Improving the Speed and Scalability of Distributed Simulations of Sensor Networks

Zhong-Yi Jin, Rajesh Gupta
Dept. of Computer Science & Eng, UCSD
{zhjin, rgupta}@cs.ucsd.edu

January 12, 2009

Abstract

Distributed simulation techniques are commonly used to improve the speed and scalability of simulators for wireless sensor networks. However, accurate simulations of dynamic interactions of sensor network applications incur large synchronization overheads and severely limit the performance of existing distributed simulators. In this paper, we present two novel techniques that significantly reduce such overheads by minimizing the number of sensor node synchronizations during simulations. These techniques work by exploiting radio and MAC specific characteristics without reducing simulation accuracy. In addition, we present a new mechanism that makes it possible to exploit any potential application specific characteristics for synchronization reductions. We implement and evaluate these techniques in a cycle accurate distributed simulation framework that we developed based on Avrora, a popular parallel sensor network simulator. In our experiments, the techniques achieve a speedup from 2 to 3 times in our simulator and 3 to 6 times compared to Avrora in simulating 1-hop networks with 32 to 256 nodes. In our multi-hop flooding tests, they achieve a speedup of 1.5 to 1.8 times in our simulator and 3.9 to 5.6 times compared to Avrora. The experiments also demonstrate that the speedups can be significantly larger as the techniques scale with sensor network sizes and radio off times.

1 Introduction

Wireless sensor network (WSN) simulators are important in developing and debugging WSN applications. By running sensor network programs on top of simulated sensor nodes inside simulated environments, the states and interactions of sensor network programs can be inspected and studied easily and repeatedly. In addition, the properties of the simulated entities (simulation models) such as the locations of sensor nodes and the inputs to the sensor nodes can be readily changed before or during simulations. By choosing appropriate simulation mod-

els, one can build an entire WSN application, including the underlying operating system, using simulations.

There are two key requirements to WSN simulators: fidelity and speed. Fidelity reflects bit and temporal accuracy of events and actions. High fidelity often leads to low simulation speed, defined as the ratio of simulation time to wallclock time. Simulation time is the virtual clock time in the simulated models [6] while wallclock time corresponds to the actual physical time used in running the simulation program. A simulation speed of 1 indicates that the simulated sensor nodes advance at the same rate as real sensor nodes and this type of simulation is called real time simulation. In general, real time speed is required to use simulations for interactive tasks such as debugging or testing.

The fidelity of WSN simulators is rapidly increasing with the use of high fidelity simulation models [12, 17, 16, 10]. However, most of these improvements of fidelity are largely achieved at the cost of decreased simulation speed and scalability because significant computational resources are required to execute high fidelity simulation models. For example, even with TOSSIM [12], a popular and one of the fastest sequential WSN simulators, one can only simulate, without cycle accuracy, about 32 nodes in real time on fast computers [12, 19]. As alternatives, parallel and distributed WSN simulators [19, 20, 3] leverage the combined computing resources of multiple processors/cores on the 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 causality of events and ensure correct simulation results.

Synchronization of two sensor nodes is illustrated in Figure 1 which shows the progress of simulating in parallel two sensor nodes that are within direct communication range of each other. After starting the simulation for T_{W0} seconds of wallclock time, Node A advances to simulation time T_{S2} while Node B only advances to T_{S1}. The speed of simulating the two nodes could be different because the threads or processes used to simulate the nodes might either run on processors of different speeds or receive different number of 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 paths based on whether there are active wireless transmissions or not. However, at T_{W0}, Node A may not be simulated any further than T_{S2} because Node A does not know, at T_{W0}, whether Node B is going to transmit at T_{S2} or not (since $T_{S1} < T_{S2}$). In other words, at T_{S2}, the input of Node A depends on the output of Node B.

To maintain such dependencies, simulated sensor nodes often need to synchronize with each other during simulations. There are two general approaches to handle synchronizations: conservative [2] or optimistic [8]. Conservative synchronization requires that Node A wait at T_{S2} until the simulation time of Node B reaches T_{S2}. The optimistic 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

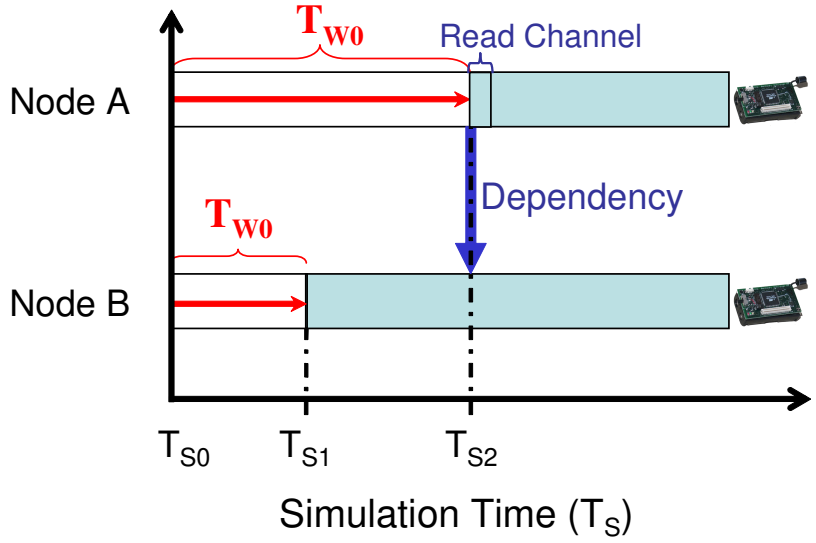

Figure 1: The progress of simulating in parallel a wireless sensor network with two nodes that are in direct communication range of each other

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. To the best of our knowledge, almost all distributed WSN simulators are based on the conservative approach as it is simpler to implement and has a lower memory footprint.

Synchronizations bring significant overheads to distributed simulations. With the conservative approach, the overheads can be divided into management overheads and communication overheads. Management overheads come from managing the threads or processes that simulate sensor nodes. For example, to maximize parallel use of computational resources, the thread or process simulating Node A in Figure 1 needs to be suspended while waiting for Node B so another thread or process simulating another node can be swapped in for execution. Suspended nodes also need to be swapped back in for simulation later on. These usually involve context switches and large numbers of them would significantly reduce simulating speed and scalability. Communication overheads arise because nodes need to communicate their progresses to each other during simulations. For example, in Figure 1, Node B must notify Node A after it advances past T_{S1} so that Node A can continue. Communicating across processes or threads is generally expensive. In the case where nodes are simulated on different computers, the communication overheads could be very high as messages have to be sent across slow networks.

The performance gains of existing distributed WSN simulators are often compromised by the rising overheads due to inter-node synchronizations. For example, with Avrora [19], a cycle accurate parallel WSN simulator, it is faster

to simulate 32 nodes with 1 processor than using all 8 processors of a parallel computer in some of our tests (Figure 4 and 12) because of large overheads in synchronizing threads (1 node/thread) across processors. In the case of DiSenS [20], a cycle accurate distributed WSN simulator, if all nodes are within communication range, DiSenS needs 16 computers to simulate 16 nodes in real time despite of the fact that each of the computers can simulate 8 nodes in real time [20]. 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 propose two novel techniques that greatly reduce synchronization overheads by cutting the number of synchronizations. The techniques work by exploiting radio and MAC specific characteristics without reducing simulation fidelity. They are effective even when node processors or radios are active and can significantly improve simulation speed and scalability. We also present a new mechanism that enables exploiting any potential application specific characteristics for synchronization reductions. We validate our approaches by their implementations in PolarLite, a cycle accurate distributed simulation framework that builds upon Avrora and serves as the underlying simulation engine [9].

We discuss related work in Section 2. The speedup techniques are presented in Section 3 and their implementations are described in Section 4. In Section 5 we present the results of our experiments followed by our conclusion and future work in Section 6.

2 Related work

Previous work in improving the speed and scalability of WSN simulators can be broadly divided into two categories. The first category focuses on reducing the computational demands of individual simulation model without significantly reducing fidelity. For example, it is very computationally expensive to emulate an actual sensor node processor in a simulation model for cycle accurate simulations [16]. To reduce the large computational needs, TimeTossim [10] automatically instruments applications at source code level with cycle counts and compiles the instrumented code into the native instructions of simulation computers for fast executions. This significantly increases simulation speed while achieving a cycle accuracy of up to 99%. However, maintaining cycle counts also slows down TimeTossim to about 1/10 the speed of TOSSIM (Section 1) which TimeTossim is based on. Our work makes this type of effort scalable on multiple processors/cores.

The second category of work focuses on reducing overheads in parallel and distributed simulations. DiSenS reduces the overheads of synchronizing nodes across computers by using the sensor network topology information to partition nodes into groups that do not communicate frequently and simulating each group on a separate computer [20]. However, this technique only works well if most of the nodes are not within direct communication range as described in the paper. In [9], we describe a technique that uses sensor node sleep time to

reduce the number of synchronizations. As demonstrated in the paper, using the node-sleep-time based technique can significantly increase speed and scalability of distributed WSN simulators. However, the technique is only able to exploit the time when both the processor and radio of a sensor node are off for speedup because it is not possible to predict the exact radio wakeup time when the processor is running. The techniques we propose in this paper are orthogonal to the previous techniques and work even when all the nodes of a WSN are within direct communication range or the processors or radios of the nodes are active.

There is a large body of work on improving the speed and scalability of distributed discrete event driven simulators in general. Among them, exploiting lookahead time is a commonly used conservative approach [5, 13]. Our techniques belong to this category in the sense that we also improve speed and scalability by increasing the lookahead time. However, our techniques are fundamentally different as we use different and application specific characteristics in a different context to increase lookahead time.

3 Reducing synchronizations in distributed simulations of WSNs

Sensor node synchronizations are required for enforcing dependencies between simulated sensor nodes due to their interactions over wireless channels. Our approach to increase simulation speed and scalability is based on identifying and exploiting parallelism by reducing synchronizations between nodes according to their communication capabilities. A node can "lookahead" for a minimum guaranteed interval to identify periods when no transmitting or receiving events in a discrete event driven simulator would occur. Such a lookahead is often closely tied to the event scheduling algorithm used in the simulators. Based on the distributed scheduling algorithm in [9], we present two techniques for reducing the number of synchronization events at the radio-level or at the MAC-level.

3.1 Radio-level Speedup Technique

Radio-level duty cycling works by selectively turning radios on and off. Since radios are one of the most power consuming components of sensor nodes, radio-level duty cycling is ubiquitously used in WSNs to reduce energy consumptions and extend working life of energy constrained sensor nodes. Due to its wide applications and complex tradeoffs in energy savings and communication overheads, radio-level duty cycling is commonly built into energy efficient WSN MACs such as S-MAC [21] and B-MAC [14].

Our radio-level speedup technique is illustrated in Figure 2 which shows the progress of simulating two sensor nodes in parallel. In this simulation, Node B turns its radio off at time T_{S1} and puts it back on at time T_{Sx}. With existing distributed simulators, after running the simulation for T_{W0} seconds of wallclock time, Node A has to wait at T_{S3} for Node B to catch up from T_{S2} despite of

the fact that node B will not transmit any packets at T_{S3}. Ideally, we can avoid this unnecessary synchronization by having Node B notify node A at time T_{S1} that its radio is off until T_{Sx}. However, this will not work as it is not possible for Node B to predict the exact radio wakeup time T_{Sx} at T_{S1}. This is because while the radio is off at T_{S1}, the sensor node processor is still running and it can turn the radio back on at any time based on current states, application logics and sensor readings. In other words, it is just not possible for Node B to predict when the radio will be turned on in the future.

Instead of predicting the exact radio wakeup time, our radio-level speedup technique exploits the radio off period by calculating the earliest possible communication time, $T_{EarliestCom}$. $T_{EarliestCom}$ is the earliest time that a turned off radio can be used to send or receive data over wireless channels and can be calculated based on T_{Act}, the amount of time to fully activate a turned off radio. A turned off radio can not be activated instantly for sending or receiving data. It takes time for the radio to be initialized and become fully functional [15]. For example, the CC1000 radio of Mica2 nodes [4] needs 2.45ms to be activated and the CC2420 radio of Telos nodes [15] needs about 1.66ms without counting the SPI acquisition time [11]. The exact delays in terms of numbers of clock cycles are hard coded into WSN MAC protocols and can be easily identified in the source code. For example, in TinyOS 1.1 [7, 18], B-MAC waits for a total of 34300 clock cycles for the CC1000 radio of MICA2 by calling the `TOSH_uwait` function. While the delays seem to be small, they are significantly larger than typical lookahead times in simulating WSNs. For example, it is about 11 times larger than the 3072 clock cycle lookahead time in simulating Mica2 nodes [19, 9]. As mentioned, lookahead time is defined as the maximum amount of simulation time that a simulated sensor node can advance freely without synchronizing with other simulated sensor nodes [9].

Our radio-level speedup technique works by tracking when sensor node radios are turned on and off. When we detect that a sensor node radio is turned off, we immediately send its $T_{EarliestCom}$ in a clock synchronization message to all neighboring nodes and then repeatedly send the latest $T_{EarliestCom}$ every T_{Act} time until the radio is detected to be turned on. $T_{EarliestCom}$ can be calculated as the sum of current simulation time and T_{Act}. As a result, neighboring nodes no longer need to synchronize with the radio-off node until the latest $T_{EarliestCom}$. For example, as shown in Figure 2, when we detect that Node B turns its radio off at T_{S1}, we immediately send its $T_{EarliestCom}$ to Node A and repeat that every T_{Act} time which is fixed according to the radio of Node B. The clock synchronization messages are shown as arrows from Node B to Node A in the figure. Upon receiving the second $T_{EarliestCom}$, the lookahead of node A increases to a time beyond T_{S3} and therefore it no longer needs to wait at T_{S3} after T_{W0} seconds of simulations. In other words, Node A knows before T_{W0} that Node B will not be able to transmit any packet at T_{S3}. Since the increase of lookahead time ($T_{Act} - OldLookAheadTime$) may just be a very small fraction of the total radio off period, it is critical to repeatedly send $T_{EarliestCom}$ every T_{Act} time to fully exploit the entire radio off period.

The radio-level speedup technique also reduces the number of clock synchro-

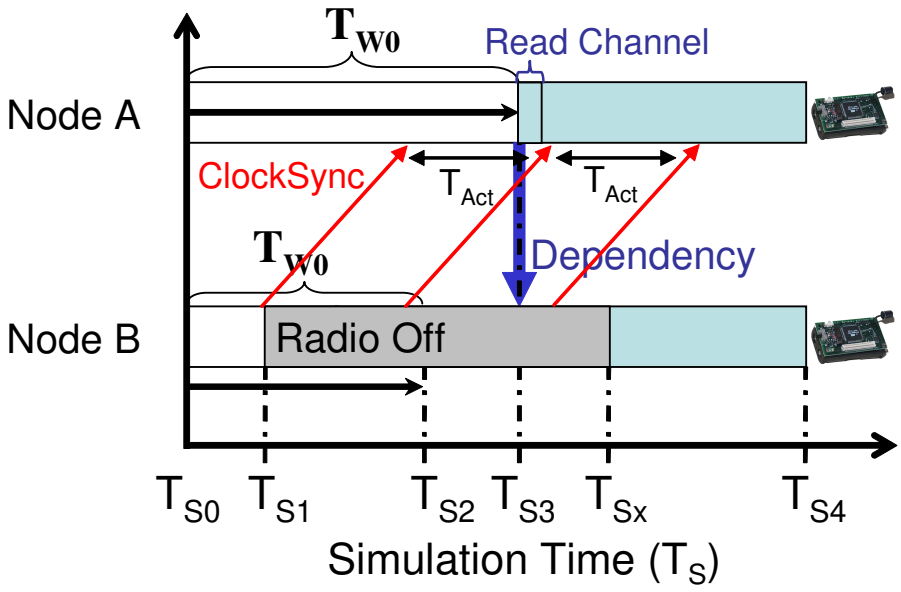

Figure 2: The progress of simulating two nodes that are in direct communication range with the radio-level speedup technique

nizations. In distributed simulations, clock synchronization messages are used to send the simulation time of a node to all its neighboring nodes so causality can be maintained and suspended waiting nodes can be revived. To maximize parallelism in simulations, a node needs to send 1 clock synchronization message for every lookahead time of its neighboring nodes [19, 20, 9]. Since our radio-level speedup technique increases the lookahead time of neighboring nodes, the number of clock synchronizations can be greatly reduced. For example, in the case of simulating Mica2 nodes, one $T_{EarliestCom}$ message can increase lookahead time by a factor of 11 and therefore eliminates 10 clock synchronization messages.

3.2 MAC-level speedup technique

While the radio-level speedup technique takes advantage of physical delays in WSN radios, our MAC-level speedup technique exploits the random backoff behaviors of WSN MACs. Almost all WSN MACs need to perform random backoffs to avoid concurrent transmissions [21, 14]. For example, before transmitting a packet, B-MAC would first perform an initial backoff. If the channel is not clear after the initial backoff, B-MAC needs to repeatedly perform congestion backoffs until the channel is clear. Because a MAC will not transmit any data during backoff periods, we are able to exploit the backoff times for speedups. Although the backoff times are random and MAC specific, they are usually a lot

longer than typical lookahead times in simulating WSNs. For example, in the case of B-MAC, the default initial backoff is 1 to 32 times longer than the 3072 clock cycle lookahead time in simulating Mica2 nodes. The default congestion backoff is 1 to 16 times longer in B-MAC

Our MAC-level speedup technique is illustrated in Figure 3 which is similar to Figure 2 except Node B enters into a backoff period from T_{S1} to T_{S4}. The MAC-level speedup technique works by detecting the start and the duration of a backoff period. When the start of a backoff period is identified, the end time of the backoff period is first calculated based on the duration of the period and then sent to the neighboring nodes. This effectively increases the lookahead time of neighboring nodes and helps to eliminate unnecessary synchronizations. For example, in order to avoid the unnecessary synchronization of Node A at T_{S3} after running the simulation for T_{W0} seconds of wallclock time, our MAC-level speedup technique first detects at T_{S1} the start of Node B's backoff period as well as the duration of the backoff period. Then we compute the end time of the backoff period and send that in a clock synchronization message to Node A. Once Node A knows that Node B will not transmit until T_{S4}, it no longer needs to wait at T_{S3}. Similar to the radio-level speedup technique, the MAC-level technique also reduces the number of clock synchronizations and provides additional speedup. We discuss how to detect the start and the duration of a random backoff period in Section 4.

The MAC-level speedup technique is a good complement to the radio-level technique as WSNs usually have very bursty traffic loads. Nodes in a WSN usually do not communicate frequently and can duty cycle their radios extensively until certain triggering events occur. Once trigged by those events, sensor nodes need to actively communicate and interact with each other to accomplish certain tasks. Our MAC-level technique is most effective when wireless channels are busy.

4 Implementation

The proposed speedup techniques are implemented in PolarLite, a distributed simulation framework that we developed based on Avrora [9]. Our simulation framework provides the same level of cycle accurate simulations as Avrora but uses a distributed synchronization engine instead of Avrora's centralized one. This makes it possible to implement and evaluate our speedup techniques in a distributed simulation environment. The synchronization engine of PolarLite is based on a distributed synchronization algorithm described in [9]. With this algorithm, nodes can be synchronized separately according to their own lookahead time. In other words, if a node is not accessing the wireless channel, it only needs to communicate to neighboring nodes its simulation time and does not need to wait for any other nodes.

As with Avrora, PolarLite allocates one thread for each simulated node and relies on the Java virtual machine (JVM) to assign runnable threads to any available processors on an SMP computer. We have improved the initial version

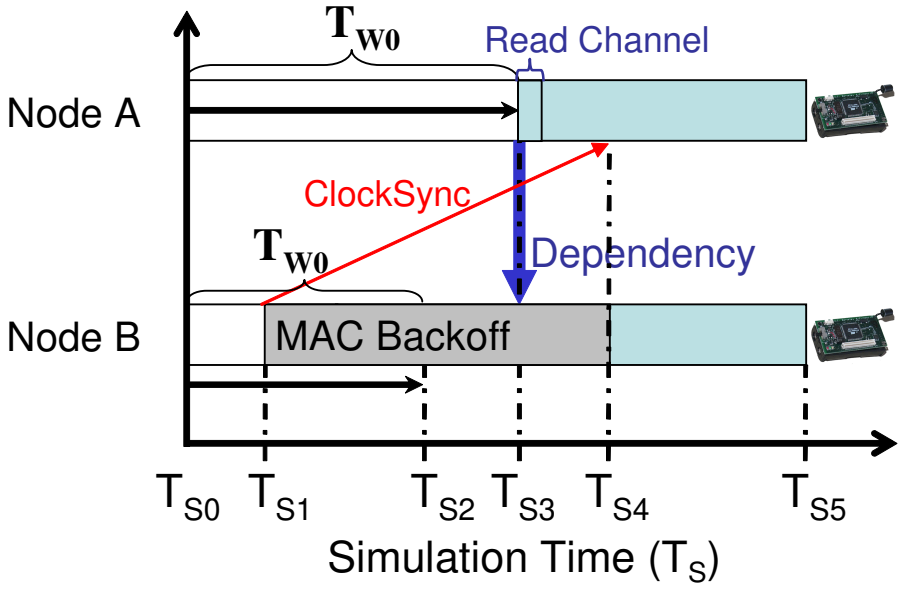

Figure 3: The progress of simulating two nodes that are in direct communication range with the MAC-level speedup technique

of our distributed synchronization engine with sorted clock lists to reduce the overheads in computing minimum channel clocks (*minclock*) [9] when a large number of nodes are in direct communication range.

To implement the radio-level speedup technique, we need to detect when radios are turned on and off. In discrete event driven simulations, the changes of radio states are triggered by events and can be tracked. For example, in our framework, we detect the radio on/off time by tracking the IO events that access the registers of simulated radios.

Detecting MAC backoff times and durations for the MAC-level speedup technique are considerately more difficult. The backoffs are MAC and application specific and generally do not correlate to any unique events or actions that can be easily tracked. In addition, the backoff durations are completely random. One possible solution to this problem is to insert special code into the source code of an application that is to be simulated. These special pieces of code are compiled into the application and used to report to the simulator the MAC backoff times and durations during simulations. However, this technique will not work for cycle accurate simulators like ours.

Cycle accurate sensor network simulators [16, 19, 20, 9] offer the highest level of fidelity among all types of sensor network simulators. They provide simulation models that emulate the functions of major hardware components of a sensor node, mainly the processor. Therefore, they take as input the same binary code (images) that are executed by real sensor nodes. To detect backoff

times and durations without changing the source code of the applications under simulations, we develop a generic mechanism based on pattern matching to expose the internal states of sensor network applications during simulations.

Our mechanism works by first using patterns to pin point from compiled applications the machine instructions that represent events of interest during the start of a simulation. The identified instructions are then replaced by corresponding "hook" instructions to report the internal states of the applications, such as the backoff durations, to a simulator during simulations. Hook instructions are artificial instructions that behave exactly the same as the original instructions they replace except they will pass to a simulator the memory locations or registers accessed by the original instructions during simulations. The values stored in those locations are the internal states of the applications that correspond to the events of interest. Since an instruction may access multiple locations, we associate with each pattern a list that indicates the operands of interest based on their order in the instruction. As a result, an instruction may be translated into different hook instructions according to the list. Note that the values stored in the reported memory locations may not correspond to lookahead time directly. In our current implementation, we simply assume the values are in units of 3072 clock cycles, the amount of time to transmit 1 byte with Mica2 nodes. (As future work, we are planning to differentiate that using different types of hook instructions.) To maintain cycle accuracy, our simulator ensures that the hook instructions consume the same number of clock cycles as the original instructions.

We do not use addresses to locate instructions in a compiled application as addresses tend to vary across compilations after source code changes, even if the changes are at places not related to the instructions. With pattern matching, we only need to create a set of patterns once, if the corresponding source code does not change. For example, if an application is written with TinyOS, the instructions that assign backoff durations to B-MAC are part of the OS, regardless of whether the backoffs are calculated by default functions in the OS or user supplied functions in the application. Therefore, we only need to create a set of patterns once for each version of TinyOS to track the backoff times in B-MAC during simulations.

We use regular expressions for pattern matching. To uniquely match an instruction, we need to match additional instructions before or after that instruction as well. We implement the MAC-level speedup technique by hard coding the backoff matching process in our simulator. It matches all three instructions (1 for initial backoff and 2 for congestion backoff) that assign backoff durations to B-MAC and replaces them with appropriate hook instructions. For simplicity, we consider that MAC backoffs start at the times that the hook instructions report the backoff durations. It is safe to do so as no data will be sent from this point on until the end of the backoff periods.

Table 1: Radio off periods under different duty cycling modes of B-MAC

Duty cycling Mode	Radio off Time (ms)
0	0
1	20
2	85
3	135
4	185

5 Evaluation

We conduct a series of experiments to evaluate the performance of our speedup techniques. All experiments are carried out on an SMP server running Linux 2.6.24. The server features a total of 8 cores on 2 Intel Xeon 3.0GHz CPUs and 16GByte of RAM. Sun's Java 1.6.0 is used to run all experiments.

All of the experiments are based on the CountSend (sender) and CountReceive (receiver) programs from the TinyOS 1.1 distribution. These programs are similar in nature to the programs used by other WSN simulators in evaluating their performance [12, 19, 20]. CountSend repeatedly broadcasts a continuously increasing counter at a fixed interval. CountReceive simply listens for messages sent by CountSend and displays the counters inside on LEDs. The programs are executed on simulated Mica2 nodes [4] and the starting time of each node is randomly perturbed between 0 and 1 second of simulation time to avoid any artificial time locks. All simulations are run for 300 seconds of simulation time and for each experiment we take average of three runs as the results.

5.1 Performance of radio-level speedup technique

For experiments in this section, we modify the sender and receiver programs slightly to enable B-MAC's built-in radio level duty cycling feature. This can be done by calling the SetListeningMode and SetTransmitMode functions of TinyOS 1.1 at start. B-MAC supports a total of 5 radio-level duty cycling modes in TinyOS 1.1 as shown in Table 1. Once enabled, B-MAC turns a radio off periodically for a duration corresponding to the duty cycling mode. The radio is turned back on either when there are data to transmit or a radio off period ends. The radio is turned off again once there are no pending packets to transmit and the channel is clear for a fixed period of time. In the case of TinyOS 1.1 [7, 18] and Mica2 nodes, the channel clear time is the amount of time to transmit 8 bytes over the radio.

5.1.1 Speed and Scalability with respect to the number of processors

Our first set of experiments evaluates the performance of our radio-level speedup technique over different numbers of processors, or cores in this case. For these experiments, we simulate a WSN of 32 nodes that are within direct communication range using 2 to 8 processors. The 32 nodes are set up in such a way

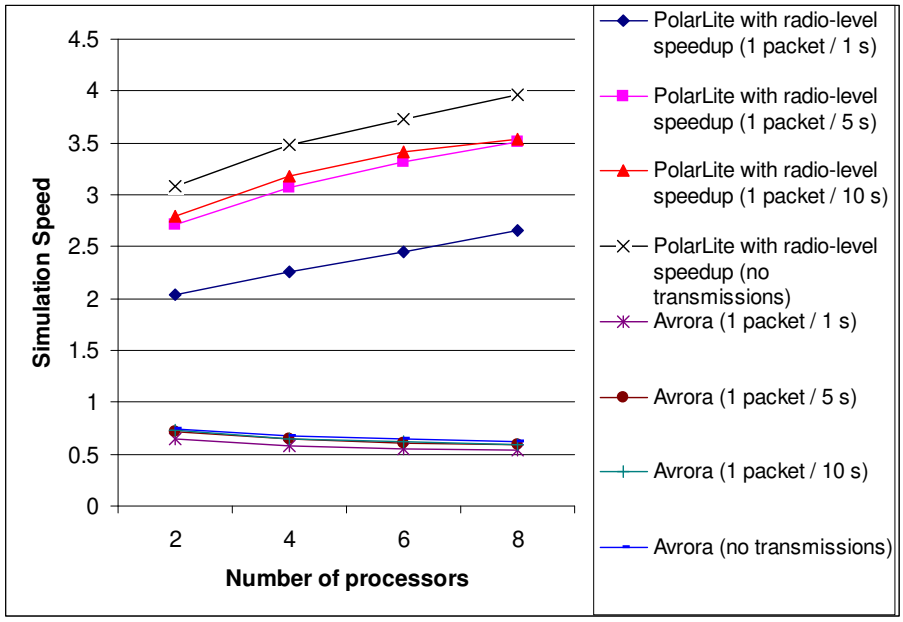

Figure 4: Speed of simulating with Avrora and PolarLite running the radio-level speedup (1 sender 31 receivers, mode 3)

that one node is configured as a sender and the rest as receivers. Since all nodes are within direct communication range, any one of the nodes can be chosen as the sender. The frequency that the sender transmits packets is varied for different experiments. The radio duty cycling mode of all nodes is set to 3 for both transmitting and receiving. For comparisons, we conduct the same experiments using Avrora, PolarLite without any speedups and PolarLite with the radio-level speedup.

As a baseline, Figure 4 shows the speeds (Section 1) of simulating with Avrora and PolarLite running the radio-level speedup technique. We can see that PolarLite running the radio-level speedup technique is considerably faster than Avrora (up to 544% or 6.44 times) and scales with the number of processors. In contrast, the speeds of simulating with Avrora decrease with increasing number of processors in the experiments due to larger synchronization overheads[1]. We can also see that the speeds of simulating with our radio-level speedup technique increase with transmission intervals. This is because our radio-level speedup technique is based on exploiting radio off time and large transmission intervals increase that. In these B-MAC based experiments, the transmission interval affects the radio off time of the senders more than the receivers because a sender has to transmit a preamble longer than the radio off

[1]Our results are different from results of similar expriments in [19] as our experiments use faster 3.0GHz CPUs, compared to 900MHz ones of theirs.

period of the receivers before transmitting a packet. This is to make sure the receivers with duty cycled radios can still receive the packet.

Figure 5 shows the percentage reductions of synchronizations based on numbers collected by running with and without the radio-level speedup technique in PolarLite. For accurate evaluations, we only show the percentage reductions within our PolarLite framework because PolarLite and Avrora are based on different synchronization algorithms and our speedup techniques are only implemented in PolarLite. As shown in Figure 5, the percentage reductions of synchronizations are significant in all cases and actually grow very slowly with the number of processors. This is because more nodes can be simulated in parallel when the number of processors increases and as a result our radio-level speedup technique has more radio sleep time to exploit simultaneously for synchronization reductions. Although the reduction numbers are very close with respect to the number of processors, simulation speeds increase significantly with the number of processors in Figure 6 which shows the speed of simulating with and without the radio-level speedup technique in PolarLite using different number of processors. The reason is that per-synchronization overheads increase with the number of processors due to high inter-processor communication overheads.

As shown in Figure 6, using the radio-level speedup technique increases simulation speeds significantly (up to 111%) in PolarLite. Comparing with Figure 4, we observe that PolarLite alone without any speedup techniques is faster than Avrora in these experiments. This is because our distributed synchronization algorithm (Section 4) can provide more parallelism by allowing nodes to be synchronized separately according to their own lookahead time. Avrora on the other hand synchronizes all nodes together at a fixed time interval. However, as shown in Figure 6, even with our distributed synchronization algorithm, PolarLite alone does not scale well with the number of processors. Using the radio-level speedup technique significantly improves the scalability.

5.1.2 Speed and Scalability with respect to network sizes and radio off times

We also evaluate the radio-level speedup technique over WSNs of different sizes and radio sleep durations (radio-level duty cycling mode). Nodes in the experimenting WSNs are within direct communication range and one sender transmits a packet every 10 seconds to the rest of the receiver nodes. Figure 7, 8, 9 and 10 show the results of simulating with or without the radio-level speedup technique in PolarLite using all 8 processors.

Figure 7 shows significant percentage reductions of synchronizations using our radio-level speedup technique. We do not see any reductions when the radio is on all the time because our radio-level speedup technique works by exploiting radio off time. Figure 7 also shows that the reduction percentages scale with radio off durations as larger durations bring more radio-off time. While the reduction percentages are about the same for all network sizes at a given radio off duration, the percentage increases of simulation speed actually grow with network sizes as shown in Figure 8. This is because in a network where all

nodes are in direct communication range, the total number of synchronizations in a distributed simulation is in the order of $N*(N-1)$ where N is the network size. Therefore, the total number of reductions in the experiments grow with network sizes. This can be seen clearly in Figure 9 which shows the total number of synchronizations in logarithmic scale.

We can also see from Figure 8 that the percentage increases of simulation speed scale very well with radio off durations. The unusually high percentage increase of speed in the case of simulating 16 nodes with a 135ms radio off duration is not a result of high synchronization reductions according to Figure 7. It is caused by the slow speed in simulating that configuration in PolarLite without using the radio-level speedup technique. In fact, without the radio-level speedup technique, we observe no change of simulation speed by increasing the sleep duration from 85ms to 135ms in the 16 node test case.

We also evaluate the scalability of the radio-level speedup technique over larger WSNs under a fixed transmission rate of 1 packet/10 seconds and the results are shown in Figure 10. We can see that the radio-level speedup technique increases simulation speed in large WSNs as well. It provides a 197% increase of simulation speed when simulating 256 nodes in PolarLite. That is an additional 76% improvement over the 119% speed increase in simulating 32 nodes as shown in Figure 8. In other words, although simulation speeds decrease with network sizes due to the limited computational power of our server, the percentage increases of simulation speed using the radio-level speedup technique still grow with network sizes.

5.2 Performance of the MAC-level speedup technique

The performance of our MAC-level speedup technique depends on how busy wireless channels are and how often sensor nodes transmit around the same time. Instead of evaluating with a large number of scenarios, we study the maximum speedup that can be achieved in simulating a WSN with the MAC-level speedup technique. For experiments in this section, we allow CountSend to send as fast as possible by modifying CountSend such that it sends out a new packet as soon as it is notified by the MAC that the previous packet is sent. We also disable the radio-level duty cycling for both CountSend and CountReceive.

We simulate two WSNs that have 1 receiver and 31 or 63 senders using Avrora, PolarLite without speedups and PolarLite with the Mac-level speedup. Unless explicitly specified, the default backoff calculation functions in TinyOS 1.1 are used for the senders. The results are shown in Figure 11, 12, 13 and 14.

5.2.1 Speed and Scalability with respect to the number of processors and backoff times

As shown in Figure 11, our MAC-level speedup technique reduces synchronizations by about 44% to 47% in PolarLite. Although the reduction numbers are about the same with respect to the number of processors, simulation speeds actually increase with the number of processors in Figure 12 which shows the

speeds of simulating with the MAC-level speedup technique using different numbers of processors. Similar to Figure 5 and 6, this is because per-synchronization overheads increase with the number of processors. As shown in Figure 12, the MAC-level speedup technique brings a speedup of 14% to 31% in PolarLite and 96% to 323% compared to Avrora using the default backoff calculation functions of TinyOS 1.1. However, the backoff windows produced by the default backoff calculation functions are not large enough for our experiments as we observe a significant amount of colliding transmissions causing dropped packets. This limits the performance of our MAC-level speedup technique as nodes may transmit at the same time without backoffs. We perform the same experiments by doubling the sizes of the default backoff windows and the results are shown in Figure 13. We can see that our MAC-level speedup technique brings more significant increases of simulation speed with larger backoff windows.

5.2.2 Speed and Scalability with respect to network sizes

Figure 11 also shows that the percentage reductions of synchronizations using the MAC-level speedup technique increase with network sizes in PolarLite. As explained in Section 5.1.2, the total number of synchronizations in these experiments are in the order of the square of the network size. Therefore, the total number of reductions is very significant when the network size doubles from 32 to 64. We can see in Figure 14 that the speeds of simulating with the MAC-level speedup technique scale with network sizes even with the default backoff windows. We notice the unusually low increase of speed in simulating 64 nodes with 6 processors. Since the reductions are consistent as shown in Figure 11, we believe this is caused by the asymmetrical use of all 4 cores of 1 CPU and 2 cores of another CPU in our server. This can also be observed in Figure 12 which shows slow grows of simulation speed when using 6 processors.

5.3 Performance over multi-hop networks with both techniques

We evaluate the combined performance of our speedup techniques with a real world scenario. In this scenario, we simulate a WSN service that floods data to every node in a WSN. This service works by having every node in a WSN relay, by broadcasting, messages it receives. To avoid sending duplicate messages, a node only relays messages with IDs greater than the largest IDs of the messages it has already sent. For experiments in this section, we modify CountReceive to relay messages the way we just described.

In our experiments, we simulate WSNs that have nodes laid 3 meters apart on square grids of different sizes. For each of the WSNs, a corner node is configured as the sender and the rest of nodes are configured as relaying nodes running the modified CountReceive program. The sender transmits a new packet every 20 seconds with an increasing ID. The radio-level duty cycling mode of all nodes is set to 4 (Table 1) and the backoff windows are doubled from TinyOS 1.1 defaults. The transmit range of all nodes is set to 19 meters. The results of our

experiments are shown in 15 and Figure 16.

As shown in Figure 15 and 16, PolarLite running both speedup techniques is significantly faster and provides a speedup of 51% to 75% over PolarLite alone and 289% to 462% compared to Avrora. We can also see from Figure 16 that the speedup techniques reduce synchronizations significantly by 58% to 70%. However, we observe that the reduction percentages decrease with increasing network sizes. This is caused by our simple flooding protocol. As the network size increases, the number of relaying messages grows as well. Although a node does not transmit the same message twice, it can be forced to receive the same message multiple times from different neighboring nodes. In other words, a transmitting node can keep all nodes within its communication range from turning off their radios. This significantly reduces the sleep times of the nodes and lowers the performance of our radio-level speedup technique. However, even under this setup, our speedup techniques still provide significant increases of simulation speed. This is because our MAC-level speedup technique benefits from an increasing number of backoffs in the larger networks. In practice, more advance protocols are usually used to reduce the number of unnecessary relaying messages. Therefore, we expect significant better performance with the speedup techniques in these cases.

.

6 Conclusion and Future work

We have described two speedup techniques that significantly improve the speed and scalability of distributed sensor network simulators by reducing the number of sensor node synchronizations. We implemented these techniques in PolarLite, a cycle accurate distributed simulation framework based on Avrora. The significant improvements of simulation performance on a multi-processor computer in our experiments suggest even greater benefits in applying our techniques to distributed simulations over a network of computers because of their large overheads in sending synchronization messages across computers during simulations.

We have also developed a general mechanism that can expose the internal states of any sensor network applications during simulations. By knowing the internal states of a sensor network application during simulations, one can exploit any application specific characteristics for the increase of lookahead time and as a result, improve simulation speed and scalability.

As future work, we are planning to use the mechanism to exploit scheduled transmission slots in TDMA type of MACs such as S-MAC [21]. With this type of MAC, a node can only send data in scheduled transmission slots. By knowing the time of the slots during simulations, we can identify the periods that a node does not transmit any data and therefore increase the lookahead time. We also plan to merge our implementation with the latest development branch of Avrora. This would make it possible to simulate TinyOS 2.0 [1] based applications with our speedup techniques. Although the techniques are currently implemented in a cycle accurate simulator, they can also be applied to other simulators like

TOSSIM to make them more scalable over multiple processors.

References

[1] Tinyos 2.0. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 320–320, New York, NY, USA, 2005. ACM.

[2] K. M. Chandy and J. Misra. Asynchronous distributed simulation via a sequence of parallel computations. *Commun. ACM*, 24(4):198–206, 1981.

[3] G. Chelius, A. Fraboulet, and E. Fleury. Worldsens: a fast and accurate development framework for sensor network applications. In *SAC '07: Proceedings of the 2007 ACM symposium on Applied computing*, pages 222–226, New York, NY, USA, 2007. ACM.

[4] Crossbow. *MICA2 Datasheet*, 2008.

[5] D. Filo, D. C. Ku, and G. D. Micheli. Optimizing the control-unit through the resynchronization of operations. *Integr. VLSI J.*, 13(3):231–258, 1992.

[6] R. M. Fujimoto. Parallel and distributed simulation. In *WSC '99: Proceedings of the 31st conference on Winter simulation*, pages 122–131, New York, NY, USA, 1999. ACM.

[7] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. *SIGPLAN Not.*, 35(11):93–104, 2000.

[8] D. R. Jefferson. Virtual time. *ACM Trans. Program. Lang. Syst.*, 7(3):404–425, 1985.

[9] Z. Jin and R. Gupta. Improved distributed simulation of sensor networks based on sensor node sleep time. In *International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 204–218, 2008.

[10] O. Landsiedel, H. Alizai, and K. Wehrle. When timing matters: Enabling time accurate and scalable simulation of sensor network applications. In *IPSN '08: Proceedings of the 2008 International Conference on Information Processing in Sensor Networks (ipsn 2008)*, pages 344–355, Washington, DC, USA, 2008. IEEE Computer Society.

[11] P. Levis. *TinyOS Programming*, June 2006.

[12] P. Levis, N. Lee, M. Welsh, and D. Culler. Tossim: accurate and scalable simulation of entire tinyos applications. In *SenSys '03: Proceedings of the 1st international conference on Embedded networked sensor systems*, pages 126–137, New York, NY, USA, 2003. ACM Press.

[13] 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*, pages 79–88, Washington, DC, USA, 2002. IEEE Computer Society.

[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*, pages 95–107, New York, NY, USA, 2004. ACM.

[15] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power wireless research. In *the Fourth International Conference on Information Processing in Sensor Networks*, 2005.

[16] 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*, pages 145–152, 4-7 Oct. 2004.

[17] V. Shnayder, M. Hempstead, B. rong Chen, G. W. Allen, and M. Welsh. Simulating the power consumption of large-scale sensor network applications. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 188–200, New York, NY, USA, 2004. ACM.

[18] TinyOS-Alliance. Tinyos 1.1.15.

[19] 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*, page 67, Piscataway, NJ, USA, 2005. IEEE Press.

[20] 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*, pages 24–34, New York, NY, USA, 2007. ACM Press.

[21] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *Infocom '02,*, pages 1567–1576, New York, NY, 2002.

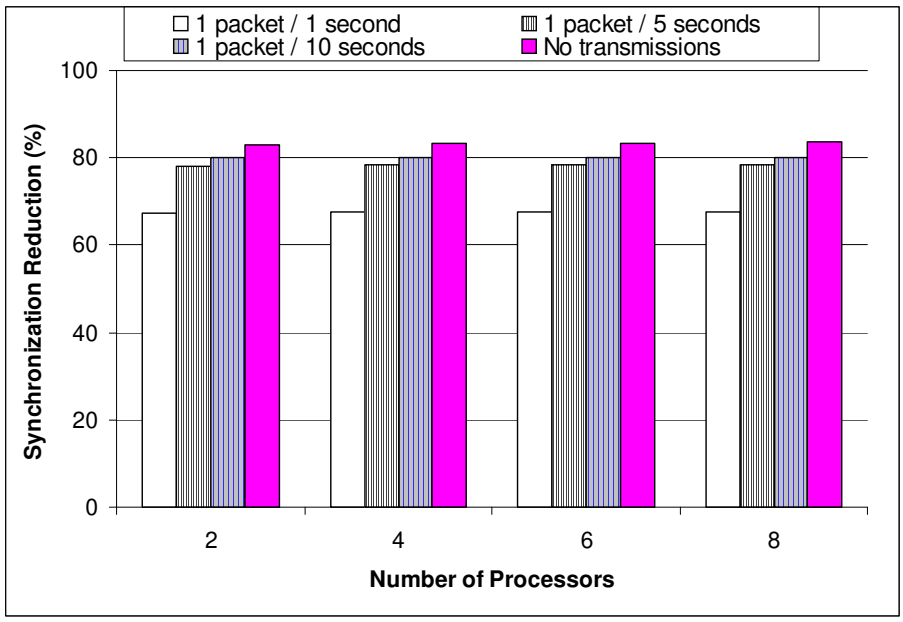

Figure 5: Percentage reductions of synchronizations using the radio-level speedup technique in PolarLite (1 sender 31 receivers, mode 3)

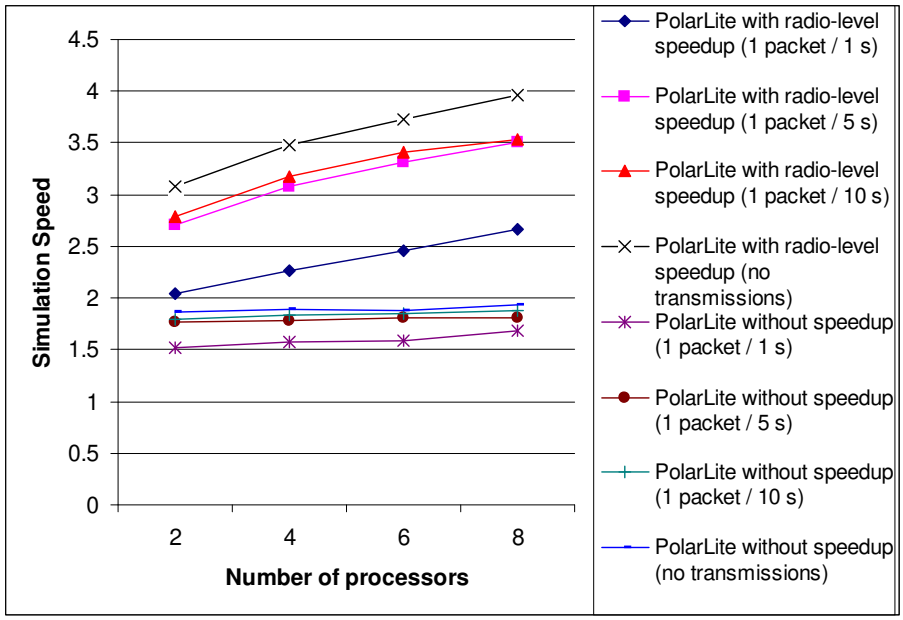

Figure 6: Speed of simulating with and without the radio-level speedup technique in PolarLite (1 sender 31 receivers, mode 3)

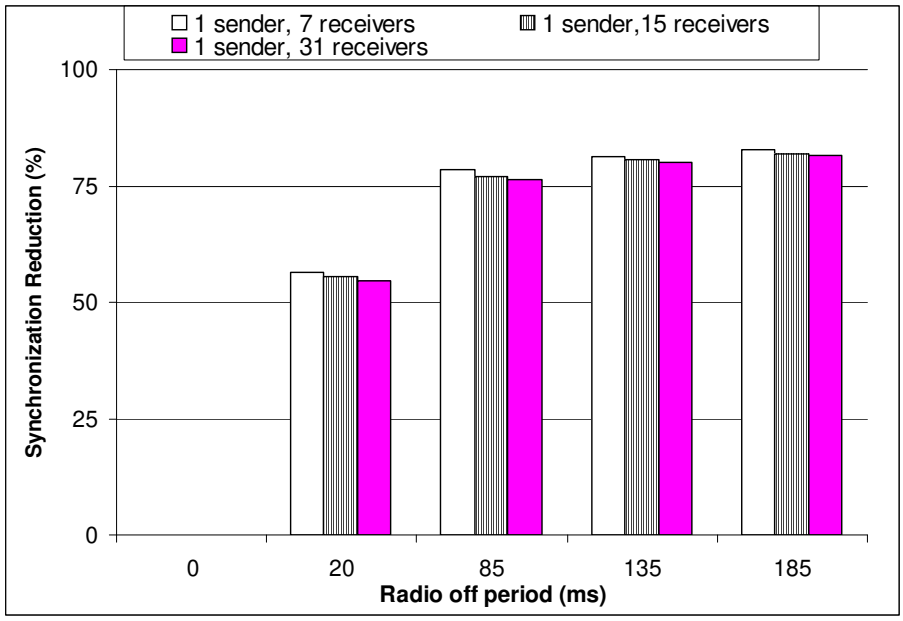

Figure 7: Percentage reductions of synchronizations using the radio-level speedup technique in simulating WSNs of different sizes and radio off times in PolarLite (1 packet/10 seconds, 8 processors)

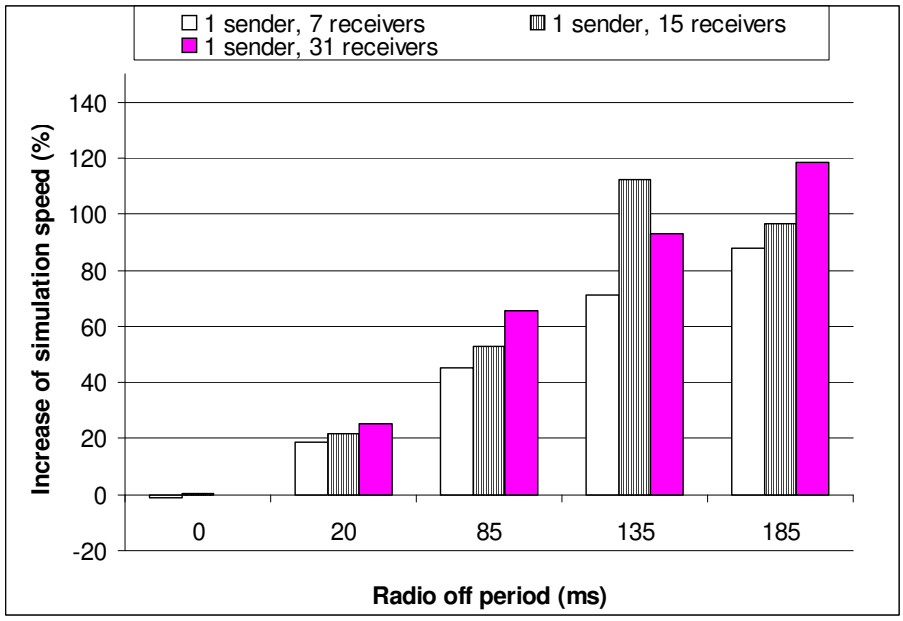

Figure 8: Percentage increases of simulation speed using the radio-level speedup technique in simulating WSNs of different sizes and radio off times in PolarLite (1 packet/10 seconds, 8 processors)

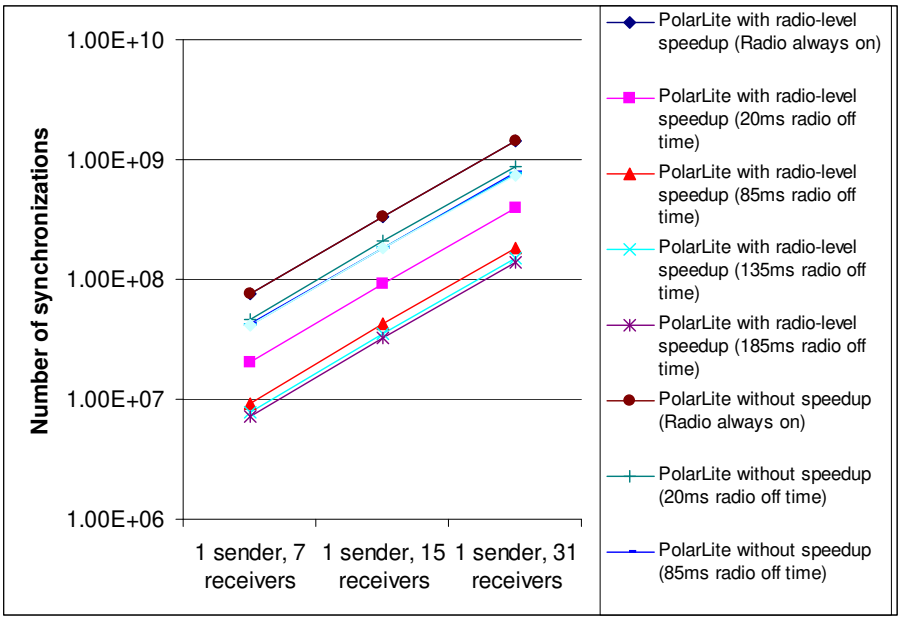

Figure 9: Total number of synchronizations in simulating WSNs of different sizes and radio off times with and without the radio-level speedup technique in PolarLite (1 packet/10 seconds, 8 processors)

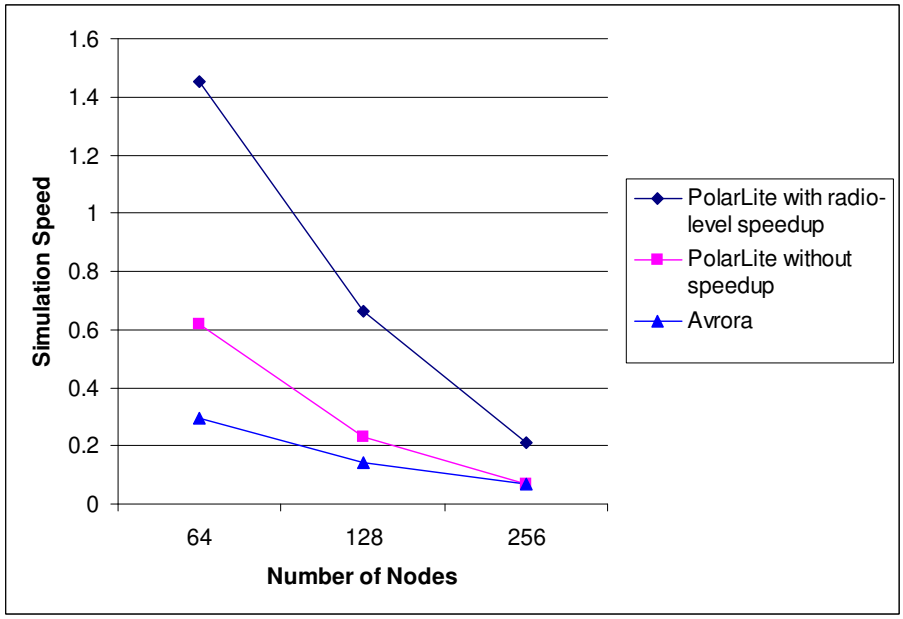

Figure 10: Speed of simulating WSNs of different sizes (1 packet/10 seconds, 8 processors, mode 3)

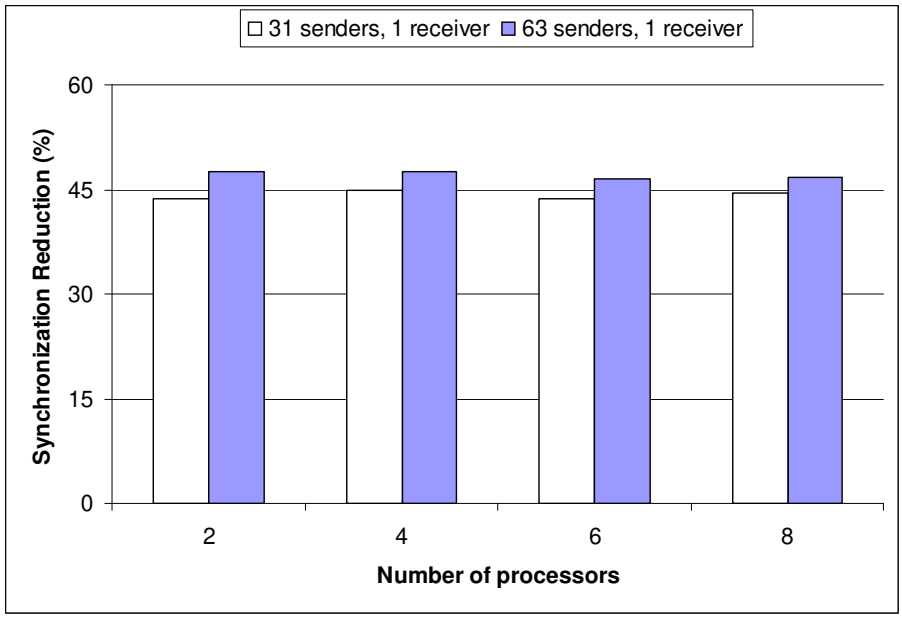

Figure 11: Percentage reductions of synchronizations with MAC-level speedup on WSNs of different sizes in PolarLite (No duty cycling)

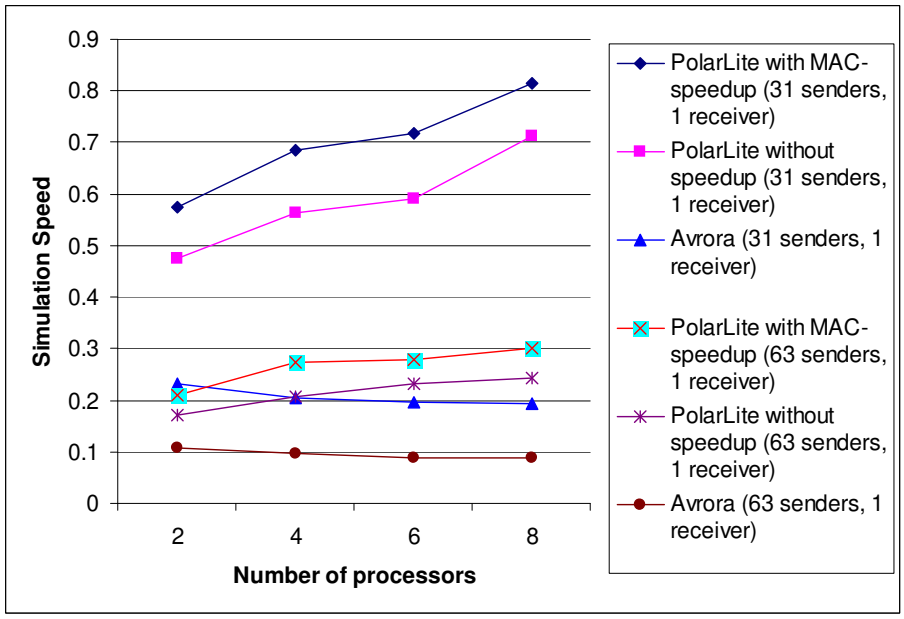

Figure 12: Speed of simulating 2 WSNs with Avrora, PolarLite and PolarLite+MAC-speedup (No duty cycling)

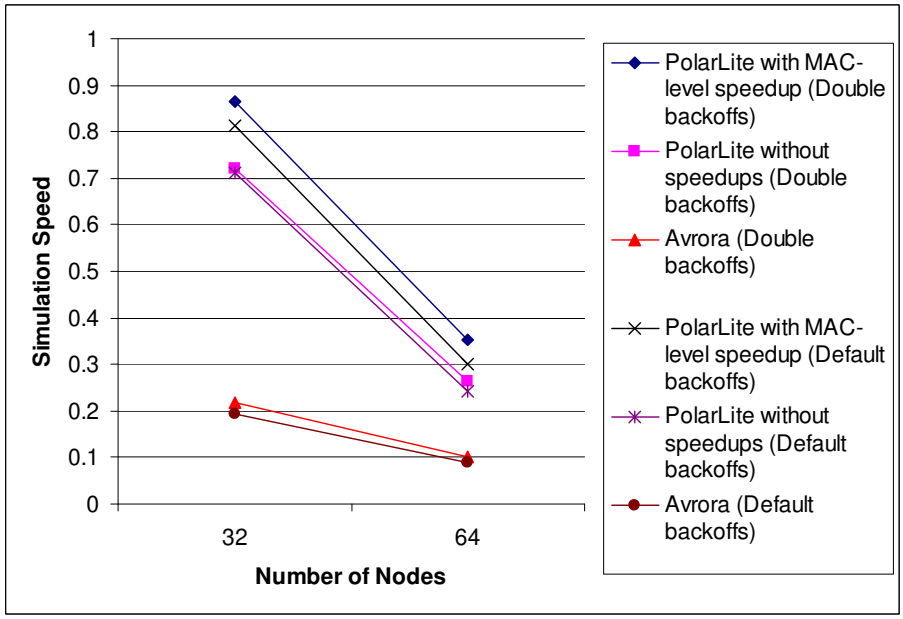

Figure 13: Speed of simulating with MAC-level speedup on WSNs using default and double sized backoff windows (No duty cycling, 8 processors)

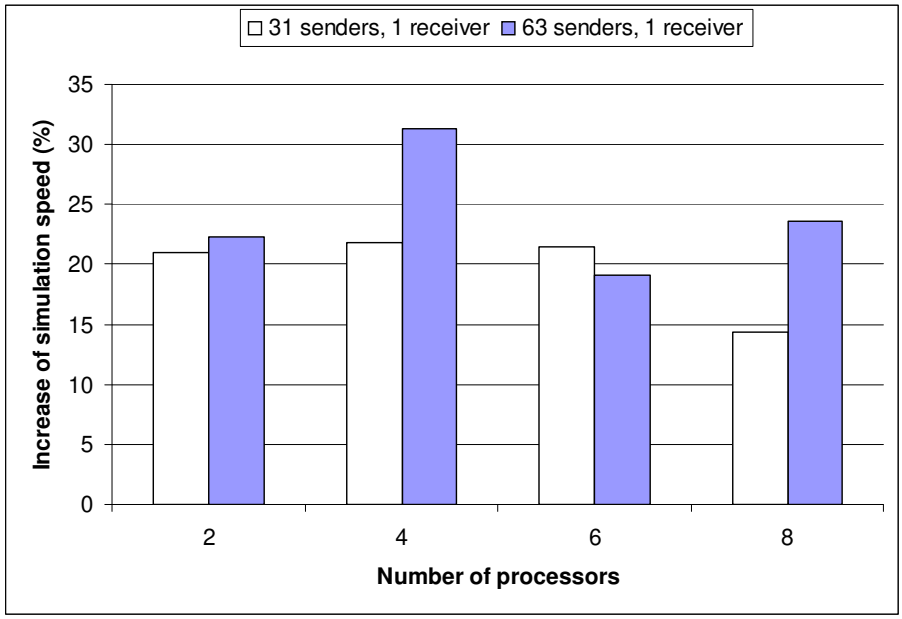

Figure 14: Percentage increases of simulation speed with MAC-speedup on WSNs of different sizes in PolarLite (No duty cycling)

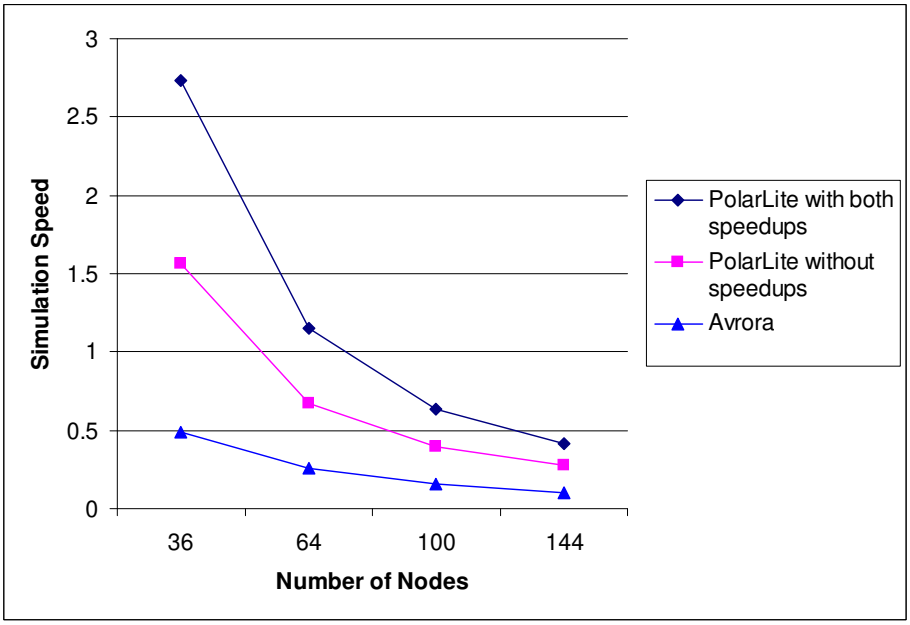

Figure 15: Speed of simulating with Avrora, PolarLite without speedups and PolarLite with both speedup techniques

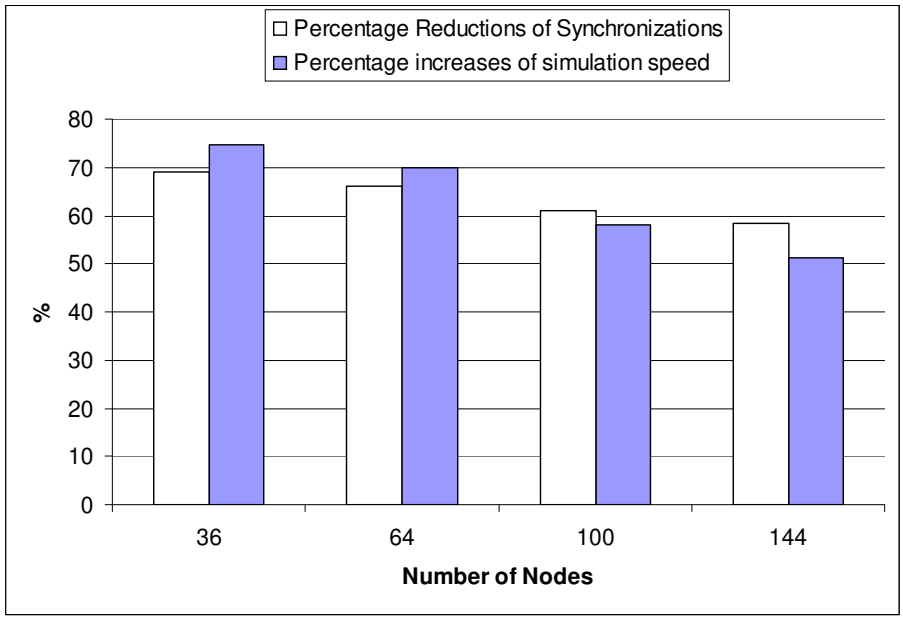

Figure 16: Percentage increases of simulation speed and percentage reductions of synchronizations with both speedup techniques in PolarLite