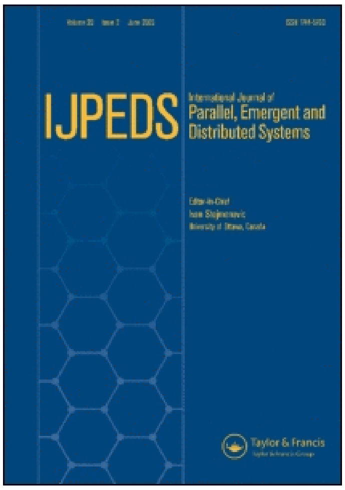

International Journal of Parallel, Emergent and Distributed Systems
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713729127

Performance tradeoffs among percolation-based broadcast protocols in wireless sensor networks

Vijay Raman [a]; Indranil Gupta [b]
[a] Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA [b] Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA

First published on: 12 February 2010

To cite this Article Raman, Vijay and Gupta, Indranil(2010) 'Performance tradeoffs among percolation-based broadcast protocols in wireless sensor networks', International Journal of Parallel, Emergent and Distributed Systems,, First published on: 12 February 2010 (iFirst)

To link to this Article: DOI: 10.1080/17445760903548309

URL: http://dx.doi.org/10.1080/17445760903548309

International Journal of Parallel, Emergent and Distributed Systems
iFirst article, 2010, 1–22

Performance tradeoffs among percolation-based broadcast protocols in wireless sensor networks

Vijay Raman[a]* and Indranil Gupta[b]

[a]*Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA;* [b]*Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*

(*Received 9 September 2009; final version received 12 September 2009*)

Broadcast of information in wireless sensor networks is an important operation, for example, for code updates, queries and membership information. In this paper, we analyse and experimentally compare the performance of vanilla versions of several well-known broadcast mechanisms: flooding, site percolation, bond percolation and modified bond percolation. While flooding and some percolation-based approaches have been compared in the literature, there is no all-to-all comparison among all schemes. We carry out our comparison for several network topologies defined by node locations: random, grid and clustered. Our analysis is performed at the link layer level, where we use a propagation model based on real experiments from the literature. The link model used is independent of the medium access control (MAC) layer and, therefore, helps us in arriving at the best possible values for the metrics that we compare in our analysis. Our main metrics are bandwidth, energy usage and broadcast latency. Our analytical and experimental results show that, given a desired high reliability for all topologies, flooding has the lowest latency but consumes the most energy per broadcast. For dense networks, site percolation achieves comparable latency and reliability to flooding, while lowering energy consumption. Modified bond percolation further lowers energy consumption compared to site percolation, while basic bond percolation leads to a latency increase. For sparse networks, results are similar to a dense network except that site percolation consumes lower energy than modified bond percolation. We briefly discuss implications for different broadcast applications.

Keywords: gossiping; percolation; code updates; query propagation; membership management; broadcast

1. Introduction

In wireless sensor networks, operations such as code updates, queries, membership management and routing updates all rely on broadcast. Several mechanisms have been designed in the literature to address the broadcast problem [9,10,13,17]. In this paper, we abstract out the basic broadcast mechanisms from these papers into four classes: flooding, site percolation, bond percolation and modified bond percolation. We analytically and experimentally compare these four approaches to each other. Although flooding has been individually compared to some of the percolation approaches [9,17], there is a lack of an all-to-all comparison among these approaches, and thus the tradeoffs among them have not been comprehensively studied.

*Corresponding author. Email: vraman3@illinois.edu

ISSN 1744-5760 print/ISSN 1744-5779 online

DOI: 10.1080/17445760903548309
http://www.informaworld.com

Abstractly, flooding is a brute force approach where each node sends out data to all its neighbouring nodes, each of which then forwards it similarly. Site percolation introduces, at each node, a probability of forwarding data to all its neighbours; thus, only a fraction of nodes forwards the data to their neighbours. Studies [9] have shown that when this fraction is above a threshold, all nodes receive the broadcast with high probability. Bond percolation is a flavour where each link between neighbouring nodes is associated with a forwarding probability [17]. Thus, even if a node decides to forward the broadcast, only some of its neighbours may receive it. Once again, studies have shown that when the fraction of a node's neighbours receiving the broadcast is higher than a threshold, all nodes receive the broadcast with high probability. In this paper, we first consider a vanilla version of bond percolation where the receiving neighbours are selected at random. We also consider a modified bond percolation variant, where the neighbour selection (for forwarding) arises out of the sleep schedules of the nodes. This is based on approaches [17] where nodes have a periodic sleep schedule (e.g. IEEE 802.11 Power Save Mode or B-MAC/S-MAC [19,23] sleep schedules), and these schedules are modified in order to enable probabilistic forwarding of broadcasts to some neighbours. We will elaborate on all these broadcast schemes in Section 2.

Broadcast is useful for varied purposes: Trickle [13] uses it for sending out continuous code updates. Aggregation using either Tiny AGgregation Service [15], directed diffusion [11] or synopsis diffusion [18] partly relies on broadcast, for example, to initiate the aggregation epochs. Routing protocols utilise broadcasts to spread routing updates. For example, [1,3,12] provide a detailed survey of various routing protocols for sensor networks. Additionally, [2] provides a survey of various energy-efficient routing protocols. Membership management within small or large groups of sensor nodes has also been looked at by researchers and utilises broadcast, e.g. [14]. As we discuss in Section 4, these applications have highly varied requirements and thus a one-size-fits-all broadcast solution may not work for all. This motivates our study of the tradeoffs among different broadcasting approaches.

The remainder of the paper is organised as follows: in Section 2, we briefly discuss the various broadcast mechanisms that are the focus of this paper and derive their energy consumption. Section 3 presents simulation results of these broadcast schemes for three different network layouts. Section 4 discusses the implications of our results for practical broadcast operations. Finally, we conclude our paper in Section 5.

2. Broadcast schemes and energy analysis

In this section, for each of the broadcast algorithms considered, we (1) briefly describe the algorithm and then (2) mathematically analyse the energy consumption of these techniques. We describe the analytical assumptions in Section 2.1 and the individual algorithms in the subsequent subsections.

2.1 Assumptions for the analysis

For the mathematical analysis of the broadcast mechanisms, we make the following assumptions:

(1) Every sensor node in the network is assumed to have a globally unique identity through which it can be uniquely addressed. If nodes do not have unique pre-assigned identities, then the ID can be assigned either based on their location [16,21,24] or randomly.

(2) The sensor nodes are static and do not move.
(3) All sensor nodes are homogeneous with respect to their transmit, receive and idle power consumption characteristics as well as wireless bandwidth.
(4) We only worry about failure in transmissions that can occur due to the propagation characteristics of the medium so as to obtain a best estimate of the metrics. We, therefore, do not model collisions in this paper. We will discuss why this is a valid assumption in Section 4.2. Furthermore, we assume a MAC layer that takes care of retransmissions and contentions.
(5) The network of sensor nodes is assumed to be connected. In other words, we assume that any pair of non-faulty nodes can communicate via at least one path consisting of all non-faulty nodes. We consider only crash-recovery failures that include sleep schedules.

We have used a realistic propagation model that is independent of the MAC layer as done in [7]. Instead, the propagation model uses the link layer characteristics. We choose to not emulate the MAC layer since sensor nodes may use a wide variety of MAC protocols (e.g. B-MAC [19], S-MAC [19] and T-MAC [6]) and we wish our results to apply generically. This choice allows us to focus on incorporating loss due to realistic propagation models such as [7]. A further evaluation with individual MAC protocols could be a fruitful future work. However, as we discuss at the end of the paper, our results are likely to hold for specific MAC protocols too.

We denote by N the number of nodes in the sensor network and by k the average number of $(1 - \text{hop})$ neighbours of a node in the network. We analyse only the power consumed by a broadcast packet, which can be interpreted as the energy consumed by a normalised packet that takes 1 second to transmit on a link. Larger broadcast packets will consume proportionally more energy.

2.2 Flooding

This is the simplest of the three mechanisms. Each node that receives a message simply re-broadcasts it to all its 1-hop neighbours. Further, the node keeps track of recently received broadcasts and does not re-broadcast duplicates. Thus, each node re-broadcasts a received message only once.

2.2.1 Analysis

In flooding, all the nodes broadcast message to their neighbours once, and the number of messages per broadcast is just 1. Further, since each node has k neighbours, we can write the total number of messages sent per node S_{flooding} and the average number of messages received per node R_{flooding} as

$$S_{\text{flooding}} = 1, \quad R_{\text{flooding}} = k.$$

To calculate the total energy consumed, we use real measurements for Mica (2001) motes from [20], which showed that the transmit power required for these types of motes for a packet (36 mW at 0 dBm) is about thrice the power required for receiving a packet (12 mW). Furthermore, since the idle listening power is the same for all nodes (which is 8 mW), we ignore that in our calculations. Accordingly, if the power required by a sensor node for receiving a packet is denoted by Pow, the average power consumed (or energy consumed in one second) by the sensor nodes for both receiving and transmitting of the broadcast process P_{flooding} is

$$P_{\text{flooding}} = \frac{(3N + kN)\text{Pow}}{N} = (3 + k)\text{Pow}.$$

2.3 Site percolation-based gossiping

In this scheme, each node that receives a message broadcasts the message to its neighbours with a probability p called *gossip probability*. A node does not broadcast the message with a probability $(1 - p)$. Thus, only a fraction of all nodes forwards the broadcast message, thereby reducing the number of messages in the network and hence the power consumption of the nodes.

We call this variant as site percolation; Haas et al. [9] proposed such a gossip protocol and used it for transmitting routing updates in AODV. The terminology is derived from *percolation theory* [8], an old branch of mathematics that studied the residual connectivity of a graph when some fraction of nodes is removed at random from it. Haas et al. [9] showed that a gossip probability p that is higher than a threshold regime of between 0.6 and 0.8 is sufficient to ensure that almost every node in the network gets the broadcast message with high probability. In other words, if p is lower than this threshold, very few nodes receive the broadcast. However, as p crosses this threshold, there is an abrupt phase transition and all nodes receive the broadcast with high probability. Note that a gossip probability of $p = 1.0$ makes site percolation equivalent to flooding.

2.3.1 Analysis

If p is the gossip probability with which each of the nodes decide independently to transmit a message, then the expected number of messages sent per node in the network and the average number of messages received per node during broadcast are given by

$$S_{\text{site}} = \frac{1}{N}\sum_{i=0}^{N} i\binom{N}{i}p^{i}(1-p)^{N-i} = p, \quad R_{\text{site}} = \sum_{i=0}^{k} i\binom{k}{i}p^{i}(1-p)^{k-i} = kp.$$

The power consumed by the nodes is thus

$$P_{\text{site}} = \frac{(3Np + kpN)\text{Pow}}{N} = (3 + k)p\,\text{Pow}.$$

2.4 Bond percolation-based gossiping

In percolation theory, the counterpart of site percolation is bond percolation. If k is the number of neighbours for a node in the network, then a node receiving a message will forward the message to only $m(m \leq k)$ of these neighbours, picked at random. Thus, only some neighbours of each node receive the message. Like site percolation, bond percolation also shows a phase transition behaviour. Specifically, only when the quantity m/k is higher than a threshold value of about 0.5, does the broadcast reach all nodes with high probability [8].

In reality, however, bond percolation cannot be implemented as is, since a message transmitted by a node will be received by all its neighbours that are *awake* at that point of time. Thus, the implementation of this scheme, called probability-based broadcast forwarding [17], modifies the periodic sleep schedule of each node, so that it stays awake with some probability during the sleep part of the schedule. Such a node can immediately re-broadcast any message it receives (with some probability), and this message will be received only by those neighbours that also happen to be opportunistically awake at that time. We call this variant as modified bond percolation.

For the sake of comparison with the other schemes, we assume that in modified bond percolation, m out of k neighbours receive the broadcast message. The main difference between modified bond percolation and bond percolation is that each node transmits only 1 broadcast message in the former and m copies in the latter.

2.4.1 Analysis

In bond percolation, since each node sends m messages, we have

$$S_{\text{bond}} = m.$$

However, for the case of the modified bond percolation, each node will send only one message. Therefore, the number of messages sent per node is

$$S_{\text{mod-bond}} = 1.$$

The average number of messages received per node is the same for both the above variants of bond percolation and is given by

$$R_{\text{bond}} = R_{\text{mod-bond}} = \sum_{i=0}^{k} i \binom{k}{i} \left(\frac{m}{k}\right)^{i} \left(1 - \frac{m}{k}\right)^{k-i} = m.$$

The power consumed by each of the bond percolation mechanisms can be given as

$$P_{\text{bond}} = \frac{(3mN + mN)\text{Pow}}{N} = 4m\text{Pow},$$

$$P_{\text{mod-bond}} = \frac{(3N + mN)\text{Pow}}{N} = (3 + m)\text{Pow}.$$

2.5 Conclusions from analysis

The tradeoffs among the four different schemes with respect to energy consumed are illustrated in Figure 1 and explained below. In Figure 1, each arrow is directed from the left-hand side to the right-hand side of an inequality. The expressions next to the arrows indicate the conditions under which the inequality is satisfied.

1. $P_{\text{site}} \leq P_{\text{flooding}}$, if $p \leq 1$. This is always true as p is a probability.
2. $P_{\text{mod-bond}} \leq P_{\text{flooding}}$, if $m \leq k$. This is always true, as modified percolation chooses an m that is always lower than k.
3. $P_{\text{mod-bond}} \leq P_{\text{bond}}$, if $m \geq 1$. This is always true for all practical purposes, since at least one neighbour receives the broadcast message in bond percolation mechanisms.

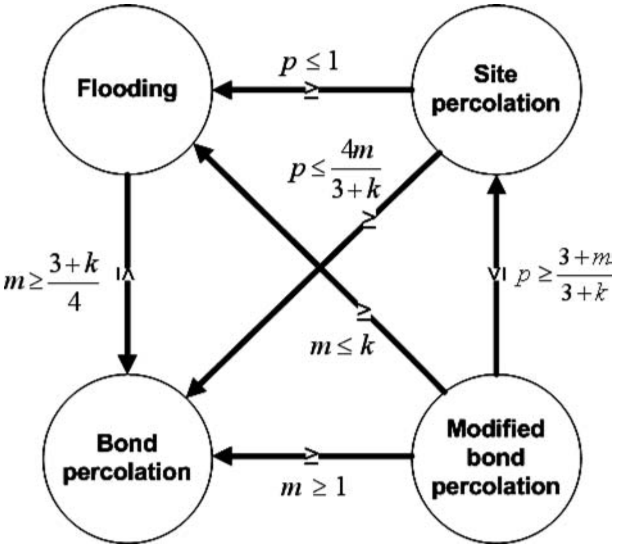

Figure 1. Conclusions on energy consumed (see Section 5 for details).

(4) $P_{\text{site}} \leq P_{\text{bond}}$, if

$$(3+k)p\text{Pow} \leq 4m\text{Pow} \quad p \leq \frac{4m}{3+k}.$$

Choosing $p = 0.7$ (a value between 0.6 and 0.8, as suggested by [9]), and $m/k = 0.5$ (as suggested by [17]), we require $k \geq 2$ for $P_{\text{site}} \leq P_{\text{bond}}$ to be true. However, $k < 2$ implies a network of disjoint node pairs, which is unlikely in actual deployments. Therefore, $P_{\text{site}} \leq P_{\text{bond}}$ for all practical networks.

(5) $P_{\text{flooding}} \leq P_{\text{bond}}$, if

$$(3+k)\text{Pow} \leq 4m\text{Pow} \quad m \geq \frac{3+k}{4}.$$

Again, by choosing $p = 0.7$ and $m/k = 0.5$, we have $k \geq 3$ for $P_{\text{flooding}} \leq P_{\text{bond}}$ to be true. Examples of networks with $k < 3$ can be either a ring or a linear topology. We believe these are unlikely to occur in real deployments. Thus, for all other practical networks, we can conclude that $P_{\text{flooding}} \leq P_{\text{bond}}$.

(6) $P_{\text{mod-bond}} \leq P_{\text{site}}$, if

$$(3+m)\text{Pow} \leq (3+k)p\text{Pow} \quad p \geq \frac{3+m}{3+k}.$$

For $p = 0.7$ and $m/k = 0.5$, we get $k \geq 4.5$ for $P_{\text{mod-bond}} \leq P_{\text{site}}$. In other words, site percolation consumes lower energy than modified percolation when the network is sparse, i.e. has typical node degree ≤ 4.

From our discussions, we conclude that modified bond percolation consumes the lowest power per broadcast for dense networks, while site percolation consumes the lowest power for sparse networks. Furthermore, irrespective of the density of the networks, bond percolation consumes the most power, followed by flooding.

3. Simulation results

While the analytical studies in the previous section looked at the energy consumption, we wish to further study the latency–energy tradeoffs among the schemes. Further, sensor network layouts differ from one deployment scenario to the next – some deployments distribute sensor nodes randomly in order to have a good coverage of the monitored area (e.g. indoor locations [4]), others have a more regular layout where nodes are deployed in a grid formation (e.g. biomedical sensors [22] or table-top deployments), while some others are likely to be highly clustered (e.g. Burrow sensors in the Great Duck Island [5]).

In this section, we compare the four broadcast mechanisms using custom simulations in MATLAB across these different network topologies. Before proceeding to the discussion on our simulations, we would like to emphasise a key aspect arising from our analysis. We observe from our analytical results that the power consumption values are only dependent on the density of the network. The link layer characteristics are, therefore, vital for our analysis. For modelling the link layer, we have used the actual traces from [7], which is based on real-world experiments. According to this link layer model, the effective neighbourhood of a node is determined by the probability to receive a transmission of the surrounding nodes.

We have replicated the probability of successful packet reception for the medium- and the low-power scenarios from [7] in Figure 2. The plot shows the probability of successful packet reception as a function of distance and transmit power. While the term ‘power’ actually implies the transmit power in [7] (that is controlled using a potentiometer), the power consumption values used for comparison in this paper (both in this section and in Section 2) are written in terms of the received power, Pow. We can, therefore, consider *Pow* to be the received power, when the transmit power is set as in [7].

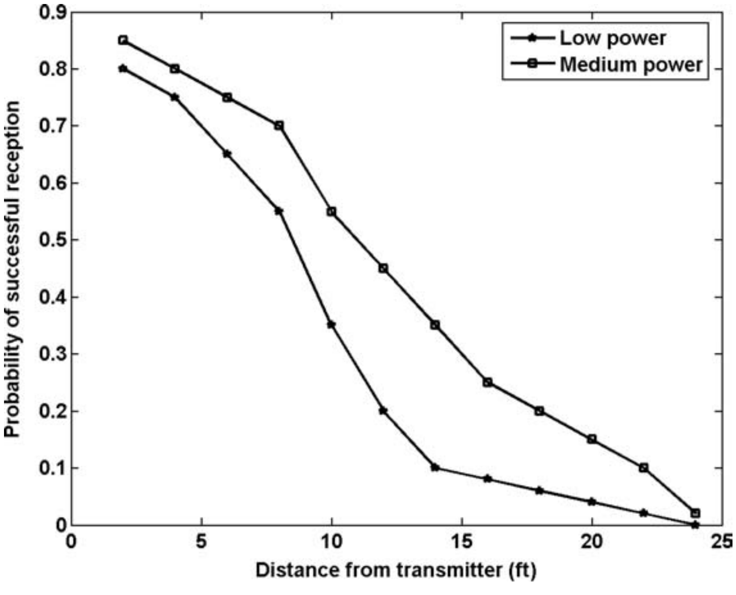

Figure 2. Probability of successful packet reception as a function of distance and power.

3.1 Experimental methodology

We simulate three different topologies, namely, (1) random, (2) clustered and (3) grid. For the random topology, we distribute the nodes uniformly at random on a 50 × 50 area. For the clustered topology, we create four square clusters, each with a side 20 on the four corners of a 50 × 50 square. Nodes are distributed uniformly at random within each cluster and are shared evenly by the four clusters. The number of nodes, for both the random and the clustered topology, is varied from 1000 to 2500 in steps of 250. In the case of the grid topology, we vary the grid size from 30 to 50 in steps of 5. The number of nodes in the grid topology is the square of a side of the grid and hence varies from 900 to 2500 (which is comparable to that for the random and the clustered topology). While the random and clustered topologies are unplanned networks, a grid topology is a planned network. We performed our simulations for both the medium- and low-power scenarios (corresponding to a potentiometer setting of 66 and 69, respectively) discussed in [7]. We have provided detailed results only for the low-power scenario in this paper. The results for the medium-power characteristics follow a similar trend as that of the low-power case. We, however, provide some interesting observations obtained from our simulations performed at the medium-power setting.

The minimum node degree (average number of neighbours) for the low-power model turned out to be 15 for random topology, 10 for the clustered topology, (corresponding to 1000 nodes) and 30 for the grid topology (corresponding to grid size 30). Since these are dense networks, we also simulate a sparse grid topology with a smaller node degree (around 4 for a grid size of 30) by choosing a transmission range of 1. For site percolation, we choose a gossip probability p of 0.7, as it lies between the threshold regime of 0.6 and 0.8, suggested in [9]. For the two bond percolation schemes, we choose the m/k ratio (where m is the number of nodes that receive the broadcast message out of k neighbours) to be the exact threshold value of 0.5, as mentioned in [17]. We generate 100 different realisations of the network for each of the three topologies, each with a different seed. During each realisation, we pick a node uniformly at random and generate a broadcast message from this node for propagation throughout the network. Furthermore, we generate different link characteristics for each realisation.

We simulate all the four broadcast mechanisms explained in Section 2. Our metrics of interest are the number of messages sent, number of messages received, average energy consumed and latency of receiving broadcasts (since broadcast initiation time). We initiate broadcasts from only one sender node. Our baseline for the comparison is to set parameters for the four broadcast schemes, so that they just barely ensure 100% reliability for all broadcast messages.

3.2 Message and bandwidth consumption

Figures 3–5 show the plots for the logarithm of the number of messages sent in the network during the broadcast process in the random, clustered and the grid topologies, respectively. We can see from the plots that for bond percolation, the number of messages sent grows very quickly as the number of nodes increase. This is because each node has to send m individual messages. However, the other three percolation schemes are almost identical, except that site percolation sends slightly fewer messages.

The average number of messages received per node is plotted for the three topologies in Figures 6–8. We observe that the average number of messages received per node is the lowest for the two bond percolation schemes. We can also note that the average number of

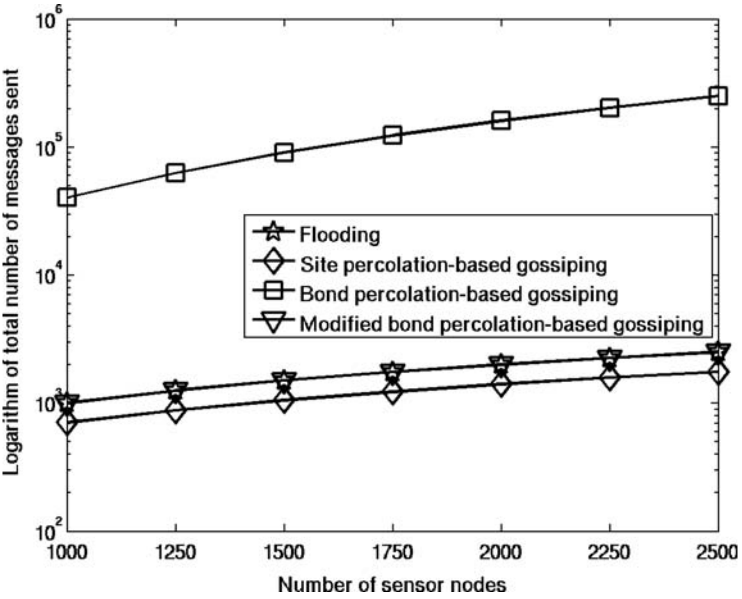

Figure 3. Number of messages sent in the network – random topology.

messages received for the site percolation is lower than that of flooding, due to the probabilistic broadcast nature of the site percolation scheme.

3.3 Latency

Next, we plot in Figures 9–11, the latency involved in broadcast for all the three topologies. The time values are normalised to the units of *T*, where *T* is the time taken

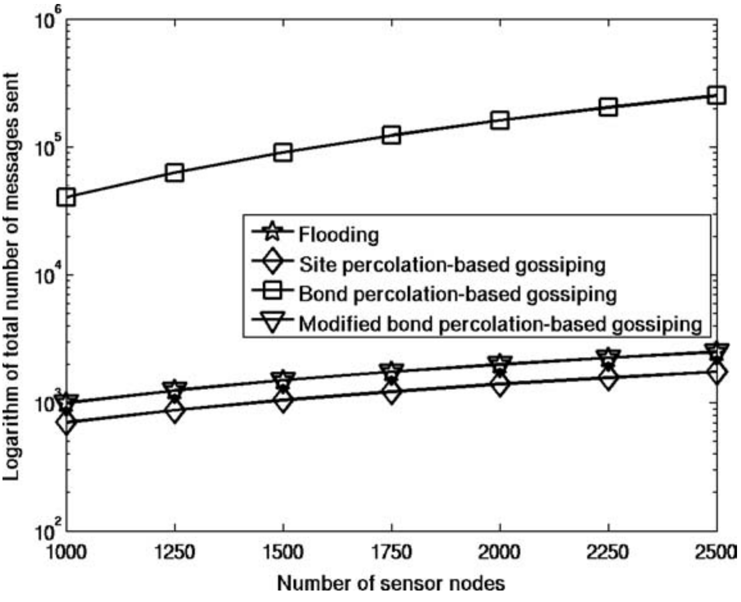

Figure 4. Number of messages sent in the network – clustered topology.

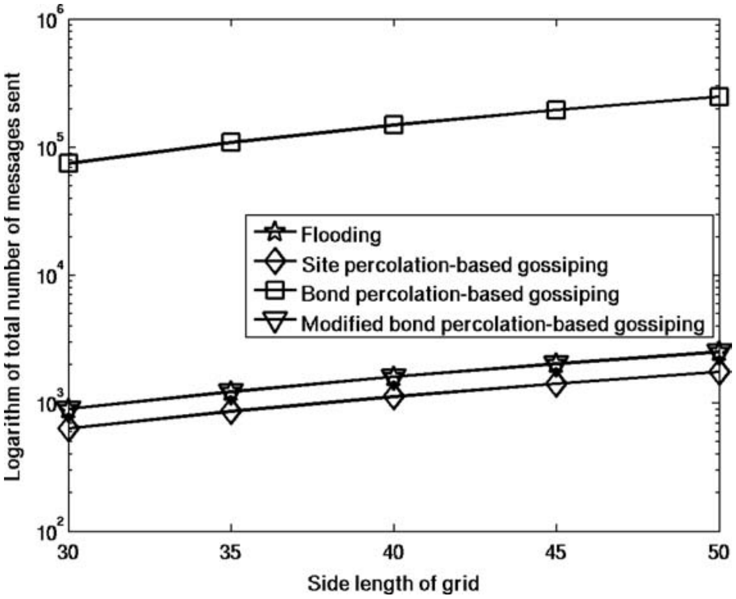

Figure 5. Number of messages sent in the network – grid topology.

to complete one transmission from one node to its neighbour and, therefore, are not absolute latency values. We observe that flooding is the fastest, while site percolation is only slightly slower. However, the latency for the two bond percolation schemes is high since fewer nodes receive the message per broadcast. This leads to longer propagation times for broadcasting the message to the entire network. We further clarify this by plotting the time at which each of the nodes received the broadcast message, normalised

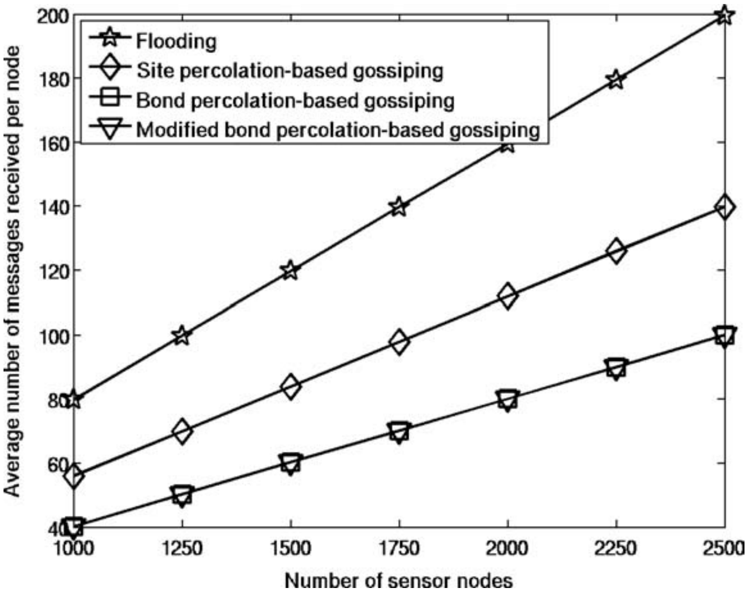

Figure 6. Average number of messages received per node – random topology.

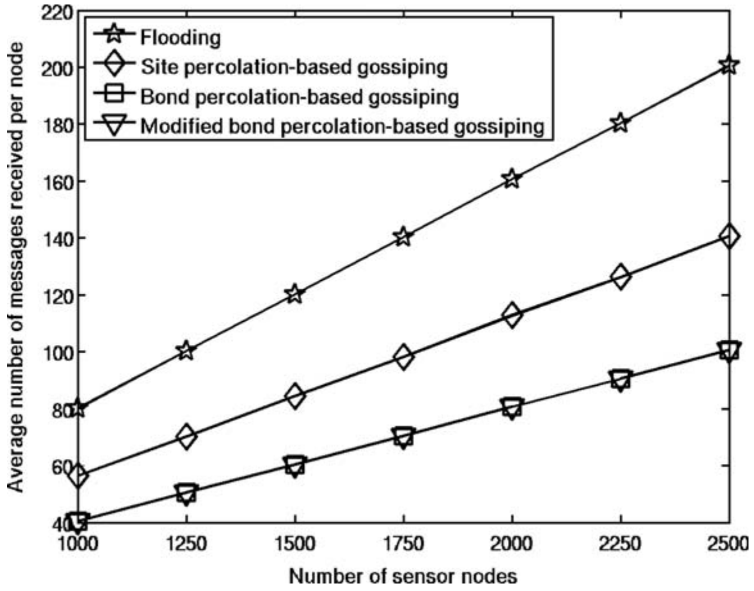

Figure 7. Average number of messages received per node – clustered topology.

to the maximum time taken for the packet to reach the destination. We plot this for a 2500-node network. Figures 12–14 show the corresponding plots. We observe from the plot that very few nodes receive a broadcast at any point in time in the case of the two bond percolation schemes when compared to flooding and site percolation schemes. This suggests that the bond percolation schemes require many more transmissions to reach the destination (than the flooding and site percolation schemes) leading to a higher latency.

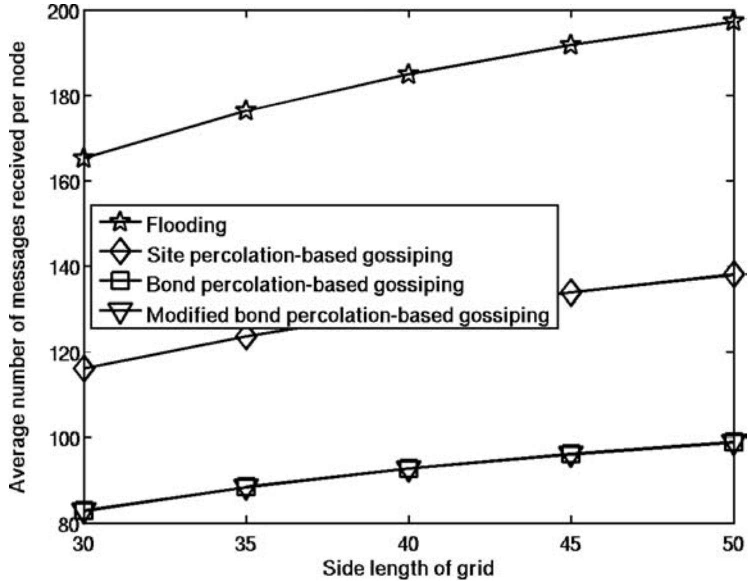

Figure 8. Average number of messages received per node – grid topology.

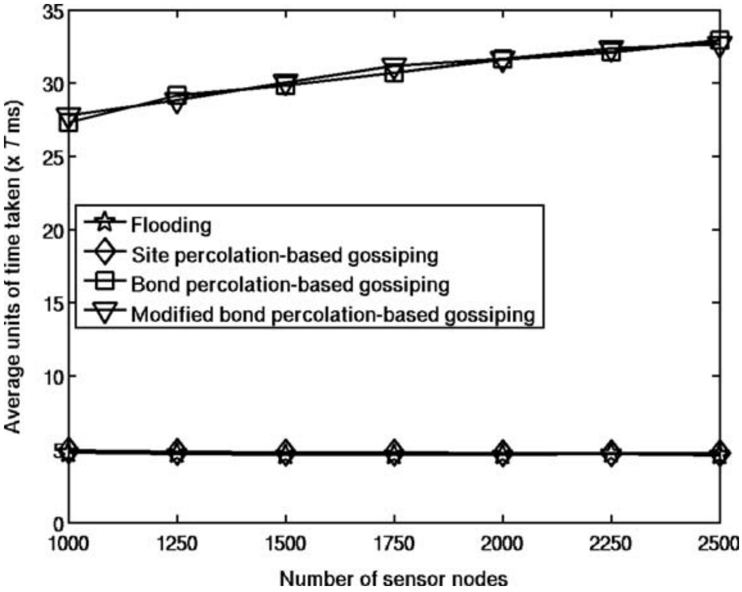

Figure 9. Average units of time taken – random topology.

3.4 Energy consumption and tradeoffs

Similar to our analysis (Section 2.1), we measure the power consumed by a broadcast packet, which is the energy consumed by a normalised 1 s long packet. Figures 15–17 show the units of average received energy (or power) and the total (sum of received and transmit) energy consumed by the nodes during the broadcast process (in units of Pow = P, see Section 2.1). The sum of the transmit and received powers gives the total

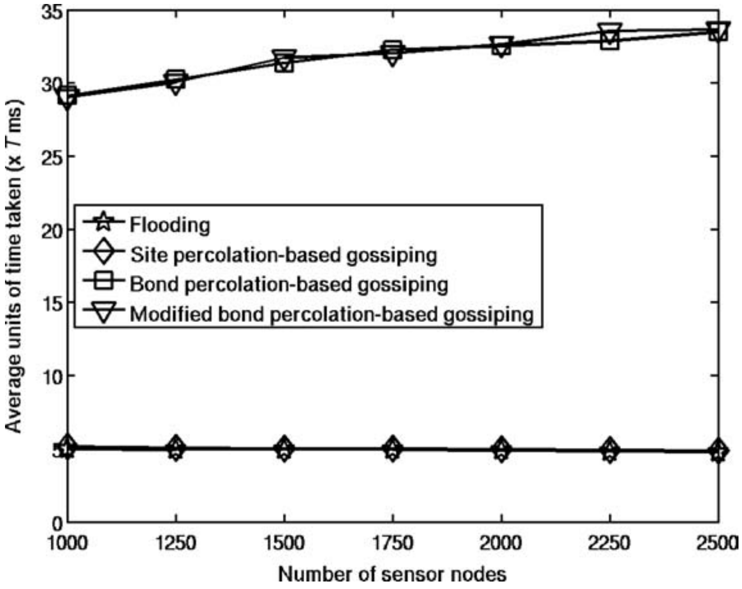

Figure 10. Average units of time taken – clustered topology.

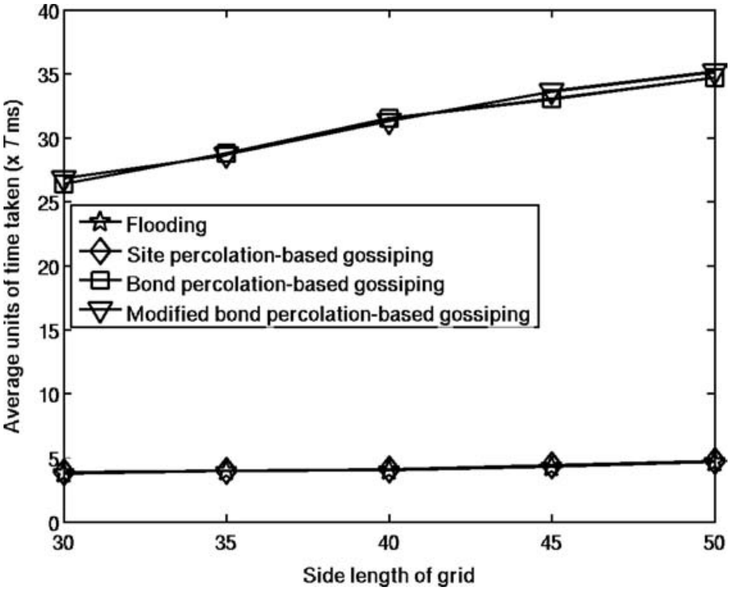

Figure 11. Average units of time taken – grid topology.

average power consumed by the nodes (not explicitly shown in the figures). As shown by the figure, the bond percolation scheme consumes more transmit power, as it is required to make *m* individual transmissions. This is evident from the plot for the total power consumed. However, the received power is very low for both the bond percolation schemes. As a result, the modified bond percolation scheme consumes the least total power among all the broadcast schemes. Site percolation consumes lower power than flooding.

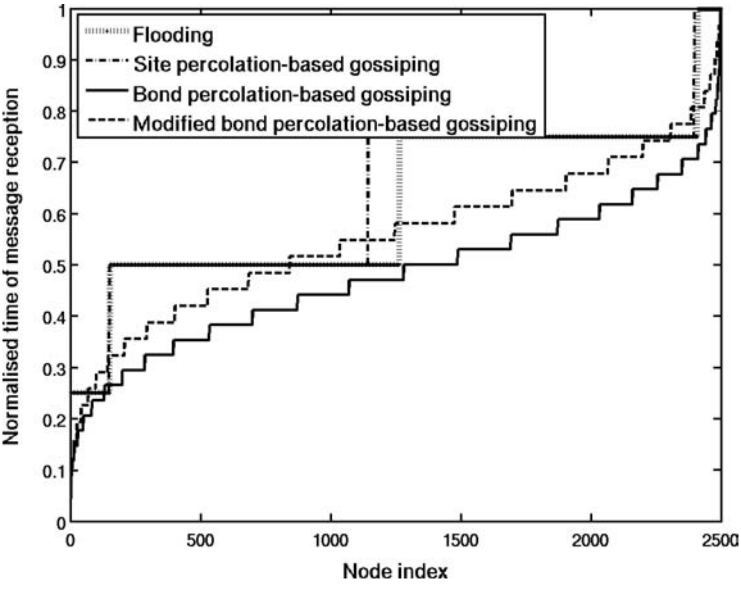

Figure 12. Normalised time taken to receive a packet – random topology.

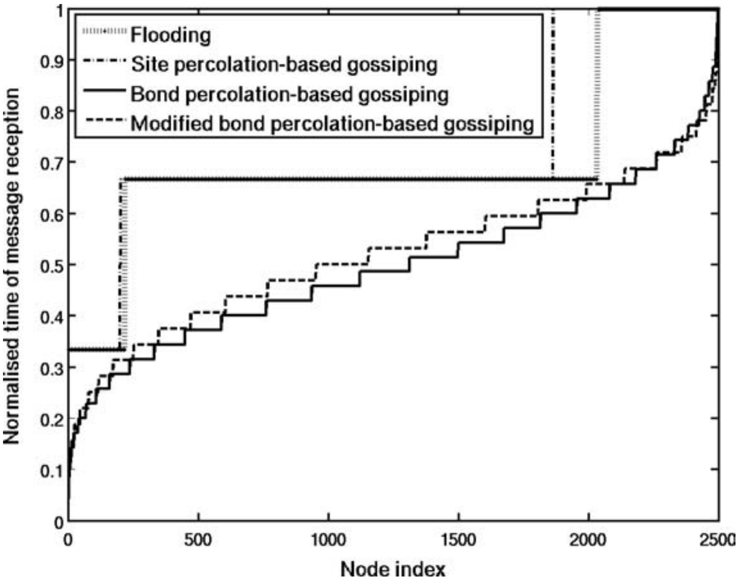

Figure 13. Normalised time taken to receive a packet – clustered topology.

To show that site percolation may consume lower energy than a modified percolation when the number of neighbours, $k \leq 4.5$, we simulate a sparse grid topology with the same number of nodes as before (900–2500). However, we reduce the transmission range for each node to just one. This ensured that each node has only four neighbours on average. The corresponding plot is shown in Figure 18. We observe from the figure that the total power consumed by site percolation is slightly lower than modified bond percolation, as predicted by our analysis.

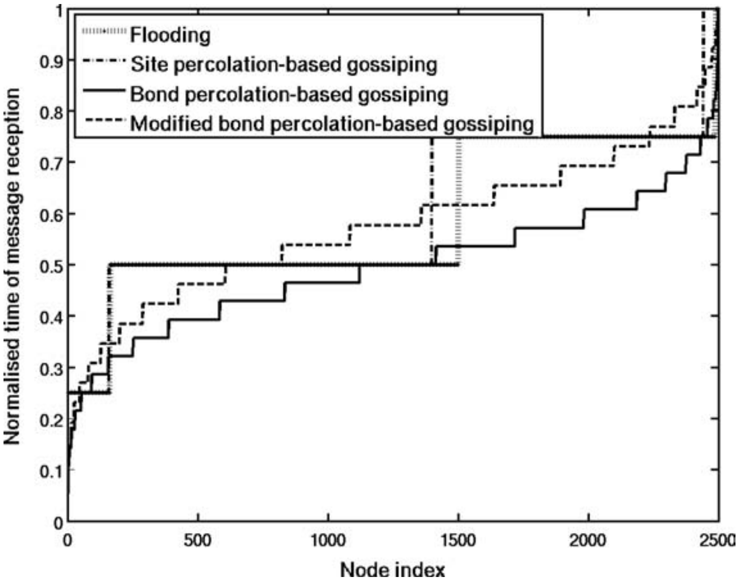

Figure 14. Normalised time taken to receive a packet – grid topology.

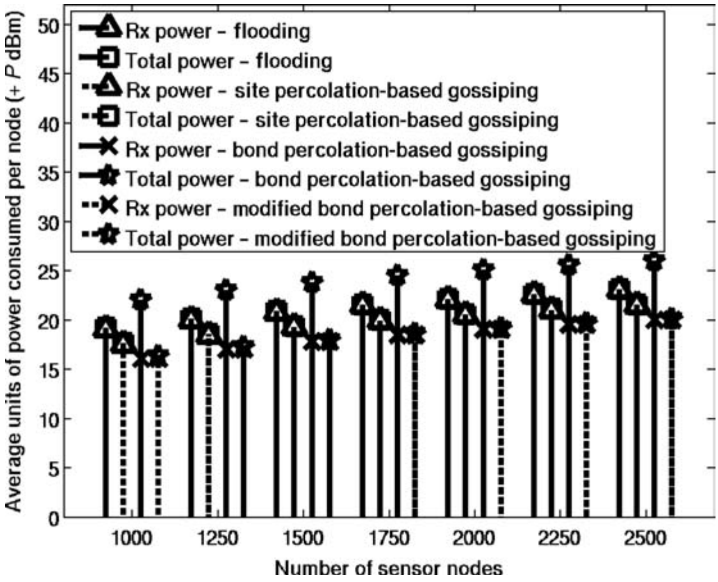

Figure 15. Average units of power consumed – random topology.

Finally, in order to map the tradeoffs among the four broadcast mechanisms, we vary the number of nodes in the system and plot the latency as well as total energy consumed per node during the entire broadcast propagation. This is shown in Figures 19–21. In these figures, each point corresponds to a varying number of nodes from left-to-right and bottom-to-top (1000–2500 for random and clustered topology and 30–50 for grid topology, as described earlier). We can see that the modified bond percolation scheme consumes the lowest power but incurs a high latency. In comparison, the flooding and site

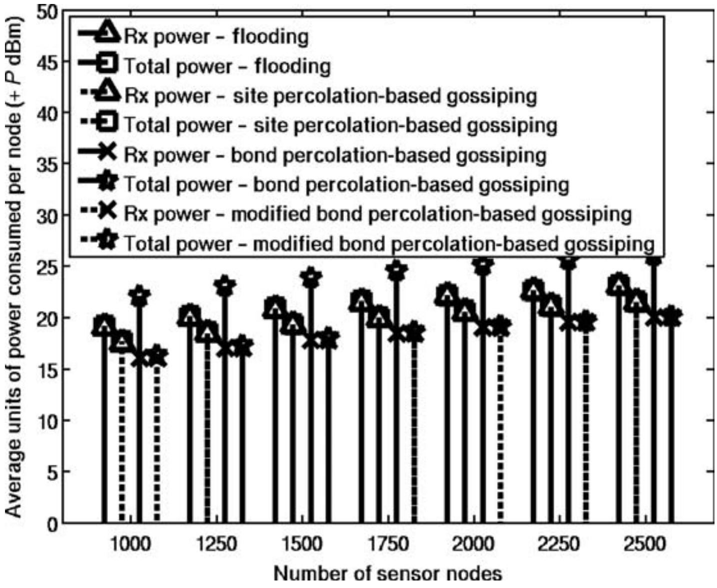

Figure 16. Average units of power consumed – clustered topology.

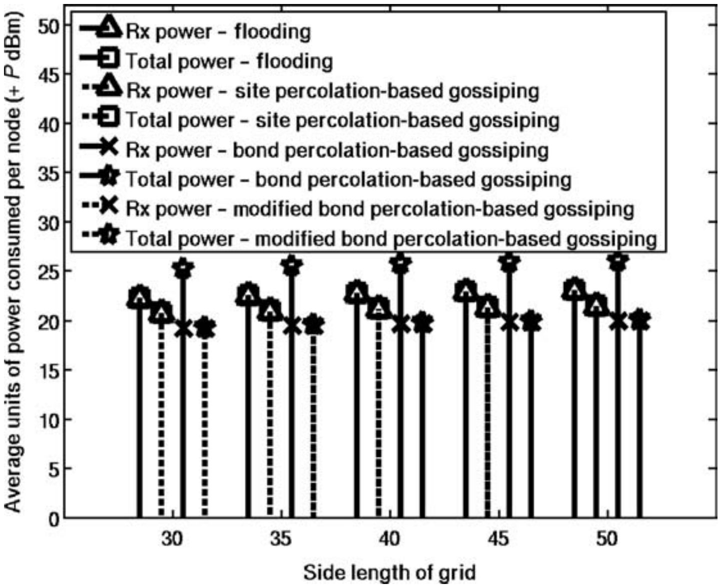

Figure 17. Average units of power consumed – grid topology.

percolation schemes consume more power but propagate broadcasts faster. Further, bond percolation consumes more energy and takes much longer to propagate the broadcast to all nodes. Figure 22 shows the energy–latency tradeoff for a grid topology with transmission range of one. We observe that site percolation consumes lower power and is faster than modified bond percolation mechanism. Flooding and bond percolation perform similar to that in Figure 21.

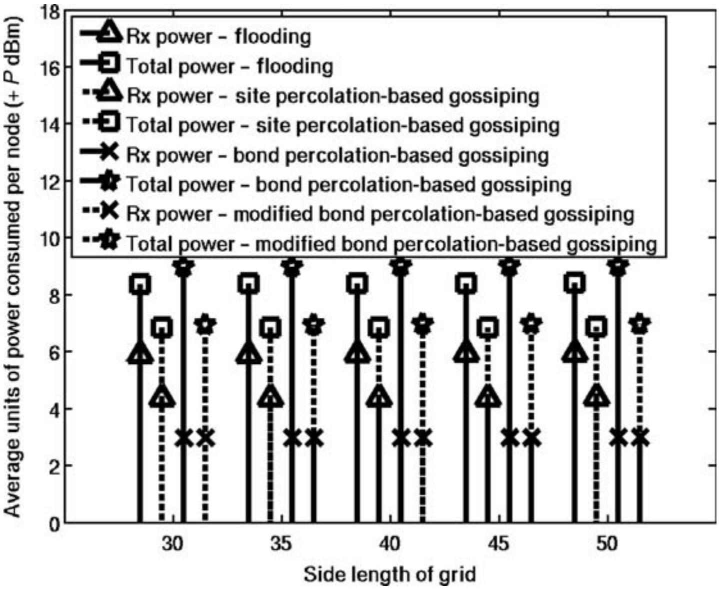

Figure 18. Average units of power consumed – grid topology, with Tx range = 1.

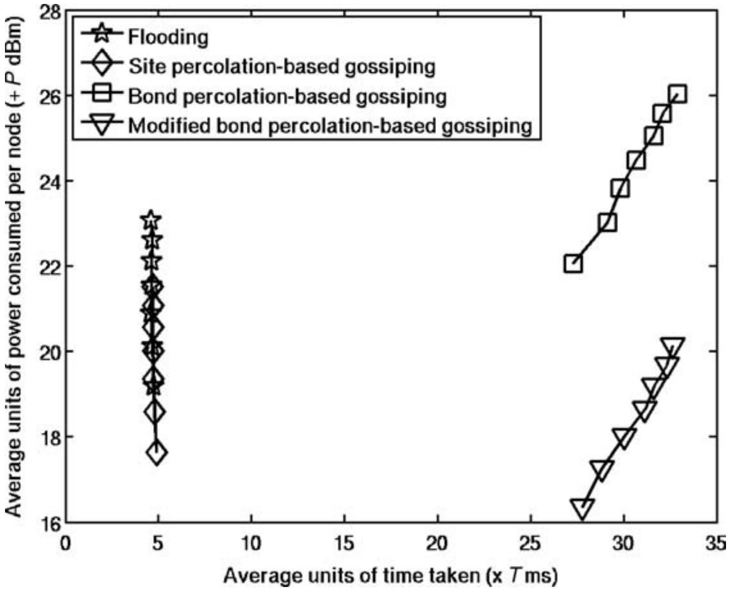

Figure 19. Latency vs. power consumed – random topology.

From the above discussion, we can conclude that the modified bond percolation (for dense networks) and site percolation (for sparse networks) are useful for applications that are required to be energy efficient, while site percolation and flooding are useful for applications that require low latency.

Before we conclude this section, we wish to compare the latency–power tradeoffs of the four schemes with medium- and low-power settings. Figure 23 shows the results for the

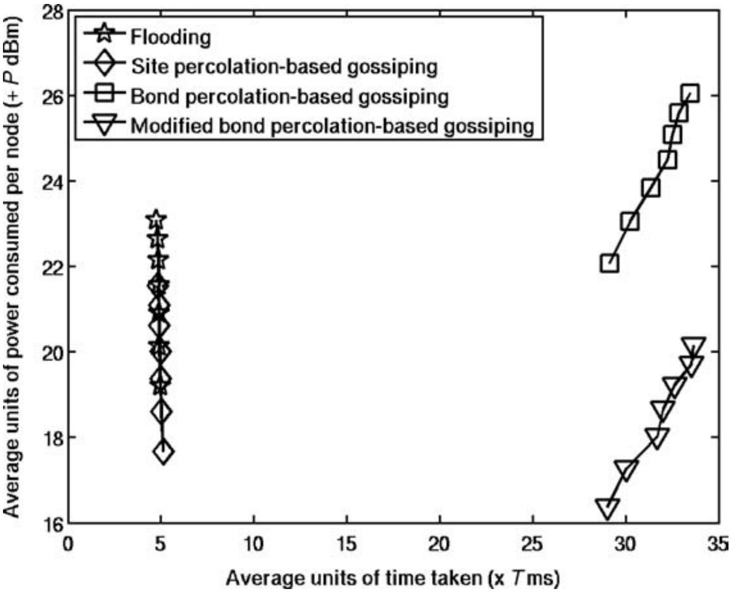

Figure 20. Latency vs. power consumed – clustered topology.

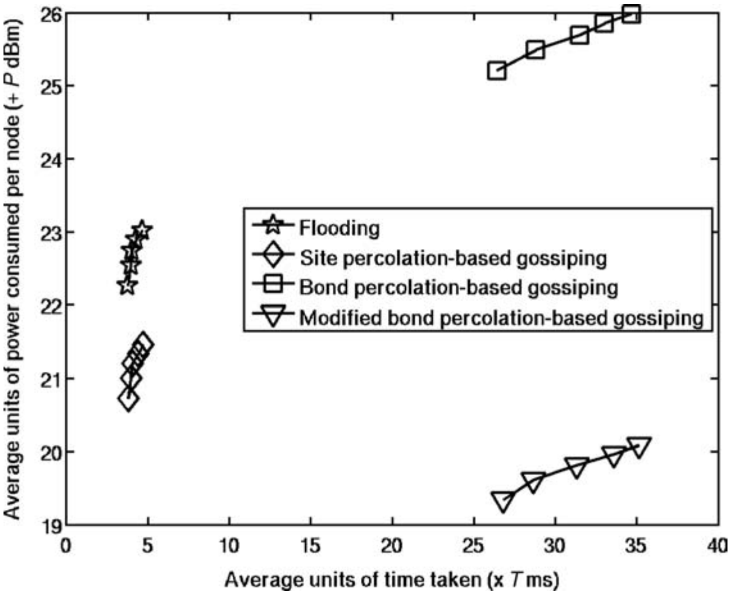

Figure 21. Latency vs. power consumed – grid topology.

random topology and Figure 24 shows the results for the grid topology. We have not provided the plots for clustered topology. However, the trend is similar to that of the random topology. From the plots we observe that, contrary to the fact that we use more power for transmission or reception in the medium-power setting, the average units of power consumed are smaller than that in the low-power setting. This trend is more clear in the case of the grid topology in Figure 24. Moreover, we also observe from Figure 24 that

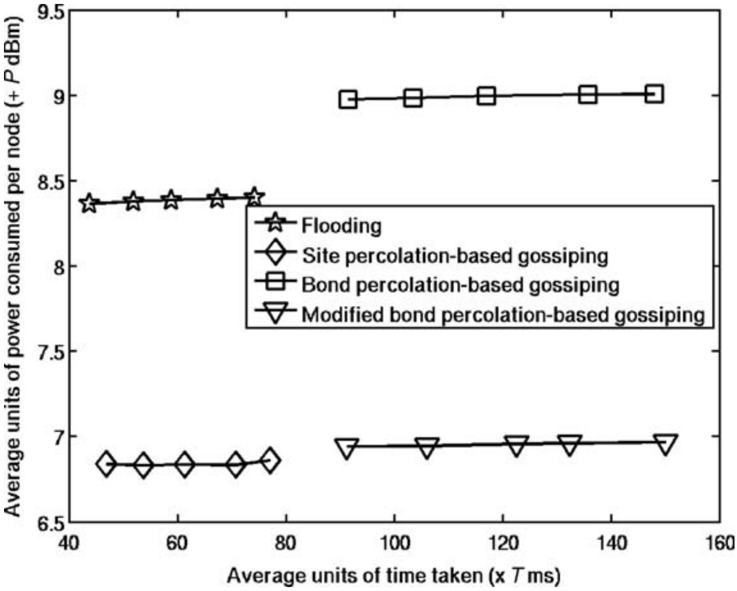

Figure 22. Latency vs. power consumed – grid topology, with Tx range = 1.

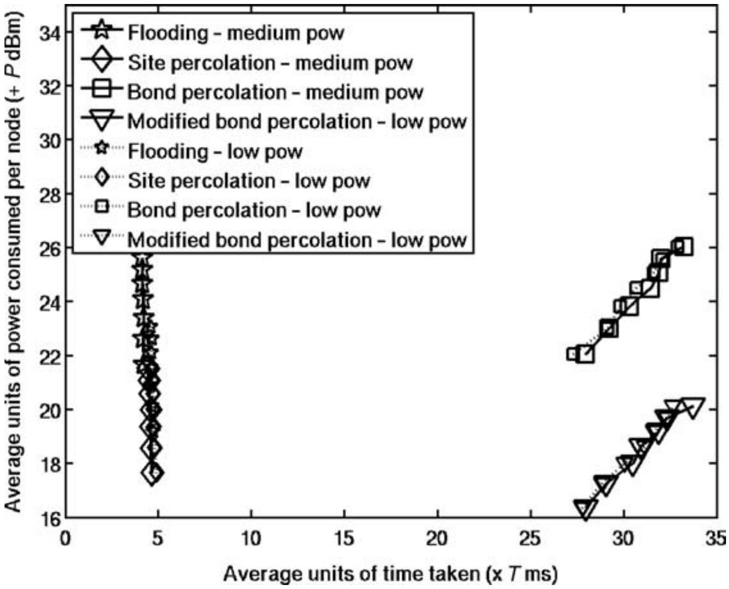

Figure 23. Latency vs. power consumed for medium- and low-power settings – random topology.

the medium-power setting case is faster than the low-power case. The reason for this behaviour is mainly because, when we use higher power, more nodes receive the broadcast message per transmission. The broadcast message can, therefore, reach the destination with fewer transmissions. This helps in reducing the overall power consumed and the latency in the network. The trend among the four broadcast schemes is the same irrespective of the power setting.

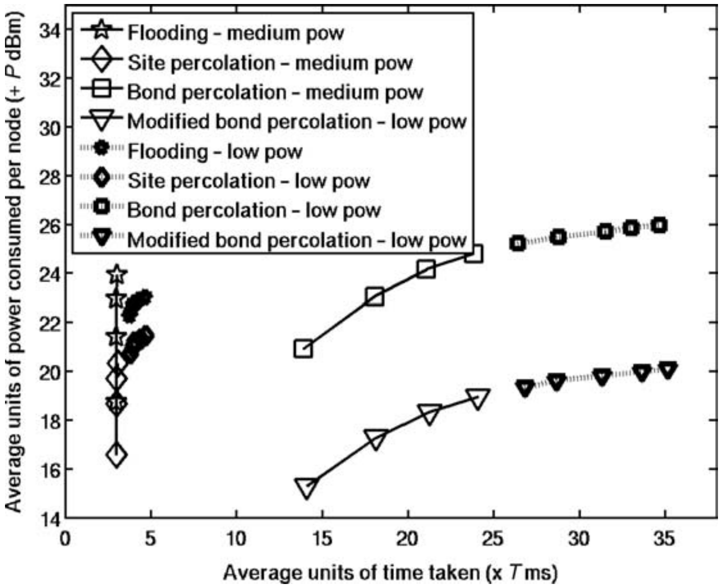

Figure 24. Latency vs. power consumed for medium- and low-power settings – grid topology.

4. Implications and caveats of our study for practical broadcast operations

4.1 Implications

We now discuss three different applications of the broadcast mechanisms in sensor networks and the implications of our study for these applications. Our findings are summarised in Table 1.

4.1.1 Code updates

The sensor nodes in a network may require frequent code updates for changing the tasks that these nodes have to execute [13]. During the code update process, each node with a stale code has to update itself and propagate the update to its neighbours. Such a code update process requires high degree of reliability, as the sensing task may fail otherwise. However, a code update process is not time sensitive, although reducing the broadcast propagation latency and the total energy consumption always helps.

From our analysis, we can see that by choosing suitable parameters, high reliability can be obtained using any of the four broadcast schemes. However, to reduce the energy consumption, percolation-based protocols may be preferable compared to flooding for dense networks. Flooding may be good if the network is sparse and clustered, or if the sensor nodes are not energy constrained, or if the code updates are infrequent.

4.1.2 Query and information propagation

Sensor networks need to be queried about their current status, for example, when average temperature is required to be calculated across a region for a histogram of temperature readings [15]. Such instantaneous querying requires timely propagation of query messages to the sensor nodes. Furthermore, the information sent by the sensor nodes in response to the query messages has to be routed back to the query source with soft real-time requirements. This means that, irrespective of the density of the network, flooding or site percolation-based gossiping is preferable compared to other approaches, for query propagation. Flooding may be more preferable if query correctness is highly sensitive to propagation time and energy is aplenty.

4.1.3 Membership management

Since sensor networks need to be self-sufficient, groups of sensor nodes may need to run a membership management protocol. For instance, this group may be comprised of all nodes that have a particular type of sensor (e.g. acoustic sensor, for tank detection) or sensors that

Table 1. Appropriate broadcast mechanism for each of the applications – further choices depend on tradeoffs involved in the application and the network layout.

Applications	Requirements	Flooding	Site percolation	Bond percolation	Modified bond percolation
Code updates	High reliability		√	√	√
Query propagation	Low latency	√	√		
Membership management	Low power		√		√

are in a specific geographical region. Each node in the group maintains a (partial) list of other nodes from the group which are currently awake. This information can then be used for multicast or aggregation within the group. The bootstrapping and maintenance of the membership list at each node can be maintained via broadcasts from each node as it goes to sleep and wakes up. The requirement for continuous and timely broadcasting means that energy consumption needs to be minimised – this points to modified bond percolation (for dense networks) and site percolation (for sparse networks) being the best choices for membership management.

4.2 Caveats

4.2.1 Effect of MAC collisions

Since we ignore MAC layer collisions and contentions in our simulations and analysis, the results are in a sense positively biased to flooding when compared to the different percolation approaches. This is because flooding generates many more messages, both overall and collision-wise. However, since our final results show that percolation approaches are preferable to flooding, we believe that the same conclusions will hold even if specific MAC protocols are taken into account.

4.2.2 Effect of on–off duty cycle

Our simulation results assume a common on–off duty cycle for all the sensor nodes and across all the broadcast approaches compared. However, in reality, the duty cycles may vary across all the sensor nodes. Furthermore, depending on the duty cycles, different broadcast mechanisms may be preferred. For instance, bond percolation is suitable for low-duty cycle sensor networks in which few neighbours wake-up simultaneously. The sender can, therefore, send the broadcast message to a subset of neighbours. On the other hand, site percolation will be suitable for sensor networks with faster duty cycles, as a single broadcast is intended for all the neighbours of the sender. The results compared above may thus be different depending on the duty cycles. A common duty cycle for all the mechanisms is, however, chosen to provide a common base for comparison.

5. Conclusion

In this paper, we have analysed four fundamental percolation-based broadcast mechanisms in sensor networks, namely flooding, site percolation-based gossiping, bond percolation-based gossiping and modified bond percolation algorithms. Our analytical and simulation results show that flooding and site percolation schemes have low latency. For dense networks, modified bond percolation consumes the lowest energy per broadcast, while site percolation consumes the lowest energy for sparse networks. Furthermore, our analysis also shows that flooding involves a larger number of messages during the broadcast propagation. These conclusions hold for three types of topologies that practical deployments use – random, clustered and grid. We also identified the tradeoffs of these broadcast alternatives for code updates, query propagation and membership management.

Acknowledgements

The authors would like to acknowledge NSF for supporting this work in part by the CAREER grant CNS 04-48246, in part by NSF ITR grant CMS 04-27089 and in part by the grant CNS 05-19817.

References

[1] A. Ahmed, H. Shi, and Y. Shang, *A survey on network protocols for wireless sensor networks*, ITRE (2003), pp. 301–305.

[2] M. Akdere, C. Bilgin, O. Gerdaneri, I. Korpeoglu, O. Ulusoy, and U. Cetintemel, *A comparison of epidemic algorithms in wireless sensor networks*, Comput. Commun. 29(13–14) (2006), pp. 2450–2457.

[3] J. Al-Karaki and A. Kamal, *Routing techniques in wireless sensor networks: A survey*, IEEE Wireless Commun. 11(6) (2004), pp. 6–28.

[4] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D.A. Maltz, and M. Zhang, *Towards highly reliable enterprise network services via inference of multi-level dependencies*, ACM SIGCOMM, 2007.

[5] CITRIS, *Great duck island: Sensor deployment* (2007), http://webs.cs.berkeley.edu/related.html.

[6] T. Damvan and K. Langendoen, *An adaptive energy-efficient mac protocol for wireless sensor networks*, SenSys, 2003.

[7] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker, *Complex behavior at scale: An experimental study of low-power wireless sensor networks*, UCLA Computer Science Technical Report, 2002.

[8] G. Grimmett, *Percolation*, Springer-Verlag, Berlin, 1989.

[9] Z. Haas, J. Halpern, and L. Li, *Gossip-based* ad hoc *routing*, IEEE Infocom, 2002.

[10] W. Heinzelman, J. Kulik, and H. Balakrishnan, *Adaptive protocols for information dissemination in wireless sensor networks*, ACM Mobicom (1999), pp. 174–185.

[11] C. Intanagonwiwat, R. Govindan, and D. Estrin, *Directed diffusion: A scalable and robust communication paradigm for sensor networks*, ACM Mobicom, 2000.

[12] Q. Jiang and D. Manivannan, *Routing protocols for sensor networks*, CCNC (2004), pp. 93–98.

[13] P. Levis, N. Patel, D. Culler, and S. Shenker, *Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks*, USENIX NSDI (2004), pp. 15–28.

[14] L. Luo, T.F. Abdelzaher, T. He, and J.A. Stankovic, *Design and comparison of lightweight group management strategies in EnviroSuite*, DCOSS (2005), pp. 155–172.

[15] S. Madden, M. Franklin, J. Hellerstein, and W. Hong, *Tag: A tiny aggregation service for* ad-hoc *sensor networks*, OSDI 02, 2002.

[16] S. Meguerdichian, S. Slijepcevic, V. Karayan, and M. Potkonjak, *Localized algorithms in wireless* ad-hoc *networks: Location discovery and sensor exposure*, ACM Mobihoc, 2001.

[17] M. Miller, C. Sengul, and I. Gupta, *Exploring the energy–latency tradeoff for broadcasts in energy-saving sensor networks*, ICDCS, 2005.

[18] S. Nath, P.B. Gibbons, S. Seshan, and Z.R. Anderson, *Synopsis diffusion for robust aggregation in sensor networks*, International Conference on Embedded Networked Sensor Systems, 2004.

[19] J. Polastre, J. Hill, and D. Culler, *Versatile low power media access for wireless sensor networks*, SenSys, 2004.

[20] J. Polastre, R. Szewczyk, C. Sharp, and D. Culler, *The mote revolution: Low power wireless sensor network devices*, http://webs.cs.berkeley.edu/papers/hotchips-2004-motes.ppt

[21] S. Ray, R. Ungrangsi, F.D. Pellegrini, A. Trachtenberg, and D. Starobinski, *Robust location detection in emergency sensor networks*, IEEE Infocom, 2003.

[22] L. Schwiebert, S. Gupta, and J. Weinmann, *Research challenges in wireless networks of biomedical sensors*, ACM Mobicom (2000), pp. 151–161.

[23] W. Ye, J. Heidemann, and D. Estrin, *An energy-efficient mac protocol for wireless sensor networks*, IEEE Infocom, 2002.

[24] Y. Zou and K. Chakrabarty, *Sensor deployment and target localization based on virtual forces*, IEEE Infocom, 2003.