–322.
- [7] J. Polley, D. Blazakis, J. McGee, D. Rusk, and J. Baras, "Atemu: a fine-grained sensor network simulator," in *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, 4-7 Oct. 2004, pp. 145–152.
- [8] *MICA2 Datasheet*, Crossbow, 2008.
- [9] *ATMega128L Datasheet*, Atmel, 2003.
- [10] D. M. Nicol, "Scalability, locality, partitioning and synchronization pdes," *SIGSIM Simul. Dig.*, vol. 28, no. 1, pp. 5–11, 1998.
- [11] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister, "System architecture directions for networked sensors," *SIGPLAN Not.*, vol. 35, no. 11, pp. 93–104, 2000.
- [12] "Tinyos 1.1.15." [Online]. Available: <http://www.tinyos.net/>
- [13] P. Levis, D. Gay, V. Handziski, J.-H. Hauer, B. Greenstein, M. Turon, J. Hui, K. Klues, C. Sharp, R. Szewczyk, J. Polastre, P. Buonadonna, L. Nachman, G. Tolle, D. Culler, and A. Wolisz, "T2: A second generation os for embedded sensor networks," Telecommunication Networks Group, Technische Universität Berlin, Tech. Rep. TKN-05-007, Nov. 2005.
- [14] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*. New York, NY, USA: ACM, 2004, pp. 95–107.
- [15] Z. Jin, C. Schurgers, and R. Gupta, "An embedded platform with duty-cycled radio and processing subsystems for wireless sensor networks," in *International Workshop on Systems, Architectures, Modeling, and Simulation (SAMOS)*, 2007.
- [16] R. M. Fujimoto, *Parallel and Distribution Simulation Systems*. New York, NY, USA: John Wiley & Sons, Inc., 1999.
- [17] S. Park, A. Savvides, and M. B. Srivastava, "Sensorsim: a simulation framework for sensor networks," in *MSWIM '00: Proceedings of the 3rd ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems*. New York, NY, USA: ACM Press, 2000, pp. 104–111.

- [18] P. Levis, N. Lee, M. Welsh, and D. Culler, "Tossim: accurate and scalable simulation of entire tinyos applications," in *SensSys '03: Proceedings of the 1st international conference on Embedded networked sensor systems*. New York, NY, USA: ACM Press, 2003, pp. 126–137.
- [19] L. Girod, J. Elson, A. Cerpa, T. Stathopoulos, N. Ramanathan, and D. Estrin, "Emstar: a software environment for developing and deploying wireless sensor networks," in *Proceedings of the 2004 USENIX Technical Conference*, Boston, MA, 2004.
- [20] F. Osterlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt, "Cross-level sensor network simulation with cooja," in *Local Computer Networks, Proceedings 2006 31st IEEE Conference on*, Nov. 2006, pp. 641–648.
- [21] D. Filo, D. C. Ku, and G. D. Micheli, "Optimizing the control-unit through the resynchronization of operations," *Integr. VLSI J.*, vol. 13, no. 3, pp. 231–258, 1992.
- [22] J. Liu and D. M. Nicol, "Lookahead revisited in wireless network simulations," in *PADS '02: Proceedings of the sixteenth workshop on Parallel and distributed simulation*. Washington, DC, USA: IEEE Computer Society, 2002, pp. 79–88.