

# TARA: Topology-Aware Resource Adaptation to Alleviate Congestion in Sensor Networks

Jaewon Kang, *Student Member, IEEE*, Yanyong Zhang, *Member, IEEE*, and  
Badri Nath, *Member, IEEE*

**Abstract**—Network congestion can be alleviated either by reducing demand (traffic control) or by increasing capacity (resource control). Unlike in traditional wired or other wireless counterparts, sensor network deployments provide elastic resource availability for satisfying the fidelity level required by applications. In many cases, using traffic control can violate fidelity requirements. Hence, we propose the use of resource control: increasing capacity by enabling more nodes to become active during periods of congestion. However, a naive approach to increase resources without a careful consideration of the type of congestion, traffic pattern, and network topology will make the situation worse. In this paper, we present TARA, a topology-aware resource adaptation strategy to alleviate congestion. The core of TARA is our capacity analysis model, which can be used to estimate capacity of various topologies. Detailed performance results show that TARA can achieve data delivery rate and energy consumption that is close to an ideal offline resource control algorithm.

**Index Terms**—Congestion control, sensor networks, resource control, fidelity, energy efficiency, simulations.

## 1 INTRODUCTION

A broad range of remote sensing applications are being enabled by the rapid development of sensor networks. Typical remote sensing applications include target recognition and tracking, forest fire detection, structural and habitat monitoring, and traffic condition surveillance. Before the target events occur, sensor nodes report data at a lower rate (for example, heartbeat messages) to save energy. As soon as these events are detected, however, a high reporting rate is necessary to generate sufficient data to accurately depict the phenomena of interest. As a result, the lifetime of a sensor network alternates between long periods of low traffic volume (referred to as *dormant state*) and short periods of high traffic volume (referred to as *crisis state*). For most applications, a long dormant state dominates the network lifetime. Unlike wired networks or other wireless counterparts, it is not viable to make all battery-powered nodes active all the time in sensor networks because it bounds the network lifetime by the lifetime of a battery (which is in the order of months) attached to each sensor node. Thus, a popular way of conserving energy is to keep only a minimum number of nodes active (for example, the nodes on the routing path) during dormant periods. This strategy, though energy efficient, can cause a serious problem—as soon as the network enters a crisis state, congestion is likely to occur because the data rates may

exceed the capacity available from the currently active nodes. Once congestion occurs, many packets that contain valuable data may be dropped. Although “traffic control” strategies that reduce the incoming traffic have been effective to alleviate congestion in traditional networks, they are unsuitable for our purposes because reducing source traffic during a crisis state is unacceptable. The data generated during a crisis state are of utmost importance, and loss of these data can defeat the very purpose of deploying an unattended sensor network. To meet this so-called fidelity requirement, we propose to use the approach of “resource control”—increasing capacity by turning on more resources to accommodate high incoming traffic during the crisis state.

Activating (or awakening) additional resources in response to increased data volume is a challenging problem. The main challenge stems from the fact that resource control strategies require not only local knowledge, but also knowledge about the end-to-end topology, as will be shown in this paper. To address this issue, we propose TARA, a topology-aware resource adaptation strategy, which activates appropriate sensor nodes whose radio is off (that is, sleeping nodes) to form a new topology that has just enough capacity to handle the increased traffic, satisfying both fidelity (or accuracy) and energy requirements at the same time. TARA’s distinct ability to find an optimal topology in response to congestion under the fidelity and energy requirements is enabled by our capacity analysis model that can efficiently estimate the capacity of various network topologies using a graph-theoretic approach.

In the rest of this paper, we first formulate the targeted problem by comparing the frameworks of traffic control and resource control and explain the need of topology-aware resource control in Section 2. Related work is covered in Section 3. Section 4 presents the capacity analysis model and capacities of several typical topologies. The distributed topology-aware resource control protocol is presented in detail in Section 5. Section 6 shows the performance

• J. Kang and B. Nath are with the Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, NJ 08854.  
E-mail: {jwkang, badri}@cs.rutgers.edu.

• Y. Zhang is with the Department of Electrical and Computer Engineering, Rutgers University, 94 Brett Road, Piscataway, NJ 08854.  
E-mail: yyzhang@ece.rutgers.edu.

Manuscript received 17 Jan. 2006; revised 1 June 2006; accepted 3 Aug. 2006; published online 8 Jan. 2007.

Recommended for acceptance by S. Olariu.

For information on obtaining reprints of this article, please send e-mail to: [tpds@computer.org](mailto:tpds@computer.org), and reference IEEECS Log Number TPDS-0015-0106.  
Digital Object Identifier no. 10.1109/TPDS.2007.1030.

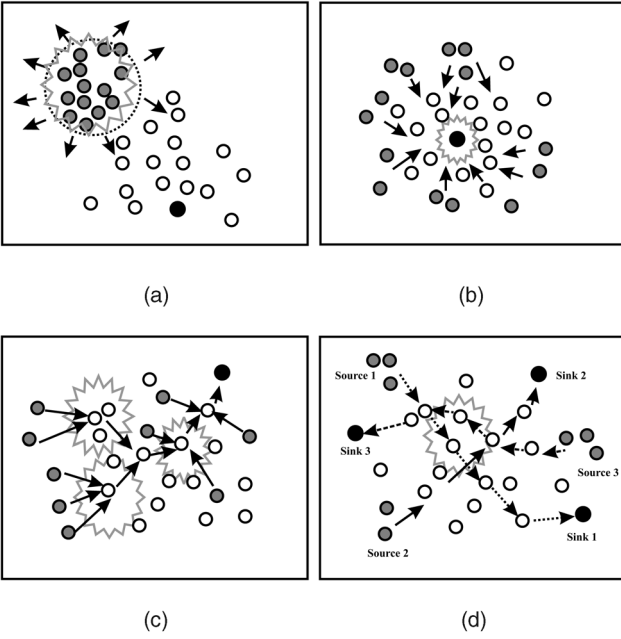


Fig. 1. Three hot spot scenarios. (a) Source hot spot. (b) Sink hot spot. (c) Intersection hot spot. (d) Intersection hot spot.

evaluation of TARA. Finally, Section 7 provides the concluding remarks.

## 2 PROBLEM FORMULATION

### 2.1 Congestion in Sensor Networks

We have identified three typical hot spot scenarios within sensor networks:

- *Source hot spot.* As soon as an event takes place, it will be detected by all the nodes whose sensing ranges (with radius  $r$ ) cover the event spot. These nodes will act as *sources*, reporting their observations to the sink(s). Since a node's radio range ( $R$ ) is usually greater than its sensing range (for example, many realistic deployments assume that  $R > 2r$  [26]), these sources are likely to be within each other's radio range as well. If all these source nodes start sending packets at a high rate to the sink, a hot spot will quickly form, and a large number of packets will be dropped around the event spot (Fig. 1a).
- *Sink hot spot.* In a crisis state, sink nodes (and the nodes around them) handle a high traffic volume (Fig. 1b) [16]. Hence, the batteries of these nodes around sinks will be drained quickly, making the sinks unreachable from the rest of the network.
- *Intersection hot spot.* The presence of multiple sources and multiple sinks results in more than one flow intersecting with each other. Due to the traffic merge at the intersection nodes, they can also become hot spots (see Figs. 1c and 1d), referred to as intersection hot spots. These intersecting flows can either share segments of routing paths (illustrated in Fig. 1c) or cross each other's routing paths (illustrated in Fig. 1d).

Although each of the three scenarios is harmful to the operations of sensor networks, some of them can be




Fig. 2. Frameworks of traffic control and resource control. (a) Traffic control. (b) Resource control.

avoided through careful planning at deployment stage. For instance, source hot spots can be eliminated by allowing only a small number of nodes to report to the sink [17], [23]. This will not degrade network services because these nodes are likely to report highly correlated observations due to their geographic proximity. At the same time, one way of alleviating sink hot spots is to deploy multiple sinks that are uniformly scattered across the sensor field and then balance the traffic among them [16]. These types of planning are feasible because sensor applications often have good estimations of their peak load (during crisis state) and node deployment density. On the other hand, intersection hot spots are far more challenging because it is very difficult to predict the intersection points due to the dynamic nature of sensor networks. Hence, intersection hot spots cannot be avoided at deployment time, but demand runtime countermeasures. This paper will thus focus on alleviating intersection hot spots.

### 2.2 Traffic Control versus Resource Control

In this section, we present the rationale behind traffic control and resource control schemes.

- *Traffic control.* The approach of controlling traffic for congestion avoidance has been extensively studied, mostly through the active queue management (AQM) in wired networks [5]. The two main criteria for traffic control policies are *resource utilization* and *fairness*. We next illustrate how these two criteria are met using the traffic control algorithm between two flows (presented in [3]). Fig. 2a depicts the offered traffic of both flows at node  $i$  ( $T_{1i}(t)$  on the x-axis and  $T_{2i}(t)$  on the y-axis). Suppose the total resource provisioning at node  $i$  is  $R_i$ , which is constant over time. Then, the efficiency line, also referred to as *resource line*, denotes the scenarios where the available resource is equal to the total traffic volume, that is,  $T_{1i}(t) + T_{2i}(t) = R_i$ . The area above the resource line, that is,  $T_{1i}(t) + T_{2i}(t) > R_i$ , corresponds to *overloaded* scenarios; the area below the resource line denotes *underloaded* scenarios. The dotted line in Fig. 2a, with  $T_{1i}(t) = T_{2i}(t)$ , denotes a fair resource allocation between the two flows. The resource line and fairness line intersect at  $\langle \frac{R_i}{2}, \frac{R_i}{2} \rangle$ , which corresponds to the optimal scenario, wherein the two flows can fully and equally utilize the available

resources. The goal of traffic control is to tune the offered load of all flows to approach the optimal point to ensure that the resource is fully utilized while each flow is allocated a fair share of resource.

- *Resource control.* The rationale behind resource control strategies is distinctively different from that of traffic control strategies. Traffic control schemes assume a fixed resource provisioning,  $R_i$ , at node  $i$ . However, this constraint does not hold in sensor networks. The available resource (or capacity) at node  $i$  is elastic due to its dynamic duty cycle, interference, and so forth, thereby represented as a time-varying function  $R_i(t)$ . As a result, resource control schemes seek to satisfy the fidelity level requirement of each flow, even during congestion, by assigning additional resources to the flow that has a higher fidelity level requirement without taking resources away from other flows as illustrated in Fig. 2b.

Traffic control strategies are effective in traditional wired networks because it is difficult, if not impossible, to adapt the available resources between any source-destination pair. However, traffic control strategies are not suitable for relieving intersection hot spots in sensor networks for the following reason: Whenever an event occurs, we first need to prevent source hot spots by limiting the number of nodes that report to the sink, leaving only a small number of sources that send data of great value to the applications. Failure to ship these data to the sink can put the motivation of deploying the sensor network at risk. In addition, there is usually an abundance of resource in sensor networks (unlike its wired or other wireless counterparts) that is not utilized during dormant periods due to the redundancy in sensor network deployment [2], [26], [27] and the nature of wireless communication.

### 2.3 Topology-Aware Resource Control

The main challenge of resource control stems from the fact that resource control strategies require not only local knowledge, but also knowledge about the end-to-end topology. For example, a naive strategy is to activate all the nodes and create multiple paths to enable parallel forwarding. Nonetheless, this approach can substantially shorten the lifetime of the network. Another strategy would be to activate a random number of nodes that are outside the congested area to detour packets, but there is no guarantee that the new topology can offer larger capacity than the existing configuration, let alone satisfy the application fidelity requirement. As a result, the goal of a resource control scheme is to increase the fidelity level observed by the deployed application while minimizing the scarce energy in the network. To meet the fidelity and energy requirements, an efficient resource control scheme should consider traffic rate, congestion level, and most important of all, network topologies.

## 3 RELATED WORK

Most of prior works have focused on traffic control. A guideline of congestion control in sensor networks was first

given in [19]. The authors suggest that congestion control must be based on not only the network capacity, but also the application fidelity requirements. CODA [23] presents the first detailed study on congestion detection and avoidance in sensor networks. In CODA, as soon as a node detects congestion, it broadcasts a backpressure message upstream. An upstream node will thus throttle the traffic volume. Sankarasubramaniam et al. [17] proposed an event-to-sink reliable transport (ESRT) protocol, which can serve as a congestion control protocol. In ESRT, the sink reduces the traffic of all sources during congestion. In [10], Hull et al. studied three congestion control techniques: hop-by-hop flow control, limiting source rate, and a prioritized medium access control (MAC). They showed that network congestion can be greatly alleviated when these three traffic control techniques operate in concert rather than in isolation. Ee and Bajcsy [4] proposed a distributed congestion control scheme based on hop-by-hop automatic repeat request (ARQ) in many-to-one routing scenarios. Woo and Culler [25] proposed to alleviate congestion by assigning bandwidth proportionally. Yi and Shakkottai [31] proposed a hop-by-hop congestion control scheme that allocates bandwidth to various users in a fair manner. They show a hop-by-hop traffic control scheme push-backs and spread congestion over space, leading to scattered small peak loads. SPEED [9], originally developed for real-time communication, handles congestion by throttling or rerouting the incoming traffic around the hot spot. The rerouted path, however, may not have a larger end-to-end channel capacity to accommodate the incoming traffic, leading to congestion.

Resource control has received little attention. In [12], [20], the authors propose congestion-adaptive routing schemes in ad hoc networks. We have, however, shown by simulation in Section 6 that these schemes cannot effectively alleviate congestion in some congestion scenarios due to their topology unawareness. Several multipath routing protocols [7], [13], [15], [18], [22], [30] are developed in the context of reliability, load balancing, energy efficiency, and failure recovery rather than congestion control.

## 4 CAPACITY ANALYSIS MODEL

The core of TARA involves a capacity analysis model, which can efficiently estimate the end-to-end throughput of different topologies. The capacity analysis model is formulated as a graph-coloring problem, and we have shown that the model results agree with simulation results and actual experimentation results.

### 4.1 Model Description

The capacity of a given topology is defined by the maximum throughput (that is, packet delivery rate) that can be observed by the sink(s). If there were no interferences between links, the capacity of a topology would be the same as the maximum throughput achievable by unlimited unidirectional transmissions in a one-hop topology. The interference between links, however, makes the overall throughput much smaller than the one-hop capacity. The objective of TARA's capacity analysis model is to capture the degree of interference of a given topology.




Fig. 3. Illustration of the capacity analysis model. (a) Topology. (b) Spatial interference graph. (c) Spatial reuse graph. (d) Colored graph. (e) Time frame assignment.

The main idea behind TARA's capacity analysis model is to map the problem of capacity estimation into a suitable graph-coloring problem [24]. Figs. 3a, 3b, 3c, 3d, and 3e illustrate this process. Fig. 3a presents a simple network topology with seven nodes. Suppose the traffic volume from each source, B and G, is  $T$  packets/second, then, the traffic for each link is as marked above the link. To produce the maximum throughput, suppose that there exists an optimal schedule that can avoid collisions within the entire topology. Under this optimal schedule, suppose that the sink receives  $m$  packets every  $n$  time frames.<sup>1</sup> Then, the capacity of this topology is calculated as  $\frac{m}{n} C^{max}$ , where  $\frac{m}{n}$  is called the *capacity fraction*, and  $C^{max}$  denotes the one-hop capacity. The capacity fraction of a one-hop topology is 1, because in every time frame, one packet is received by the sink. In order to calculate the capacity fraction using graph-theoretic methods, we construct the *spatial interference graph* as depicted in Fig. 3b, where each vertex denotes the corresponding wireless transmission in Fig. 3a, and the edges between two vertices indicate that these two transmissions are within each other's interference range. That is, any two wireless links with an edge connecting them cannot transmit concurrently under the optimal schedule. In the spatial interference graph, it is necessary to replicate some wireless links (such as IJ in this example) when their traffic load is a multiple of  $T$ , as shown in Fig. 3a. In particular, the link whose traffic is  $mT$  should be repeated  $m$  times. This is because two packet transmissions from I to J should be scheduled per one packet transmission on other links to prevent congestion around I in Fig. 3a. Calculating the capacity fraction now boils down to assigning a coloring to the spatial interference graph, where each color corresponds to the particular time frame in which each link transmits. In Fig. 3d, five colors are needed to color the vertices, and sink J receives two packets every five time frames. Thus, the capacity fraction of this topology is  $\frac{2}{5}$ .

It is well known that the problem of finding the minimum number of colors for a graph is NP-complete [24]. However, by using the following theorems, we can obtain an upper bound for colorability: 1) If  $G$  is a simple graph with the largest vertex degree  $d$ , then  $G$  is  $(d + 1)$  colorable. 2) If  $G$  is a simple connected graph and not a complete graph, and if the largest vertex degree of  $G$  is  $(d \geq 3)$ , then  $G$  is  $d$ -colorable. It should be noted that a spatial interference graph is always connected because there will be no link that does not interfere with any other link. In addition, it is very likely that the graph is not complete due

to the limited radio range, and the largest vertex degree will be greater than or equal to 3, especially in the congested area. As a result, these theorems, especially the second one, give a quick estimation on the number of time frames needed for a certain topology. Applying these methods are meant to serve as guidelines for analyzing the best case capacity capabilities of the sensor network. In reality, they provide a bound, and one can find a better coloring than what these theorems promise.

In addition to theorems that give an upper bound, many heuristic solutions have been proposed to color a certain graph using the minimum number of colors for different applications. To solve our problem of estimating capacity fraction, we also propose the following heuristic solution:

- Construct the spatial reuse graph, which is the complement of the spatial interference graph, where two vertices are connected by an edge when they do not interfere with each other, as shown in Fig. 3c.
- Sort all the vertices in the ascending order of their degrees in the spatial reuse graph. For example, after sorting the vertices in Fig. 3c, we have CD, DI, HI, IJ, IJ, GH, and BC, where the tie is broken alphabetically.
- Start from the first vertex in the list and find the largest complete subgraph that contains this vertex. All the vertices in this subgraph comprise a *concurrent transmission set*, which includes all the links that can transmit within the same time frame. Then, remove all these vertices and the corresponding edges from the graph and repeat the process until the graph is empty. In this example, we have the following five concurrent transmission sets: {CD, GH}, {DI}, {HI, BC}, {IJ}, and {IJ}. The reason why we start forming concurrent sets from the vertices with least number of edges is because those nodes are the most constrained.

The number of concurrent sets is the minimum number of time frames needed to deliver a packet from each source to the sink. In this example, the result of our heuristic matches the result from the second theorem. One side product of using this heuristic approach is a packet schedule based on the time frame assignment shown in Fig. 3e. The time frame number for each transmission is the color number in Fig. 3d.

## 4.2 Estimating Capacity for a Large-Scale Network

The time complexity of our model is greatly affected by the number of nodes in the topology. As a result, it may be hard

1. One time frame corresponds to the time interval for a node to transmit one packet to its immediate neighbor.




Fig. 4. Finding a bottleneck zone. (a) Two merging flows, each with traffic load  $T$ . (b) A bottleneck link IJ and its bottleneck zone.

to apply our model over a large topology. To deal with this, we take the viewpoint that *the throughput of a topology is limited by the throughput of the bottleneck link(s)*. Therefore, one only needs to calculate the capacity fraction for the portion of the network topology that contains the bottleneck link(s), referred to as bottleneck zone. Extracting the bottleneck zone from a topology involves two steps: identifying the bottleneck link and then identifying all the links that interfere with it.

Usually, a bottleneck link handles higher traffic loads and/or experiences a higher degree of interference. To capture these two factors, we introduce the term *congestion sum* of a link, which is the sum of that link's traffic volume and the traffic volumes of all the links that cannot transmit concurrently with this link. The bottleneck link, therefore, has the largest congestion sum. Fig. 4a shows a simple topology in which two source nodes, A and E, each generate a traffic of  $T$  packets/second. The traffic load of each subsequent link is shown in Fig. 4a. To calculate the congestion sum, for each link, we need to identify its interfering links. For example, if B and D transmit at the same time, then C will not receive correctly. For another example, G's data packet may collide with I's ACK packet (to D) at node H. As a result, even under the assumption that the radio interference range equals the communication range, link DI interferes with the following links: BC, CD, GH, HI, IJ, and JK, which are also highlighted in Fig. 4a. By adding up the traffic volumes of all the interfering links, one can calculate the congestion sum of a link. Fig. 4b shows the congestion sum of all the links, and the bottleneck links are IJ and JK. Between these two, we focus on the one closer to the source to infer the capacity fraction of the entire topology. Therefore, we pick IJ and isolate its bottleneck zone, which is shown in Fig. 4b.

After the step above, we can focus on the appropriate bottleneck zone, which is much smaller than the entire topology, and apply our model to calculate the capacity for the bottleneck zone, which is also the capacity of the entire topology. Identifying bottleneck links from a large and complex topology, however, will be explored as part of our future work.

### 4.3 Capacity of Typical Hot Spot Topologies

The model can be used to estimate capacities for arbitrary topologies, under different system assumptions. Next, we validate our model results against both simulation and actual experimental results by studying how to increase capacities for string, merging, and crossing topologies. Through these studies, we have learned four valuable lessons, which can guide the design of TARA. For each topology scenario, we have measured the capacity by using the following methods:

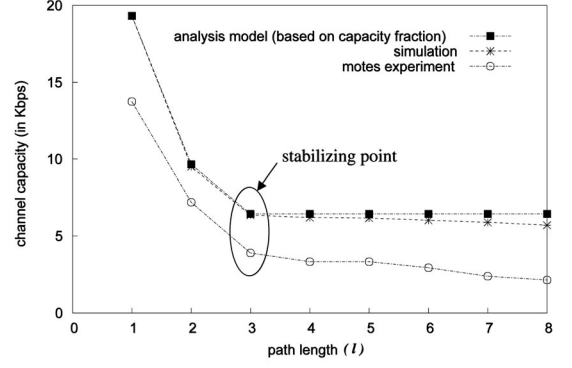


Fig. 5. End-to-end channel capacity of a string topology.

- *Mote experimentation.* The Mica2 motes (MPR400CB) [23], [25], running TinyOS v1.0, are used for sensor nodes. The motes are carefully placed so that only adjacent ones can communicate with each other. The interference range is irregular and may span multiple hops. The raw data rate of Mica2 radio is 38.4 Kbps. The motes adopt a simple MAC protocol that samples the channel activity several times before sending a packet.
- *NS-2 simulation.* We configure the NS-2 simulator by using the same parameters as in Mica2 motes (that is, raw radio data rate and packet size). We set up the network in such a way that only adjacent nodes can communicate with each other. The interference range is within two hops. A simple carrier sense multiple access (CSMA) MAC protocol, with RTS/CTS, is used. ACKs are sent to acknowledge the reception of packets.
- *Our model.* The capacity from our model has the format of  $\varepsilon C^{max}$ , where  $\varepsilon$  is the capacity fraction, and  $C^{max}$  is the one-hop channel capacity obtained from simulations. In our model, we assume that ACKs are used.

The capacity under a topology is measured by varying the source's sending rate until the network is saturated. We would like to emphasize that our focus is not to match the absolute capacity figures from different methods, which is impractical given the extreme dynamic nature of a wireless experiment. However, we would like to examine whether the trends from different methods agree with each other.

#### 4.3.1 String Topology

In this set of experiments, we study how the capacity of a string topology<sup>2</sup> varies with different hop counts between the source and the sink, that is, path length  $l$ . The results are presented in Fig. 5. For the string topology, our analytical model has the following capacity fractions:

$$\varepsilon = \begin{cases} 1 & \text{if } l = 1 \\ 1/2 & \text{if } l = 2 \\ 1/3 & \text{if } l \geq 3. \end{cases} \quad (1)$$

$C^{max}$  is set to the one-hop throughput obtained using the simulations. Among the three methods, mote experiments

2. All nodes are serially placed with a source on one end and a sink on the other end.




Fig. 6. Aggregate end-to-end channel capacity of a two-flow merging topology.

have lower capacities due to the simple MAC protocol and larger degree of interference. However, all three methods demonstrate the similar trend: as  $l$  increases from 1 to 3, the capacity quickly drops, but stabilizes as  $l$  further increases. Similar observations have also been found in [14], where the authors have pointed out that the capacity fraction of a string topology is constant if local communications dominate and path lengths do not grow arbitrarily with the network dimensions (which is different from the assumption in [8]). Similarly, we have an implicit assumption that path lengths in sensor networks stay more or less constant regardless of the network size, which can be achieved by deploying more sinks as the network size grows. Using the model, we can quickly draw the following conclusion:

*Lesson 1.* Minimizing the path length in a string topology does not increase the capacity if the resulting topology has a path length of more than two hops.

The results also indicate that even though a string topology has a large hop count, it can provide a certain level of channel capacity (denoted as  $C^{min}$ ) since the capacity fraction is stabilized. Therefore, we have the following approach of alleviating congestion:

*Lesson 2.* If the node whose incoming traffic volume is less than  $C^{min}$  experiences congestion due to interference with other flows, the congestion can be eliminated by rerouting the incoming traffic onto the *noninterfered* path.

#### 4.3.2 Merging Topology

A merging topology is more complicated than a string topology (illustrated in Fig. 4), and an  $n$ -flow merging case (assuming that they merge at the same point) can be characterized by the following  $n+1$  parameters:  $l_1, \dots, l_n$ , and  $h$ , where  $l_i$  is the path length for the  $i$ th flow, and  $h$  is the number of hops between the merging point and the sink. In Fig. 4a, for example,  $l_1$ ,  $l_2$ , and  $h$  are 8, 8, and 4, respectively. Since the merging point is the bottleneck of the topology,  $h$  plays a more important role than individual path lengths, especially when each individual flow length is reasonably long. Our analysis model derives the following capacity fractions:

$$\varepsilon = \begin{cases} n/(n+1) & \text{if } h = 0 \\ n/(2n+1) & \text{if } h = 1 \\ n/3n = 1/3 & \text{if } h \geq 2. \end{cases} \quad (2)$$




Fig. 7. Four scenarios of two crossing flows. (a) I. (b) II. (c) III. (d) IV.

Fig. 6 presents the results from the model, simulations, and mote experiments for a two-flow merging case. In this set of experiments, we have  $l_1 = l_2 = 8$ , and we vary  $h$ . When two flows merge at the sink, that is,  $h = 0$ , we have maximum capacity. Among the three, the capacity from our model is the best, and the capacity from the mote experiments is the worst. This is because actual experiments have more interference than simulations, which, in turn, have more complicated interference than the model. Though the absolute numbers are different, the trends from all three methods are the same. When  $h$  increases from 0 to 2, the resulting capacity quickly decreases, and it stays constant after  $h$  exceeds 2. From these observations, we have the following lesson to boost the capacity for a merging topology:

*Lesson 3.* The capacity of a merging topology can be increased by moving the merging point within two hops away from the sink.

#### 4.3.3 Crossing Topology

A crossing topology consists of multiple flows that have distinct sinks. A two-flow example is shown in Fig. 7a. Unlike the merging case, moving the crossing point, as shown in Fig. 7b, will not increase the capacity. Instead, one may want to have multiple paths for either of the flows and split the traffic of that flow onto these paths, as shown in Figs. 7c and 7d. These points are further validated by the calculated capacity fractions of the following scenarios:  $\frac{2}{5}$ ,  $\frac{2}{5}$ ,  $\frac{1}{2}$ , and  $\frac{6}{11}$ . We have measured the capacities of these four scenarios by using all three methods and presented them in Fig. 8. The percentage of capacity increase in scenarios 2, 3, and 4, with respect to scenario 1 are shown on top of each bar. From the results, we observe that the capacity of the crossing topology can be increased by having multiple paths in either flow. This is because the resulting incoming traffic volume at each crossing node is reduced, which are  $2T$ ,  $2T$ ,  $\frac{3}{2}T$ , and  $\frac{4}{3}T$  in the four scenarios above, when each source generates a traffic volume of  $T$ . Therefore, we have the following conclusion:




Fig. 8. Channel capacity of the four scenarios shown in Fig. 7.

*Lesson 4.* To increase the capacity of a crossing topology, at least one flow should have multiple paths and split its traffic onto these paths.

## 5 TOPOLOGY-AWARE RESOURCE ADAPTATION SCHEME

Building upon the capacity analysis tool, we develop TARA, a topology-aware resource adaptation protocol, which can adapt network resources based on the congestion level. Hence, it can serve the dual purposes of alleviating congestion during crisis states and conserving energy during dormant states.

### 5.1 Increasing Resources during Crisis States

The framework of TARA is illustrated in Fig. 9, where the hot spot is around node B. As soon as the hot spot node detects that its congestion level is above the upper watermark, it needs to quickly locate two important nodes: the distributor (node G) and the merger (node J). Then, a detour path can be established, starting at the distributor and ending at the merger. As suggested by their names, the distributor distributes the incoming traffic between the original path and the detour path, whereas the merger merges these two flows. The choice of these two nodes, the establishment of the detour path, and the distribution of loads between the two paths are topology-aware, thereby challenging research topics.

#### 5.1.1 Congestion Detection and Hot Spot Node

Congestion detection in sensor networks has been studied in [11] and [23]. As pointed out in both papers, the actual congestion level around a node cannot be accurately measured by the buffer occupancy alone, which is a popular metric in traditional wired networks [1], [5], [6]. Therefore, TARA measures not only the buffer occupancy but also the channel loading,<sup>3</sup> as done in [11] and [23]. A node's congestion level is an aggregation of these two metrics. As soon as the congestion level hits the upper watermark, it declares congestion and becomes a *hot spot node*.

As shown in [21], [28], [32], and [33], it is a common practice in sensor networks to make extra nodes, referred to as *backup nodes*, into sleep mode (by turning off their radio)

3. The channel loading of a node indicates how busy the wireless channel is perceived by the node for a given time window. The channel loading of a node falls between 0 and 1.




Fig. 9. Illustration of TARA.

so that the overall resource (for example, energy) is saved during a dormant state. However, these nodes need to wake up periodically to check whether they are needed, such as when a node on the routing path runs out of battery. The portion of the time a node spends in an active state (that is, with radio on) is called *duty cycle*. The choice of the duty cycle is of great importance because it is undesirable to wake up too frequently or too infrequently. Numerous studies have focused on how to calculate an optimal duty cycle. In order to make the backup nodes congestion conscious, each node might adapt its duty cycle based on their congestion level. Every time a backup node wakes up, it measures its congestion level. If it observes that the current congestion level is greater than the previous measurement, it will shorten its sleep interval and wake up more often. As soon as its local congestion level is above a threshold, it will significantly reduce its sleep interval and becomes “alert.” During a dormant state, the congestion level around a backup node is extremely low because no traffic is routed through it. As soon as there is a hot spot, the backup nodes around the hot spot will first become active. Once they become active, they communicate with each other, and the generated traffic will increase the perceived congestion level of those backup nodes that are nearby. Similarly, a large number of backup nodes that are not too far away from the hot spot can thus be woken up and become ready to be picked to form detour paths.

#### 5.1.2 Traffic Distributor

In order to select the traffic distributor, every sensor node keeps track of the incoming traffic volume from each of its neighbors. This can be achieved by adding one more column in the neighbor table, a data structure that is already maintained by each node in most sensor networks. Using this information, the hot spot node can select the upstream neighbor that injects the most packets and send an *upstream control packet* to that neighbor. If that neighbor node is also congested, it repeats the process until it reaches a node with a low congestion level. That node will then become the traffic distributor.

#### 5.1.3 Traffic Merger

To locate a traffic merger, the distributor traces downstream by sending a *downstream control packet* to the sink that the most traffic is destined for. Therefore, the merger should be located on the routing path from the distributor toward the sink, with a low congestion level. We would like to point out that congestion level should not be the sole criteria for choosing a merger; a node's location is also important. For instance, in the example shown in Fig. 9, nodes C, D, and J may all have low congestion levels as many packets get dropped at node B. However, in this particular example,

TABLE 1  
Three Types of Intersection Zones

Intersection Zone	Congestion Topology	Traffic Distribution
Braided		
Crossing		
Merging		

only J is qualified to be a merger: building a detour path connecting G and C (or G and D) does not help distribute the traffic as C handles exactly the same packets as B does and will become the new hot spot node. Therefore, the choice of the merger is dependent on the topology of the *intersection zone*, which includes all the nodes the two intersecting flows have in common on their routing paths. Table 1 gives a detailed summary of possible topologies for interference zones for two flows (the second column) and illustrates the possible detour path topologies and traffic distribution (the third column). The traffic distributions of the braided, crossing, and merging intersection zones are based on lessons 2, 4, and 1 and 3, respectively.

A braided intersection zone is formed when two flows with different sinks interfere with or without sharing node(s) and do not cross each other. Selecting a merger for a braided intersection zone is similar to selecting a merger for a crossing intersection zone in that the merger has to be located after the two flows split. As a result, those nodes that forward packets to both sinks are not qualified to be a merger, and they should pass the control packet further downstream.

In the merging intersection zone, the abovementioned problem is no longer an issue since every node on the routing path only serves one sink, but the location of the merger is still of critical importance because it bears a great impact on the resulting capacity based on lesson 3. In order to apply this lesson, each node on the routing path keeps track of the hop count between itself and the sink, which is readily available from the routing table. Therefore, as the downstream control packet reaches a node, which contains the desired capacity level based on the traffic rate at the distributor, the node can decide whether or not to become a merger based on its distance to the sink. For instance, if the desired capacity exceeds the resulting capacity by this node becoming the

merger, then it should just forward the control packet to the next downstream node on the routing path that is closer to the sink. In parallel with the effort of tracing the merger, we also need to notify the distributor's ID to the merger so that the merger can establish a detour path.

#### 5.1.4 Detour Path

The merger tries to establish the detour path by locally flooding an REQ packet including the time-to-live (TTL) field toward the distributor. A node may receive multiple REQ packets from a merger due to the nature of flooding. In order to limit the flooding overhead, we have adopted the following optimizations:

- A node discards all the REQ messages if its local congestion level is above a threshold.
- A node discards all the REQ messages if it is already on the original routing path. This information is usually made available from the routing table.
- A node decrements the TTL value of the packet before forwarding and discards the packet if the TTL value falls below 0.
- A node keeps track of the largest TTL value it has seen. It drops the REQ messages whose TTL values are lower because these messages have traveled a longer path to reach this node.

When the REQ message reaches the distributor, a candidate detour path is established. Usually, the distributor will receive more than one candidate detour path, and it chooses the one whose REQ message has the largest TTL value, which corresponds to the smallest path length. To break the tie, the REQ message also records the path congestion level, which is the highest congestion level among all the nodes on the candidate detour path, and the distributor chooses a path with a lower congestion level by sending an ACK message toward the merger through the selected path.

#### 5.1.5 Traffic Distribution

To alleviate congestion, the distributor should split the outgoing traffic between the original path and the detour path. It should be careful in this process because some packets should not be routed via the detour path. In the example illustrated in Fig. 9, the distributor (node G) may need to forward packets to both sinks, and packets that are destined to sink F should not be assigned to the detour path because it is costly for the merger (node J) to connect to node F. To address this challenge, we introduce the concept of *streams*. As far as an intermediate forwarding node is concerned, all the flows that are destined to the same sink belong to the same stream. After introducing this concept, we propose a stream-based traffic distribution strategy. This strategy requires that the distributor checks each packet's destination sink before routing it. Each detour path has a corresponding sink, and if a packet's destination sink does not match the detour path sink, the distributor will only send that packet to the original path.

Among the packets that have the matching sink, we adopt the *weighted fair-share scheduling* to split traffic between two routes. Suppose the original path's congestion level is  $l_1$ , and the detour path's congestion level is  $l_2$ , with  $0 < l_1 < 1$  and  $0 < l_2 < 1$ . Then, the traffic rates for these two paths,  $t_1$  and  $t_2$ , should follow  $\frac{t_1}{t_2} = \frac{l_2}{l_1}$ . To further understand this algorithm, let us look at the detour



topologies shown in Table 1. Among the three types of intersection zones, the detour path in the crossing case has a higher congestion level than the detour path in the other two cases because the detour path intersects with another routing path. As a result, the detour path in the crossing case takes a smaller share of traffic compared to the detour paths in the other two types of zones.

## 5.2 Convergence

To converge into an equilibrium state, TARA needs to ensure that the duration of each resource adaptation is upper-bounded, and “thrashing” effect by frequent resource adaptations should be prevented. The information maintained by each node for resource adaptation is based on soft-state. The timer at each node silently expires, except for the distributor node. If TARA fails to create a detour path, the distributor’s timer expires and notifies the hot spot node of this failure. Upon receiving this message, the hot spot node resorts to the traditional traffic control approach by sending a backpressure message to its upstream neighbors. Due to TARA’s topology awareness, the frequency of resource adaptations is minimized. To further prevent TARA’s thrashing effect caused by the fluctuating traffic volume, TARA lowers the lower watermark. This will delay the resource shrinkage by allowing a temporary congestion level jitter.

## 6 PERFORMANCE EVALUATION

In this paper, we have conducted a detailed simulation-based study to compare TARA and several other strategies that cope with congestion in sensor networks. The results show that TARA clearly outperforms others in satisfying application fidelity requirements while spending less energy for each packet that is delivered from source to sink.

### 6.1 Performance Metrics

In this study, we propose to measure the effectiveness of TARA using the following metrics:

- *Fidelity index.* The primary goal of TARA is to satisfy application fidelity requirements. If the application intends to receive  $F^o$  packets per time unit, and a congestion control strategy actually delivers  $F$  packets per time unit to the applications (across all the sinks to which this application is connected), then we define the fidelity index of this strategy as the ratio of  $F/F^o$ .
- *Total energy consumption:* The total energy consumption,  $E^{total}$ , is represented as follows:

$$E^{total} = \sum_{i=1}^n (p^{xmit} \times t_i^{xmit} + p^{recv} \times t_i^{recv} + p^{idle} \times t_i^{idle}), \quad (3)$$

where  $p^{xmit}$ ,  $p^{recv}$ , and  $p^{idle}$  are the power consumption in transmission, reception, and idle modes, respectively, and  $n$  is the number of deployed nodes.

- *Bit energy consumption:* This metric is the ratio of the total energy consumption with respect to the total number of bits successfully delivered to the sink(s). It measures how effective is the energy usage from the end application’s point of view.

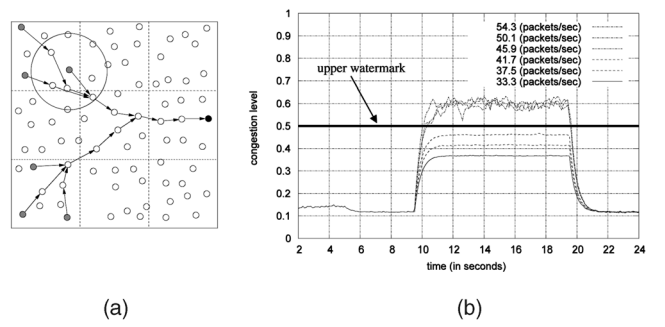


Fig. 10. Simulation setup. (a) Sensor field. (b) Traffic model.

### 6.2 Simulation Environment

First, we introduce the network model in our simulation studies. A sensor network has two main functionalities: sensing (data collection) and networking (data delivery), and our network model addresses both of these. We have 81 sensor nodes that are uniformly randomly distributed over a  $160 \times 160m^2$  field. The radio has a communication range of 30 m and an interference range of 50 m. On the average, each node has about nine neighbors, resulting in a rather dense network with a reasonable degree of redundancy. As far as sensing is concerned, we partition the entire field into nine grids, as shown in Fig. 10a, and all the sensors that belong to the same grid can detect any event within that grid.

We use the NS-2 simulator in our studies, and we have tuned many parameters according to the parameters of an actual Mica2 radio [28]. Each packet is 100 bytes long, and each node can hold at most 10 packets in its outgoing buffer. A node consumes 13.5 mW, 13.5 mW, 24.75 mW, and 15 uW in the idle listening, receiving, transmitting, and sleeping mode (that is, the radio is off), respectively.<sup>4</sup> The MAC protocol is a 2-megabit-per-second 802.11 distributed coordination function (DCF). Upon collision, a packet will have up to seven retransmissions. To exclude the impact of a routing protocol, the initial routing topology is hardwired.

In our simulations, we focus on two types of intersection hot spots: merging and crossing, as shown in Figs. 1c and 1d. In particular, we have considered the following traffic pattern to simulate the lifetime of a sensor network. At the beginning of the lifetime, the network is in its dormant period, and the offered load is 1 packet per second per source. After a random amount of time, the network enters a crisis state, and the crisis lasts for 10 seconds. In our studies, we vary the traffic rate during the crisis state from 33.3 to 66.9 packets/second to result in congestion with different severity levels. Fig. 10b plots the resulting congestion levels at the intersection node with increasing traffic rate. In our simulations, the upper watermark is set to 0.5, unless otherwise noted; therefore, Fig. 10b shows that TARA will be triggered when the traffic rate reaches around 45.9 packets/second in the merging topology.

In this study, we compare the following five strategies that can be adopted during a crisis period:

- *No congestion control.* This represents the baseline scenario, in which no congestion control is performed.

4. The parameters are based on those used in [29].

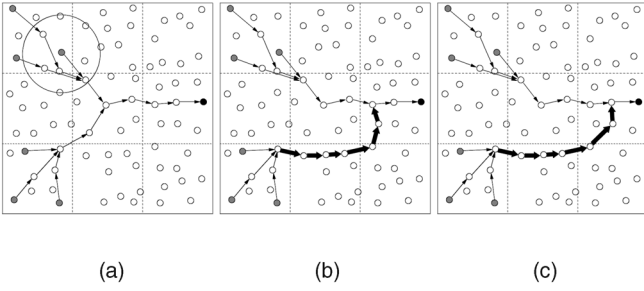


Fig. 11. Topologies for various congestion control strategies in the case of merging intersection. (a) No congestion control and traffic control. (b) Topology-unaware resource control. (c) Ideal resource control and TARA.

- *Traffic control.* In this case, when a node detects congestion, it sends a backpressure message to the upstream nodes on the routing path so that they can reduce the traffic load. In the simulations, we have carefully tuned the parameters such as the frequency of backpressure messages and the reduced traffic volume to ensure that it delivers the best possible fidelity index.
- *Ideal resource control.* This corresponds to an optimal offline resource control algorithm. Given a certain traffic load, this algorithm always finds the minimum topology for it. Though it cannot be implemented in a real system, we include this algorithm to investigate the gap between TARA and an optimal algorithm.
- *Topology-unaware resource control.* This corresponds to an ad hoc resource control algorithm. Specifically, it chooses the first downstream node with a low congestion level as the merger to form the detour path and blindly routes all the packets to the detour path.
- *TARA.* TARA is a topology-aware resource control algorithm.

### 6.3 Simulation Results for the Merging Intersection

In the merging scenario, six sources report to one sink: three of them in the upper left grid, and the other three in the lower left grid. The topology is illustrated in Fig. 11a. In the simulation, topology-unaware resource control selects the first downstream node whose congestion level is below the threshold as the merger (the resulting topology is shown in Fig. 11b) and routes all the packets to the detour path. On the other hand, TARA can construct the same detour path as an ideal resource control, which is shown in Fig. 11c. Based on lesson 3, a one-hop difference between the sink and the merging point in Figs. 11b and 11c can make a big difference on the channel capacities in the two figures.

Fig. 12 plots the fidelity index under different offered load for each strategy. When the load is lower than the network capacity, the fidelity index is closer to 1. After the load exceeds the capacity, the fidelity index for no congestion control degrades significantly. Traffic control shows a similar trend, but the fidelity index is even lower because it reduces the incoming traffic below the available capacity to prevent congestion. On the other hand, the three resource control schemes produce much better fidelity indexes. Ideal resource control has a fidelity index of 1 at all load ranges. Under the topology-unaware resource control

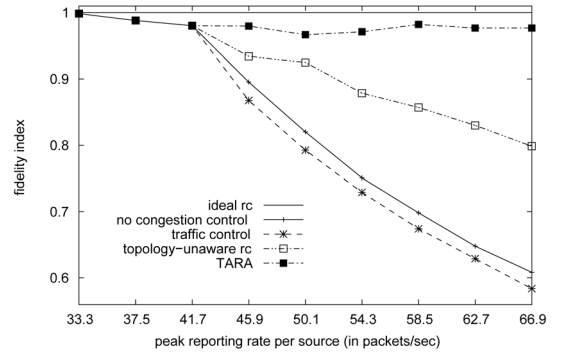


Fig. 12. Fidelity index during the crisis period in the merging congestion scenario.

scheme, since the construction of the detour path ignores the network topology, the resulting capacity is not enough to satisfy high traffic loads. As a result, though it performs much better than no congestion control and traffic control, the gap between this scheme and ideal resource control is still large. However, as shown in Fig. 12, TARA, due to its topology awareness, results in a fidelity index of 0.97 on the average regardless of the offered load level, close to ideal resource control. The gap between TARA and the ideal case is caused by the latency for TARA to establish a new topology, though both strategies end up having the same detour paths.

Fig. 13 shows the total energy consumption of the network. When the offered load is below the network capacity, the total energy consumption increases slowly with the load. To explain the sublinear increase, let us look at the total energy consumption  $E$ , which can be calculated as the sum of the energy consumption in transmission, reception, and idle states, that is,  $E = p^{xmit} \times t^{xmit} + p^{recv} \times t^{recv} + p^{idle} \times t^{idle}$ . When the load is low, a sensor node spends more time idling than transmitting, and idling energy thereby dominates the total energy consumption. As a result, as the traffic load increases, the increase in total energy consumption is rather modest. After the load is higher than the capacity, however, the total energy consumption increases dramatically. As expected, the traffic control scheme consumes the least amount of energy because fewer packets are transmitted. In the case of no congestion control, the energy consumption increases considerably as congestion

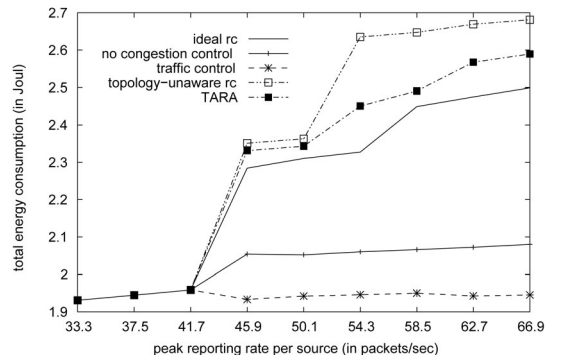


Fig. 13. Total energy consumption by network during the crisis period in the merging congestion scenario.

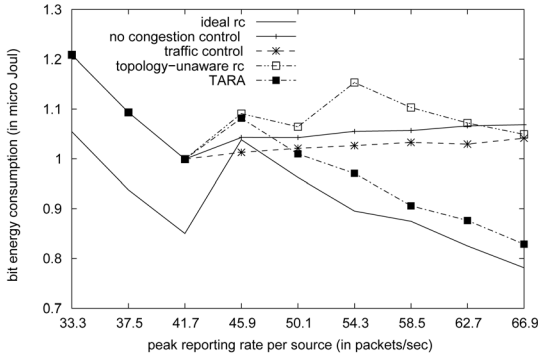


Fig. 14. Bit energy consumption during the crisis period in the merging congestion scenario.

first occurs because many packets now need multiple transmissions, but after some time, the total energy consumption stabilizes because the system has reached the maximum energy budget of the current topology. Resource control schemes consume more energy because more nodes that were previously turned off become actively involved in packet transmissions, and many more packets are delivered to the sink, as shown earlier. However, the amount of extra energy budget by resource control strategies is rather modest compared to the number of additional packets they can handle: at the traffic load of 66.9 packets/second, TARA can deliver  $(398 - 234)/234 = 70$  percent more packets than traffic control schemes by consuming  $(2.589 - 1.945)/1.945 = 33$  percent more energy. Among the three resource control schemes, ideal resource control consumes least energy. Topology-unaware resource control spends the most energy because congestion is not fully eliminated so that many packets still incur multiple transmissions (in Fig. 13, the total energy consumption significantly increases when a new detour path is created). However, we point out that TARA's energy consumption is only marginally higher than that of ideal resource control, that is, 4 percent most of the time.

Fig. 14 shows the bit energy consumption when varying the traffic load. Before congestion occurs, the bit energy consumption decreases with the traffic rate. This is because within this range, the total energy consumption only increases marginally with the traffic load. During light congestion (with a traffic load around 45.9 packets/second), the bit energy consumption of three resource control schemes are larger than the other schemes because the created detour path is not yet fully utilized. The reason why this happens even in ideal resource control is that the network capacity is not always a continuous function. As traffic rate further increases, the bit energy consumption in TARA and ideal resource control significantly decreases because the detour path gets more utilized. Again, the bit energy consumption of TARA is only slightly more than that of ideal resource control. Whether to enforce the resource controlling during light congestion is dependent on the application requirement. To balance the bit energy consumption and fidelity index, one can carefully decide when to enable resource controlling methods by configuring the watermark values appropriately.

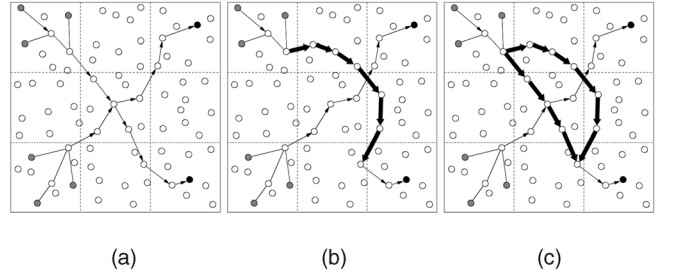


Fig. 15. Topologies for various congestion control strategies in the case of crossing intersection. (a) No congestion control and traffic control. (b) Topology-unaware resource control. (c) Ideal resource control and TARA.

#### 6.4 Simulation Results for Crossing Intersection

In this section, we study the performance of different congestion control strategies in a crossing scenario. In our experiments, we have six sources and two sinks, each sink connected to three sources, as shown in Fig. 15a. These two streams cross each other in the center of the sensor field. In this case, all three resource control strategies will find the same merger and establish the same detour path, but topology-unaware resource control blindly forward all the packets to the detour path, as shown in Fig. 15b, whereas the other two split them between the original path and the detour path, as shown in Fig. 15c.

Fig. 16 shows the fidelity indexes for all five strategies. The trend is similar to what we have observed in the merging case, with the exception that topology-unaware resource control now performs the worst. This is because topology-unaware resource control simply forward all the packets to the detour path, which cannot alleviate congestion, but only shift the hot spot. What makes matters worse is that it will result in the “ping-pong” effect in which a new detour path continues to be created. On the other hand, TARA effectively accommodates the incoming traffic by splitting them into the two paths at the same time.

Figs. 17 and 18 show the total energy consumption and bit energy consumption for the five strategies. The trend in the crossing case is similar to that in the merging case, except that the topology-unaware resource control strategy performs much worse here. The bit energy consumption for topology-unaware resource control is significantly higher

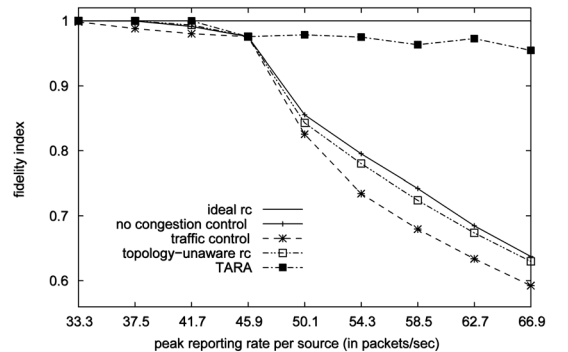


Fig. 16. Fidelity index during the crisis period in the crossing congestion scenario.




Fig. 17. Total energy consumption by network during the crisis period in the crossing congestion scenario.

because it does not alleviate congestion at all, though a large amount of resources are provided.

### 6.5 Discussion: Resilient to Transient Congestion

Setting the appropriate value for the upper watermark is a tricky issue. Ideally, we would like it to be small enough so that TARA can quickly respond to congestion. A small upper watermark, however, will make the network unstable: a short “glitch” may trigger the algorithm, and the glitch may disappear even before the detour path is established, causing much energy to be wasted. In order to allow quick recovery from real congestion without responding to glitches, we introduce an application-specific parameter, called the *transient threshold*, to filter out glitches. Using this threshold, resource-increasing methods will only take effect when the instantaneous congestion level has been above the upper watermark for a certain period of time. As a result, we can make the upper watermark small without undermining the stability of the network. We next conduct an experiment to illustrate our point. Fig. 19 shows the bit energy consumption with different upper watermark values. In this experiment, we considered a merging intersection case with six sources. The sources first generate a transient traffic surge, which lasts 0.9 second, and then 3 seconds later, a crisis period that lasts 10 seconds. The average congestion level during the transient traffic surge is 0.32. If the transient threshold is not used, then the upper watermark of 0.35 results in the least bit energy consumption since it can ignore the transient congestion and react to

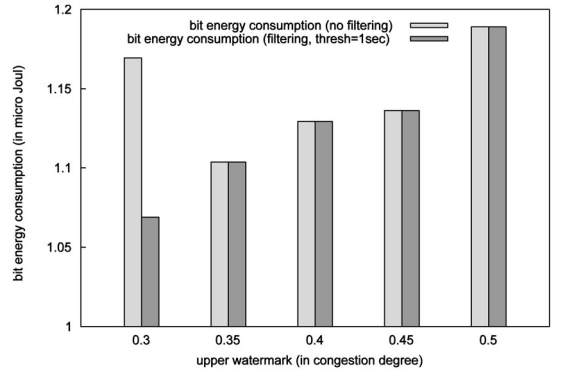


Fig. 19. Bit energy consumption during the crisis period.

the persistent congestion early. However, when the transient threshold of 1 second is used, the transient surge is filtered out, and the bit energy consumption is lowest when the upper watermark is 0.3. Clearly, the introduction of the transient threshold can lead to a smaller upper watermark, and more importantly, the bit energy consumption is lower than that of the no transient threshold case.

## 7 CONCLUDING REMARKS

The special nature of sensor networks calls for a new approach to alleviating congestion that can satisfy the application fidelity requirements. Therefore, TARA, a topology-aware resource adaptation strategy, was designed, and its performance was experimentally evaluated. TARA offers several advantages: 1) it is topology aware, 2) it is energy efficient, and 3) it is distributed. TARA uses a capacity analysis model to determine the needed topology. This model is formulated using a graph-coloring problem. The model results are compared against simulation (NS-2) and experimental results (using Mica2 motes). Detailed simulation results have shown that TARA can energy-efficiently absorb incoming traffic load. The results have also demonstrated that TARA performs very close to an ideal offline resource control algorithm in terms of both fidelity satisfaction and energy conservation.

## REFERENCES

- [1] S. Athuraliya, V.H. Li, S.H. Low, and Q. Yin, “REM: Active Queue Management,” *IEEE Network*, May 2001.
- [2] A. Cerpa and D. Estrin, “ASCE NT: Adaptive Self-Configuring Sensor Networks Topologies,” *Proc. IEEE INFOCOM '02*, June 2002.
- [3] D.-M. Chiu and R. Jain, “Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks,” *Computer Networks and Integrated Services Digital Network (ISDN) Systems*, vol. 17, no. 1, pp. 1-14, 1989.
- [4] C.T. Ee and R. Bajcsy, “Congestion Control and Fairness for Many-to-One Routing in Sensor Networks,” *Proc. Second Int'l Conf. Embedded Networked Sensor Systems (ACM SenSys '04)*, Nov. 2004.
- [5] S. Floyd and K. Fall, “Promoting the Use of End-to-End Congestion Control in the Internet,” *IEEE/ACM Trans. Networking*, Aug. 1999.
- [6] S. Floyd and V. Jacobson, “Random Early Detection Gateways for Congestion Avoidance,” *IEEE/ACM Trans. Networking*, Aug. 1993.
- [7] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, “Highly Resilient, Energy-Efficient Multipath Routing in Wireless Sensor Networks,” *ACM Mobile Computing and Comm. Rev.*, vol. 5, no. 4, pp. 11-25, 2001.

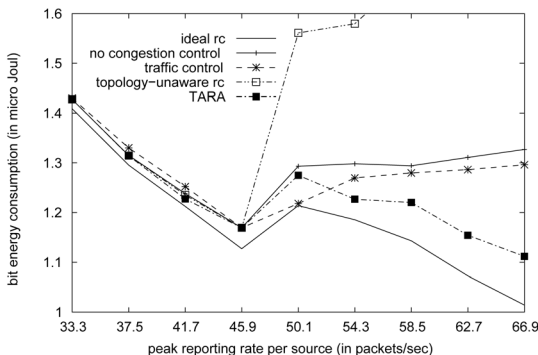


Fig. 18. Bit energy consumption during the crisis period in the crossing congestion scenario.

- [8] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Information Theory*, vol. 46, no. 2, pp. 388-404, Mar. 2000.
- [9] T. He, J. Stankovic, C. Lu, and T. Abdelzaher, "SPEED: A Stateless Protocol for Real-Time Communication in Sensor Networks," *Proc. 23rd Int'l Conf. Distributed Computing Systems (ICDCS '03)*, May 2003.
- [10] B. Hull, K. Jamieson, and H. Balakrishnan, "Mitigating Congestion in Wireless Sensor Networks," *Proc. Second ACM Conf. Embedded Networked Sensor Systems (SenSys '04)*, Nov. 2004.
- [11] J. Kang, Y. Zhang, and B. Nath, "Accurate and Energy-Efficient Congestion Level Measurement in Ad Hoc Networks," *Proc. IEEE Wireless Comm. and Networking Conf. (WCNC '05)*, Mar. 2005.
- [12] S.J. Lee and M. Gerla, "Dynamic Load-Aware Routing in Ad Hoc Networks," *Proc. IEEE Int'l Conf. Comm. (ICC '01)*, June 2001.
- [13] S.J. Lee and M. Gerla, "Split Multipath Routing with Maximally Disjoint Paths in Ad Hoc Networks," *Proc. IEEE Int'l Conf. Comm. (ICC '01)*, June 2001.
- [14] J. Li, C. Blake, D.S.J.D. Couto, H.I. Lee, and R. Morris, "Capacity of Ad Hoc Wireless Networks," *Proc. ACM/IEEE MobiCom 2001*, July 2001.
- [15] M.K. Marina and S.R. Das, "On-Demand Multipath Distance Vector Routing in Ad Hoc Networks," *Proc. Int'l Conf. Network Protocols (ICNP '01)*, Nov. 2001.
- [16] S. Ratnasamy, D. Estrin, R. Govindan, B. Karp, S. Shenker, L. Yin, and F. Yu, "Data-Centric Storage in Sensornets," *Proc. First ACM Int'l Workshop Wireless Sensor Networks and Applications (WSNA '02)*, 2002.
- [17] Y. Sankarasubramaniam, O. Akan, and I. Akyildiz, "ESRT: Event-to-Sink Reliable Transport in Wireless Sensor Networks," *Proc. ACM MobiHoc '03*, 2003.
- [18] R. Shah and J. Rabaey, "Energy Aware Routing for Low Energy Ad Hoc Sensor Networks," *Proc. IEEE Wireless Comm. and Networking Conf. (WCNC '05)*, Mar. 2002.
- [19] S. Tilak, N.B. Abu-Ghazaleh, and W. Heinzelman, "Infrastructure Tradeoffs for Sensor Networks," *Proc. ACM Workshop Wireless Sensor Networks and Applications (WSNA '02)*, Sept. 2002.
- [20] D.A. Tran and H. Raghavandra, "Routing with Congestion Awareness and Adaptivity in Mobile Ad Hoc Networks," *Proc. IEEE Wireless Comm. and Networking Conf. (WCNC '05)*, Mar. 2005.
- [21] Y. Tseng, C. Hsu, and T. Hsieh, "Power-Saving Protocols for IEEE 802.11-Based Multi-Hop Ad Hoc Networks," *Proc. IEEE INFOCOM '02*, June 2002.
- [22] A. Valera, W.K.G. Seah, and S.V. Rao, "Cooperative Packet Caching and Shortest Multipath Routing in Mobile Ad Hoc Networks," *Proc. IEEE INFOCOM '03*, Mar. 2003.
- [23] C. Wan, S. Eisenman, and A. Campbell, "CODA: Congestion Detection and Avoidance in Sensor Networks," *Proc. First Int'l Conf. Embedded Networked Sensor Systems (ACM SenSys '03)*, 2003.
- [24] R.J. Wilson, *Introduction to Graph Theory*, fourth ed. Longman Group, 1996.
- [25] A. Woo and D. Culler, "A Transmission Control Scheme for Media Access in Sensor Networks," *Proc. ACM/IEEE MobiCom 2001*, July 2001.
- [26] T. Yan, T. He, and J.A. Stankovic, "Differentiated Surveillance Service for Sensor Networks," *Proc. First Int'l Conf. Embedded Networked Sensor Systems (ACM SenSys '03)*, 2003.
- [27] F. Ye, G. Zhong, S. Lu, and L. Zhang, "PEAS: A Robust Energy Conserving Protocol for Long-Lived Sensor Networks," *Proc. 23rd Int'l Conf. Distributed Computing Systems (ICDCS '03)*, May 2003.
- [28] W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC Protocol for Wireless Sensor Networks," *Proc. IEEE INFOCOM '02*, June 2002.
- [29] W. Ye, J. Heidemann, and D. Estrin, "Medium Access Control with Coordinated Adaptive Sleeping for Wireless Sensor Networks," *IEEE/ACM Trans. Networking*, June 2004.
- [30] Z. Ye, S.V. Krishnamurthy, and S.K. Tripathi, "A Framework for Reliable Routing in Mobile Ad Hoc Networks," *Proc. IEEE INFOCOM '03*, Mar. 2003.
- [31] Y. Yi and S. Shakkottai, "A Hop-by-Hop Congestion Control over a Wireless Multi-Hop Network," *Proc. IEEE INFOCOM '04*, Mar. 2004.
- [32] R. Zheng, J.C. Hou, and L. Sha, "Asynchronous Wakeup for Ad Hoc Networks: Theory and Protocol Design," *Proc. ACM MobiHoc Conf. '03*, 2003.
- [33] R. Zheng and R. Kravets, "On-Demand Power Management for Ad Hoc Networks," *Proc. IEEE INFOCOM '03*, Mar. 2003.



**Jaewon Kang** received the BS and MS degrees in computer engineering from Seoul National University in 1992 and 1994, respectively. He is currently a PhD candidate in the Department of Computer Science, Rutgers University. Before entering the PhD program at Rutgers, he joined LG Electronics in 1994, working in the Code Division Multiple Access (CDMA) Smartphone Laboratory for the development of embedded operating system, wireless Transmission Control Protocol (TCP)/Internet Protocol (IP) suite, and CDMA data services for CDMA smart phones. His current research interest includes voice over packet wireless networks, fair packet scheduling in Internet backbone routers, and congestion control in wireless sensor networks. He received an award from the IBM Personal Systems Products division for contributions to the enhancement of the IBM Microkernel in 1995. He is a student member of the IEEE.



**Yanyong Zhang** received the PhD degree in computer science and engineering from Pennsylvania State University in 2002. She is an assistant professor in the Department of Electrical and Computer Engineering, Rutgers University and a faculty member at Wireless Laboratory (WINLAB). Her current research interests include sensor networks, sensor network security and privacy, and fault-tolerant sensor networks. She has received several US National Science Foundation (NSF) grants on these topics, including NSF's Faculty Early Career Development (CAREER) Award. She organized the Workshops on System Management Tools for Large-Scale Parallel Systems (2005 and 2006) and the Workshop on Self-Healing and Adaptive and Self-MANaged Systems (SHAMAN) in 2002. She is on the technical committee for several conferences and workshops and a member of the ACM and the IEEE.



**Badri Nath** received the PhD degree from the University of Massachusetts, Amherst. He is currently a professor in the Department of Computer Science, Rutgers University, and an associate director at Wireless Laboratory (WINLAB). He is a coprincipal investigator of the DataMan project at Rutgers University. His research interests include mobile computing and sensor networks. As part of his research, he is developing protocols and services suited for large-scale dense networks. His current interest is on developing an information architecture for sensor networks. He is also the recipient of the 10-year Best Paper Award at VLDB and the Best Paper Award at HPSR 2003. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).