Revising Quorum Systems for Energy Conservation in Sensor Networks

Daniela Tulone
EE Department, UCLA
tulone@ee.ucla.edu

Erik D. Demaine
CSAIL, MIT
edemaine@mit.edu

Abstract

Quorum systems are well-known techniques designed to enhance the performance of distributed systems, such as to reduce the access cost per operation, to balance the load, and to improve the system scalability. All of these properties make quorum systems particularly attractive for large-scale sensor applications involving coordinated tasks, such as rescue applications. In this paper we analyze quorum techniques in the specific context of sensor networks and energy conservation, and show why quorum systems designed for wired networks and their metrics fail to address the challenges introduced by sensor networks. We then redefine quorum metrics such as access cost, load balance, and capacity in a way that takes into account the limitations and the characteristics of sensor networks, and discuss some energy-efficient design strategies. In addition, we propose a family of energy-efficient quorum systems and a particular construction, called Regional Quorum system *(RQ), which reduces the quorum access cost. Finally, we propose a data diffusion protocol built on top of the RQ system, which improves energy consumption by reducing message transmissions and collisions, and increases the available bandwidth. We apply our diffusion protocol to analyze the RQ system using our novel metrics.*

1. Introduction

Quorum systems have been applied to a number of problems in distributed systems in order to enhance the performance and the scalability of the system. For instance, they have been used in data replication [20, 28, 26, 25, 31, 32, 33], dynamic routing [18, 17], overlay networks [14, 16, 15], distributed access control, and signatures [21]. A *quorum system* of a finite universe U is a set of subsets of U, called *quorums*, such that any two quorums intersect. Thus, each quorum has at least one node in common with every other quorum. Such a property is crucial to guarantee the consistency of replicated data when distributed operations (read/write operations) are performed on a quorum.

Quorum techniques have recently been applied to ad-hoc wireless networks, in particular to mobility management in cellular networks [11, 10, 6, 5], node tracking in mobile ad-hoc networks [2, 12], and data diffusion [13, 12, 5]. Some of this work, such as [11, 10, 13, 2], simply applies quorum techniques designed for wired networks to improve the system performance. However, as we show in this paper, those quorum systems and metrics do not address the challenges introduced by sensor networks such as the limited energy supply of sensors, their limited bandwidth, and their higher failure probability. For instance, the communication cost does not affect the efficiency of quorum systems in wired networks since the transmission of messages has no cost and servers have unlimited energy supply. There is some work on quorum systems in ad-hoc networks using topology and geographic information [4, 12, 3, 6]. For instance, Dolev et al. [3] propose an interesting use of quorum systems in mobile ad-hoc networks based on the *focal point* abstraction, which associates mobile nodes with a fixed geographical location, thus masking node mobility and failures. Carmi et al. [4] propose a geographic quorum system for hybrid networks (e.g., cellular networks) consisting of a static set of servers and mobile clients, which is designed to reduce the cost of accessing a quorum. However, they study the problem from a geometric perspective that does not take into account the limited energy supply and the characteristics of sensor networks (in fact, their cost of accessing a quorum Q from a location P is defined as the Euclidean distance between P and Q). For all these reasons their construction does not suit well multi-hop sensor networks. The recent work of Cheng et al. [6] proposes a location management protocol based on a novel quorum construction designed for multi-hop sensor networks. Their quorum system is built on top of a minimum dominating set, which is computed by sensor nodes in a distributed way, in order to reduce transmissions. Despite this work [3, 2, 12, 13, 11, 10] no previous work on quorum systems in ad-hoc networks provides an analysis of quorum techniques in the specific context of sensor networks and energy conservation. Our work diverges from previous proposals for proposing a systematic analysis of quorum techniques in sensor networks based on their properties and limitations, and for regarding them as a tool for energy conservation. The characteristics and limitations of sensor networks make quorum systems an appealing tool to conserve energy since operations are performed on subsets of sensor nodes. This can lead to a noticeable energy savings since the major energy consumption comes from radio communication (e.g., in [35] the cost for transmitting messages accounts for 90% or more of the total energy consumption). In addition, quorum systems can improve the *load balance* of sensor nodes, thus leading to a better energy management, and therefore to an extension of the system lifetime. In fact, a node that always relays messages is likely to fail sooner than other nodes, and this can cause network partitions or even service disruption. Moreover, reducing the number of message transmissions in the network increases the available bandwidth, which is a relevant feature since sensors have usually low bandwidth.

Although in this paper we regard quorum systems as a tool to conserve energy, it is worth to notice that they can be applied to enhance the *scalability* of the system since each operation is performed on a quorum set. This can be particularly advantageous in large-scale environmental monitoring applications that encompass several hundreds of nodes, and in high-density deployments. In these cases scalability becomes an issue. Moreover, quorum systems can be applied to improve the *fault-tolerance* and the *availability* of the service provided. For instance, they can be applied to data replication to guarantee the availability and the correctness of the data despite the high failure probability of sensor nodes. For all these reasons we think that quorum systems, if properly adapted to address sensors' constraints, can represent an efficient tool to reduce their energy consumption and improve scalability and fault-tolerance. They are particularly suitable to sensor applications involving coordination tasks, such as rescue applications in which operations performed by different teams of rescuers require coordination among them, or in medical applications with emergency care scenarios such as [23].

In this paper, we first discuss the reasons why quorum systems and quorum metrics proposed for wired networks do not suit well sensor networks, and then we redefine some quorum metrics by taking into consideration the limited energy supply of sensors, their high failure probability, and the network topology. We also discuss some design strategies, and propose a family of energy-efficient quorum systems and a particular construction, called *Regional Quorum system* (RQ), which reduces communication costs of accessing a quorum by using properties of the radio broadcast and topology information. Note that in contrast with [4, 6], the RQ system does not require coordination among nodes but exploits *locality*: each node computes a quorum based on local information and on system parameters related to the geographic system region and to the radio broadcast. We show the applicability of our RQ system by proposing an energy-efficient data diffusion protocol based on it, and analyze the RQ system using our metrics and our diffusion protocol. Our contributions can be summarized as follows:

- We analyze quorum techniques in the specific context of sensor networks and energy conservation, and show the unsuitability of quorum systems and metrics proposed for wired networks (see Section 4).
- We redefine quorum metrics such as access cost, load balancing, and quorum capacity in a way that takes into account the limitations of sensor networks (see Section 6).
- We discuss energy-efficient strategies useful when designing quorum systems in sensor networks, and apply these design strategies to propose a family of energy-efficient quorum systems with high resiliency. In particular, we show a construction, called the RQ system, that reduces the quorum access cost, and propose a data diffusion protocol built on top of it that saves energy by reducing transmissions and collisions. We analyze the RQ system using our novel metrics and our data diffusion protocol (see Section 7).

Structure of the paper. In Section 2 we compare our results with previous work relevant to our problem, and in Section 3 we briefly review quorum systems and their metrics. In Section 4 we analyze the main differences between quorum systems in wired and wireless sensor networks. We illustrate our system model in Section 5, and redefine quorum metrics in Section 6. In Section 7 we propose a family of energy-efficient quorum systems and the RQ system. Then, we analyze it using our metrics and our data diffusion protocol built on top of the RQ system.

2. Related Work

There is a large volume of work on quorum systems in wired networks [20, 28, 26, 25, 18, 17, 14, 16, 15, 21]. As we show in the paper, these solutions do not suit well sensor networks since they fail to address the challenges introduced by sensor networks, such as limited energy supply, high energy cost of radio operations, and high failure probability and communication failures. These constraints do not appear in wired networks. As a result, quorum systems proposed for wired networks, including those designed for dynamic networks [14, 16, 15, 18, 17] do not suit well sensor networks. We discuss the differences between wired and wireless sensor networks in Section 4. On the positive side, we show that sensors' properties such as the radio broadcast and the reliability of sensors, can be used to enhance efficiency in sensor networks.

Increasing attention has been given to quorum systems in wireless ad hoc network [1, 6, 3, 2, 12, 13, 11, 10]. However, some of this work [11, 10, 2, 12] simply applies previous quorum systems designed to wired networks to ad hoc networks, other work [6, 4, 3] proposes a quorum system that take into account topology information. For instance, Carmi et al. [4] study the problem of devising a quorum system that reduces the quorum access cost. However, their quorum system is designed for hybrid networks consisting of a static set of servers and mobile clients and it does not take into consideration the cost of multi-hop communication among mobile nodes (in fact, the communication cost is given by the Euclidean distance between a client and a quorum). Moreover, if high resiliency is required, the cost of computing the quorum system is $O(n^2 + n \log n)$ with n universe size, which is very expensive in case of large networks. The recent work of Cheng et al. [6] analyzes the location management problem and proposes a protocol based on a novel quorum construction designed for multi-hop sensor networks. Their quorum system is built on top a *minimum dominating set* computed by sensor nodes in a distributed way, in order to reduce transmissions. However, our RQ system built on top of clusters seems more efficient since the computation of the quorum construction is local. We attain that because of our assumption that each node is provided with an estimate of its absolute or relative position. Dolev et al. proposed in [3] an interesting use of quorum systems in mobile ad-hoc networks, which relies on the *focal point* abstraction that associates mobile nodes with fixed geographical locations. Our work uses this abstraction for its simplicity and because it masks variations in the set of sensors and a not-uniform distribution of sensors over the geographic system area. However, [3] applies previous quorums to a universe of focal points without taking into account the communication cost in multi-hop networks. Several solutions based on geographic information have been proposed to improve the efficiency of data diffusion [12, 7]. Recently, Tulone [9] has proposed a novel class of quorum systems suitable for highly mobile ad hoc networks. However, the focus of that work is on mobility and not on low-power sensor networks. Efficient quorum systems with probabilistic data guarantees have been proposed by Malkhi et al. [30], and applied to dynamic networks [18] and to mobile ad-hoc networks [13, 5].

Our work differs from all previous work for proposing a systematic analysis of quorum systems in sensor networks that takes into account both the limitations and the properties of sensor networks. We regard quorum techniques as a tool for energy conservation in sensor networks and redefine some previous quorum metrics [20] to provide a measure of the energy cost associated with a quorum system, such as access cost and load balance under the best and worst failure configuration, and quorum capacity, which is related to the system lifetime. Our family of quorum systems defined by energy functions, and our RQ construction, are novel.

3. Preliminaries

In this section we provide an overview of previous quorum metrics and show an example of quorum system construction. Quorum systems assume a static universe U of n nodes, over which quorums will be constructed. We assume that at most f nodes are faulty, where $0 \leq f \leq \frac{n}{2}$.

Definition 1 *A quorum system $\mathcal{Q} \subseteq 2^U$ is a collection of subsets of U called* quorums*, each pair of which intersects.*

Different quorum systems have been proposed in the literature based on different *failure models*. Below are the definitions of quorum metrics proposed in previous work, such as [20].

Access cost. The size of the smallest quorum of $\mathcal{Q}$ has been used to provide a measure of the access cost per operation.

Capacity. This metric indicates the number of quorum accesses that $\mathcal{Q}$ can handle during k time units normalized by k for $k \rightarrow \infty$.

Load balancing [20]. This is a measure of the inherent performance of a quorum system. A *strategy* w is a rule giving each quorum an access frequency, such that $\sum_{Q \in \mathcal{Q}} w(Q) = 1$. It induces a load on each node in U, which is the sum of the frequencies of all quorums it belongs to. This represents the fraction of the time a node is used. The load $\mathcal{L}(\mathcal{Q})$ of a quorum system $\mathcal{Q}$ is defined as the load of the busiest node minimized over all strategies. Intuitively, the load measures the quality of a quorum system since each element is accessed rarely if the load is low. The system load is a *best case* definition, which is achieved only if an optimal access strategy is used and there are *no failures*. It was defined in [20] as follows:

Definition 2 *Given a quorum system $\mathcal{Q}$, a* strategy *w is the probability distribution over all quorums Q in $\mathcal{Q}$, i.e., $\sum_{Q \in \mathcal{Q}} w(Q) = 1$, where $w(Q)$ is the access probability for quorum Q.*

Definition 3 *For an element $i \in U$, the load induced by a strategy w on i is defined as $l_w(i) = \sum_{Q \ni i} w(Q)$. The load induced by w on a quorum system $\mathcal{Q}$ is $\mathcal{L}_w(\mathcal{Q}) = \max_{i \in U} l_w(i)$. The system load of a quorum system $\mathcal{Q}$ is $\mathcal{L}(\mathcal{Q}) = \min_w \mathcal{L}_w(\mathcal{Q})$.*

Resiliency [20]. The resilience of a quorum system provides a measure of its availability, more precisely it measures how many crash failures a quorum system is guaranteed to survive. The resilience of a quorum system $\mathcal{Q}$ is defined as the largest k such that for every subset $K \subseteq U$ with $|K| = k$, there exists $Q \in \mathcal{Q}$ such that $K \cap Q = \emptyset$.

An example: a Grid quorum system [20]. Let us suppose that $|U| = k^2$ and that nodes are arranged into a $\sqrt{n} \times \sqrt{n}$ grid. A quorum of a Grid quorum system is defined as the

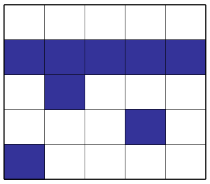

Figure 1. Grid quorum system for $|U| = 5^2$.

union of a full row and one element from each row below the full row [20, 28] (see Figure 1). Each quorum has size $O(\sqrt{n})$, the system load is $O(\sqrt{n})$, and its resilience is only $\sqrt{n} - 1$.

4. Quorum Systems in Sensor Networks

In this section we analyze quorum techniques in the specific context of sensor networks and discuss some strategies useful when designing quorum systems for sensor networks.

4.1. Wired Vs. Wireless Sensor Networks

We analyze here the main differences between quorum systems in wired and wireless sensor networks, and discuss the reasons why quorums designed for wired networks cannot address the challenges introduced by sensor networks.

Transmission costs. In wired networks messages are transmitted at *no cost* since nodes have unlimited energy source, while in sensor networks the transmission of messages is a highly energy-consuming operation. This is a relevant difference since as shown in [35], the communication cost amounts of 90% of the total energy. In addition, the transmission cost must take into account the number of intermediate nodes that relay a message since each broadcast has an energy cost.

Connectivity and network topology. Previous work on quorum systems assumes a fully connected network where each node is able to communicate to each node in the universe. However, this does not necessarily hold in sensor networks where some nodes might be isolated (e.g., because of failures, or obstacles, or intermittent connectivity). In addition, in a multi-hop sensor network most messages are forwarded by intermediate nodes whose failures (e.g., because of battery depletion) can cause network partitions and node isolation.

High failure probability. Sensor nodes have higher failure probability than servers because of their limited energy source and hardware, and often because of harsh environmental conditions (e.g., chemical sensors measuring soil moisture become unreliable after 4–5 days, as shown in [24]). Moreover, the sensor's failure probability is not constant but increases over time according to the battery usage.

Dynamic universe. Most quorum systems in wired networks assumes a static universe of servers, and rely on some *global knowledge* shared among all nodes regarding the quorum construction. This is often achieved via coordination among nodes during initialization and each time the system parameters vary. Note that in some sensor applications the universe changes over time (e.g., sensor nodes might be added, or in mobile applications nodes can move over time). It is inefficient to maintain a global knowledge in case of very large networks, such as environmental monitoring, or in case of mobile networks. As a result, quorum system should be computed at sensor nodes based on *local* information.

All of these elements show the unsuitability of quorums proposed for wired networks to address energy conservation since they do not take into consideration the actual cost associated with a quorum access. For instance, a quorum set chosen uniformly at random in the universe, or a grid quorum might not be the most convenient choice from a communication viewpoint. As discussed before, quorum systems proposed for ad-hoc networks such as [4, 6] do not address fully all these issues, although they take into account topology information.

4.2. Design Strategies

As discussed above, it is important to take into account the limited energy source of the sensors and the high communication cost when designing quorum systems for sensor networks. We discuss here some strategies that can lead to energy conservation when designing quorum systems for sensor networks. These strategies will drive our analysis in Sections 6–7.

Radio broadcast. Properties of the radio radio broadcast of sensor nodes can be used when designing quorum systems to reduce the amount of communications. Although symmetric radio links are not an entirely realistic assumption, recent publications suggest that careful neighborhood management and retransmissions can provide loss rates as low as 1–2 percent in static sensor networks [34], which should be sufficient for our purposes. Therefore, we can apply symmetric broadcast to avoid message transmissions and collisions (see Sections 5 and 7.2). Note that reducing the amount of message transmission among neighbors reduces also the probability of collisions, which represent a noticeable energy waste, and it increases the available bandwidth.

Clustering sensors. In most sensor applications nodes are densely deployed. A well-known technique to reduce the amount of transmissions exploits the high density of nodes and groups sensors together, thus requiring only the cluster head to transmit messages. In our case, we group sensors that are within their radio broadcast range to each other via focal points, as discussed in Section 5. As a result, our quorum systems are built on top of clusters.

Quorum intersection size. There is a clear trade-off between efficiency (quorum size) and fault-tolerance (minimum quorum intersection size), although small quorums do not necessary lead to energy saving in sensor networks. Most work on quorum systems, also in ad-hoc networks [6, 4, 3], simply requires that any two quorums must intersect. This implies that the minimum quorum intersection size is equal to 1. As a result, any client operation completes only if the quorum it accesses contains only correct nodes. In case of an unavailable quorum (containing some faulty node), the client has to access another quorum. A successful quorum access represents energy waste for the system, and this event is likely to occur in case of high failure probability. In order to avoid unsuccessful quorum accesses, we consider quorum systems with minimum intersection size equal to $f + 1$ where f is the maximum number of failures that is tolerated. This condition increases the resiliency of each quorum and makes it always available.

Intermediate nodes and locality. Nodes that forward messages from the sender to destination should be accounted as quorum members in order to save message transmissions. Moreover, the computation of a quorum system should be *local* to avoid additional message transmissions.

5. System Model

Our system model consists of an unknown and dynamic set of static sensor nodes displayed in a geographical region G of the plane. We denote by r the radio broadcast of each sensor node, and assume reliable broadcast and symmetric radio links. As discussed before, this assumption is reasonable although not entirely realistic (see [34]). In order to reduce the amount of message transmissions and collisions and improve the load balance, we cluster together nodes. For simplicity, we group sensors that are within their radio broadcast via *focal points* [3], however other cluster strategies can be employed. More specifically, we rely the implementation of focal points proposed in [9], which requires only *one sensor* node per focal point to transmit messages on behalf of it. Note that this approach not only conserves energy, but also improves the network bandwidth and reduces collisions. Moreover, the protocol for computing a cluster head has very low communication cost for being randomized and not requiring a leader election. Note that the focal point abstraction assumes that nodes are aware of their position. This can be attained in several ways, by considering some nodes equipped of GPS or by estimating the relative distance between nodes.

Focal points. We partition our geographic system area G into subregions $G_1, \ldots, G_n$ such that each subregion has a diameter smaller than the node's broadcast radius. Therefore, each sensor node in G_i can communicate with the other nodes in G_i. Figure 2 shows an example: a square grid of n focal points $G_1, \ldots, G_n$ each of length $\frac{r\sqrt{2}}{2}$. A focal point P_i is an abstraction that associates a subregion G_i with the sensor nodes that populate it, for $i = 1, \ldots, n$. The universe of our quorum systems is $U = \{P_1, \ldots, P_n\}$. Each focal point P_i has the same unique identification number of G_i that is represented by the coordinate of the upper leftmost corner of G_i. We refer to P_i as a static node of our universe U, and call it simply a node. Two focal points P_i and P_j are *adjacent* if their associated subregions G_i and G_j are adjacent.

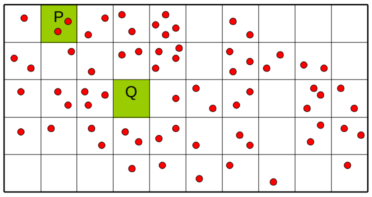

Figure 2. Partition of the geographical system region into focal point regions.

Similarly to [3], we do not limit the number of sensor failures, but rather the number of focal points that can become faulty, which cannot exceed f. A focal point P_i is *faulty* at time t if its subregion G_i does not contain *any* correct sensor node (it is either empty or contains faulty sensors). In Figure 2 focal point P is correct and Q is faulty. Note that this failure model is weaker than limiting the number of faulty sensors in most cases because of their high density nodes. To support fault tolerance, we focus on the *vertex-connectivity* of the network. In the worst case, a k-connected network requires k node failures to disconnect the network. Additionally, k-connectivity ensures that there exists at least k node-disjoint paths between every pair of nodes in the network, providing additional bandwidth. We assume that at most f focal points can become faulty, and that the network satisfies $(f + 1)$-connectivity. As shown in [8], this connectivity assumption can be attained by deploying few additional sensors. This assumption is also natural because it is difficult to define the reliability of a quorum system beyond the reliability of the network.

A failure configuration F is a subset of U containing at most f faulty focal points. We denote the set of faulty configurations containing at most f failures as $\mathcal{F}$, the *shortest communication path* from node i to node j with $i, j \in U$ under the failure configuration F as $(i, j)_F$, and the shortest path under no failures as (i, j). In case of multiple shortest paths, $(i, j)_F$ (or (i, j)) denotes a shortest path chosen uniformly at random.

6. Revised Metrics

In this section we show the unsuitability of quorum metrics defined for wired networks, and redefine metrics such as the load balance, access cost, and quorum capacity. Note that our quorum systems are built on top of clusters, more precisely focal points. As a result, from now on we will refer a focal point P_i as a node i.

6.1. Access Cost

The cost of accessing a quorum set Q, often called *probe complexity*, was defined in [29] as the size of the smallest quorum Q. This definition suits well wired networks where messages are transmitted at no cost and the communication path affects only the system latency. However, it does not capture the actual cost of accessing a quorum Q in sensor networks, which depends on the communication path used to transmit a message from the sender to each node in Q. Note that the location of the requester affects the cost of accessing a quorum since its communication cost changes according to the sender's location. We can summarize the factors that influence the cost of accessing a quorum Q as follows: (1) the current failure configuration, (2) the location of the sender, and (3) the routing protocol employed. Since the access cost is a property of the quorum system, it should be independent from the underlying routing protocol (the communication tree rooted at the sender). As a result, we define the access cost of a quorum system $\mathcal{Q}$ under the *best routing strategy* and under no failures. More precisely, we consider the shortest distance path between two nodes no failures.

We first define the cost $c(a)$ of accessing a quorum from node $a \in U$. It is the communication cost of accessing the *best quorum* (involving the minimum number of message transmissions) under the *best routing* strategy and under no failures:

$$c(a) = \min_{Q \in \mathcal{Q}} |\bigcup_{i \in Q} (a, i)|.$$

where (a, i) is the shortest path from node a to i under no failures, as defined in Section 5. Note that the routing strategy under consideration avoids retransmissions of the same message (i.e., an intermediate node forwards a message once). This is shown by the union of the shortest communication paths from the sender a to each node in Q. We will refer to it as the *minimum communication tree* of a. Therefore, $c(a)$ represents the cost of the minimum communication tree since only one sensor node per focal point transmits. The use of clusters (in our case of focal points) reduces the energy consumption associated with each quorum access. Note that there might be more than one quorum with minimum access cost and more than one minimum communication tree. The access cost of a quorum system $\mathcal{Q}$ is defined as the maximum cost over the nodes in U, that is, $c(\mathcal{Q}) = \max_{1 \leq a \leq n} c(a)$. According to this definition, the access cost is a property of the quorum system $\mathcal{Q}$ and it measures the inherent performance of $\mathcal{Q}$ for a given sensor network.

Since sensors have high failure probability, our previous definition of access cost of a quorum system $\mathcal{Q}$ is not indicative of the efficiency of $\mathcal{Q}$ in the presence of failures. We adapt the definition of $c(a)$ to capture the *worst failure configuration*. More specifically, we define the cost of quorum access from a under the worst failure configuration as $c_f(a) = \max_{F \in \mathcal{F}} \ \min_{Q \in \mathcal{Q}} | \cup_{i \in Q} (a, i)_F|$. Therefore, the access cost of $\mathcal{Q}$ under the worst failure configuration is $c_f(\mathcal{Q}) = \max_{1 \leq a \leq n} c_f(a)$.

6.2. Load Balance

The definition of load balance presented in Section 3 is not suitable for sensor networks because it does not take into account the load of the intermediate nodes serving as routers, which is relevant in multi-hop networks. We address that in our definition of load balance, which does not refer to sensor nodes but to focal points. This improves the load balance of the system since only one sensor per focal point is in charge of transmitting a message [9]. Therefore, the load of each sensor node in a focal point i never exceeds the load of a focal point i. To make this idea more concrete, if the sensor load is equally balanced within a focal point i, then the load of a sensor in i is equal to a fraction $\frac{1}{k}$ of the load of the focal point i, where k is the number of sensors in i. This approach improves resource utilization and energy consumption.

The load of a node $i \in U$ depends mainly on the failure configuration F, and on the probability of being chosen as a destination or intermediate node. We call *minimum-cost quorums* the set of quorums that have minimum access cost with respect to node a under failure configuration F, and denote this set as $MQ(a, F)$. Let us fix a sender node $a \in U$, and suppose that $M(a, F) = \{Q_1, \ldots, Q_r\}$ with $r \geq 1$ contains quorums with minimal access cost with respect to node a and under no failures. We denote the probability that node a accesses quorum Q_j under failure configuration F as $w_{a,F}(Q_j)$ for $j = 1, \ldots, r$, such that $\sum_{j=1}^{r} w_{a,F}(Q_j) = 1$. For each quorum $Q \in M(a, F)$ let us denote by $T_1, \ldots, T_{t_Q}$ with $t_Q \geq 1$ the set of minimum communication trees rooted at a to diffuse data to a quorum Q, and by $w'(i, T_k)$ the probability that node i belongs to tree T_k with $1 \leq k \leq t_Q$. We denote by $\bar{w}_{a,\emptyset}(i, Q)$ the probability that node i is chosen as a destination or intermediate node when accessing quorum Q under no failures, that is, the sum of $w'(i, T_k)$ over $k = 1, \ldots, t_Q$. Since we want the load balance to be a property of the quorum system, we define it under no failures. We define the load $l(i)$ of a node

$i \in U$ as follows:

$$l(i) = \max_{1 \leq a \leq n} \max_{Q \in M(a,\emptyset)} \bar{w}_{a,\emptyset}(i, Q).$$

Therefore, the load of node i is defined as the probability that node i transmits a message ($i = a$), or forwards a message to another node ($i \neq a$) under the best strategy and no failures. The load of a quorum system $\mathcal{Q}$ is defined as $l(\mathcal{Q}) = \max_{1 \leq i \leq n} l(i)$.

Similarly, we define the load of node i under the worst failure configuration. That is, the load $l_f(i)$ of a node $i \in U$ under the worst failure configuration is defined as follows:

$$l_f(i) = \max_{F \in \mathcal{F}} \max_{1 \leq a \leq n} \max_{Q \in M(a,F)} \bar{w}_{a,F}(i, Q).$$

The system load $l_f(\mathcal{Q})$ under the worst failure configuration is defined as $l_f(\mathcal{Q}) = \max_{1 \leq i \leq n} l_f(i)$.

Unbalanced quorum access. Note that similarly to [20], these definitions assume that the quorum accesses are uniformly distributed across the network. That is, the sensor nodes contained in each focal point issues requests at the same rate. However, this might not be the case in most sensor networks. For instance, in some applications most read/write requests are performed by some specialized nodes (e.g., the sinks). We generalize our previous definition by assigning different weights to different focal points. If c_i is the weight associated to focal point i, such that $\sum_{i=1}^{n} c_i = 1$, then the load of i is defined as follows:

$$l(i) = \max_{1 \leq a \leq n} c_a \max_{Q \in M(a,\emptyset)} \bar{w}_{a,\emptyset}(i, Q).$$

Note that the weights $c_1, \ldots, c_n$ might vary dynamically during the system lifetime. For instance, this occurs in case the system shows geographical areas of interest (e.g., active areas) that change over time.

6.3. Capacity

In [20] the quorum capacity indicates the number of quorum accesses that $\mathcal{Q}$ can handle during k time units normalized by k for $k \to \infty$. This definition is not meaningful in sensor networks because of the limited energy supply. We redefine capacity and make it strictly related to the energy consumption and to the system lifetime. More precisely, our goal is to provide a measure of the number of quorum accesses that a focal point can tolerate during a unit time interval such that the lifetime of a sensor node is at least T, and to relate this metric to the load balance defined previously. We achieve this by setting T to the mean time to failure (MTTF) of a sensor node, that is, an estimate of the average time until a sensor's first failure under some stress, and by defining the access capacity of a sensor node as follows:

Definition 4 *The access capacity C for a sensor node is defined as the average number of times it can be accessed (as a destination or intermediate node) during a unit time window Λ to meet its mean time to failure (MTTF).*

The access capacity of a sensor, along with the access cost and the load balance (defined in terms of focal points), are useful metrics to devise quorum systems suitable for a specific sensor application. In fact, the number of radio operations that a node can perform during its lifetime can be estimated as the energy budget that can be allotted to radio operations divided by the energy cost per transmission. Therefore, the access capacity C can be estimated as the average number of radio operations that a sensor node can perform multiplied by Λ/MTTF. As discussed in Section 6.2, the load $l_f(i)$ of a node i is the probability that node i is accessed as an intermediate or destination node under the worst failure configuration. Therefore, $\int_0^\Lambda l_f(i)\, dt$ represents the expected number of accesses of node i during Λ time units under the worst failure configuration. As a result, the quorum system should be designed such that $\int_0^\Lambda l_f(\mathcal{Q})\, dt \leq k \cdot C$, provided an active focal point region contains at least $k \geq 1$ sensor nodes and that the load is evenly distributed among its sensor nodes. If $\int_0^\Lambda l_f(\mathcal{Q})\, dt > k \cdot C$, other strategies should be applied to reduce the load (e.g., increase the size of the geographical cluster, or use a hierarchical approach).

7. Energy-Efficient Quorum Systems

In this section we apply our quorum design strategies and propose a family of energy-efficient quorum systems built on top of the focal points, and a specific construction, called *Regional Quorum system* (RQ). We also illustrate a data diffusion protocol that is built on top of focal points and the RQ system, and that uses information regarding the network topology and properties of the radio broadcast to reduce the number of transmissions and collisions. We apply this protocol to analyze the RQ quorum system and provide bounds on its access cost and load in the best and worst failure case, using the metrics proposed in Section 6.

7.1. A Family of Energy-Efficient Quorums

Dissemination quorum systems were introduced in [28] in the context of Byzantine failures to provide data consistency and data availability in case of self-verifying data. They are set of subsets of U such that any two quorums intersect in at least $f + 1$ nodes. An example of dissemination quorum system is the *threshold quorum system* $\mathcal{Q}_d$, which is a set of subsets of U of equal size $q = \lceil \frac{n+f+1}{2} \rceil$. We decide to adapt $\mathcal{Q}_d$ to sensor networks for its interesting features, which we summarize as follows:

1. Its minimum quorum intersection size that makes each quorum available (see Section 4.2);
2. Its flexibility since any subset of q nodes forms a quorum;
3. Its high resiliency since it tolerates up to $\lfloor \frac{n}{3} \rfloor$ failures.

Note that there is a trade-off between energy consumption and fault tolerance since the size q of the quorum increases with f. Our choice of tolerating up to one third of faulty focal points favors the robustness of the system versus its energy consumption (note that faulty nodes mask both sensor failures and the often not uniform distribution of sensors in a geographic area). However, we can reconcile fault-tolerance and efficiency by choosing a quorum system with high resiliency (e.g., $\mathcal{Q}_d$) and dynamically adapting the estimated number f of maximum failures as discussed in [9]. More precisely, nodes start with a small quorum size and increase it as the number of failures that have been detected or notified by other nodes increases.

In addition, $\mathcal{Q}_d$ allows the sender to compute a quorum based on different criteria (e.g., energy consumption, or areas of interest), and its construction does not require global information or coordination among nodes as in [4, 6], but it relies only on the number n of stationary focal points, which is an a-priori known system parameter related to the geographical system region, and on the maximum number of tolerated failures f, which can be dynamically adjusted during the system lifetime.

Since our focus is on energy conservation in sensor networks, we apply the $\mathcal{Q}_d$ system to design a family of energy-efficient quorum systems built on top of the focal point clusters. Let us consider a family $\mathcal{E}$ of *energy* functions such that each function $e : U \times U \times \mathcal{F} \rightarrow \mathbf{R}$ maps any two nodes i and j (focal points) under a failure configuration F into a real number representing the energy cost associated with i and j under the best routing strategy and failure configuration F. The $\mathcal{E}$ family induces a family of quorum systems $\{\mathcal{Q}_e\}_{e \in \mathcal{E}}$ such that for each $e \in \mathcal{E}$ quorum system $\mathcal{Q}_e$ is a refinement of $\mathcal{Q}_d$, and it is defined as follows. A quorum set $Q(i, F) \in \mathcal{Q}_e$ computed by node i under failure configuration F consists of q nodes $j_1, \ldots, j_q$ in U with qth smallest energy with respect to i such that $e(i, j_1) < \ldots < e(i, j_q) < \ldots < e(i, j_n)$.

Definition 5 *For each energy function $e \in \mathcal{E}$, we define a refinement $\mathcal{Q}_e$ of quorum system $\mathcal{Q}_d$ such that $\mathcal{Q}_e = \bigcup_{i \in U} \mathcal{Q}_i$ where $\mathcal{Q}_i = \bigcup_{F \in \mathcal{F}} Q(i, F)$ for all $i \in U$.*

We consider a specific energy function $h \in \mathcal{E}$ representing the number of hops in the communication path from a sender i to a receiver j under failure configuration F, such that $h(i, j, F) = |(i, j)_F|$. In case of communication paths of equal size we break tights by ordering nodes according to their identification numbers. We call this quorum system *Regional Quorum system* (RQ). In the following section we show the applicability of the RQ system by presenting an energy-efficient protocol for data diffusion built on top of it, and use this protocol to analyze RQ. For simplicity, we describe a basic (not optimized) version of the protocol.

7.2. An Energy-Efficient Data Diffusion Protocol

We describe here our data diffusion protocol, which is built on top of the focal points and of the RQ system. It is designed to reduce the amount of communication by exploiting properties of the radio broadcast of sensor nodes. We partition the geographic region G of the system into a square grid of n focal point regions $G_1, \ldots, G_n$ each of length $\frac{r\sqrt{2}}{4}$, where r is the radio broadcast, such that each focal point P_i communicates with its adjacent nodes. Each focal point is uniquely identified by the coordinate of the upper leftmost corner of its region. Note that smaller subregions G_i induce a larger universe size. However, this fact does not affect energy consumption of our protocol since a larger subregion G_i would imply a larger broadcast radius with higher energy consumption (energy consumption increases as r^2-r^4). Moreover, the performance gain of quorum techniques is more evident in case of large universe.

The idea underlying the protocol is to diffuse a message m across a quorum set containing the closest q focal points to the sender. Two approaches are possible: (1) the sender computes a geographical subregion R containing at least q focal points, or (2) the subregion R is dynamically adapted by intermediate nodes upon detecting faulty focal points (focal points that did not forward the request). For simplicity of presentation we present and analyze the first approach, which is not optimized, and sketch the second dynamic approach at the end of this section. Note that in order to guarantee that at least $q - f$ focal points received the request, our first version of the protocol must either assume that there is a communication path that connects the sender with at least $q-f$ correct focal points in R and that is contained in R, or it must request an acknowledgment from at least $q - f$ focal points in R.

Let us suppose that region R is computed by the sender, denoted as focal point S, before diffusing message m. For simplicity of presentation we suppose that R is a square geographical region of side $\lceil \sqrt{q} \rceil$ contained in G and such that any two correct nodes in R are connected and their shortest path is contained in R. A sender S broadcasts a request containing message m, its identification, and the coordinates of the upper rightmost point and the leftmost lower point of R that uniquely identify R. The message is recursively forwarded by intermediate nodes in R until the message is diffused across the entire region R. That is, each focal point in R receiving a request m broadcasts it to its adjacent focal points.

The key idea underlying the protocol is that not each of the intermediate nodes in R needs to forward message m. In fact, because of our communication model a message broadcast by node (i, j) reaches each of its adjacent nodes (k, l) where $i-1 \leq k \leq i+1$ and $j-1 \leq l \leq j+1$. We show in Lemma 1 that only half of them must forward the message in order to guarantee the correct diffusion of m over region R. For instance, a message broadcast by node (k, l) can be forwarded only by nodes $(k-1, l-1), (k-1, l+1), (k+1, l-1), (k+1, l+1)$. We call these nodes *landmark points* of (k, l). Note that this choice can be made dynamic to improve the load balance.

```
Sender $(i, j)$ :
1) bcast $\langle (i, j), (i, j), R, m \rangle$
```

Figure 3. Sender.

```
Intermediate node $(k, l)$:
Upon receiving request $\langle (i, j), (s, t), R, m \rangle$ from node $(s, t)$
1) if $(k, l) \in R$)
2)     if $(k, l) \in$ Landmarks$(s, t)$
3)         bcast $\langle (i, j), (k, l), R, m \rangle$
4)     else wait $\Gamma$ time
5)         if some of its adjacent Landmarks$(s, t)$ is faulty
6)             bcast $\langle (i, j), (k, l), R, m \rangle$
```

Figure 4. Intermediate node.

Figure 3 shows the sender's request $\langle (i, j), (i, j), R, m \rangle$ for diffusing a message m originated by node (i, j), across region R. Figure 4 shows the steps taken by each intermediate node (k, l) upon receiving a request $\langle (i, j), (s, t), R, m \rangle$ that has been forwarded by node (s, t). If node (k, l) is in the region R and it is one of the landmark points of (s, t), it broadcasts message $\langle (i, j), (s, t), R, m \rangle$, Figure 4:1–3. Since a landmark point can fail, each node has to monitor its adjacent nodes. An intermediate node N is in charge of transmitting a message broadcast by node A only if one of its adjacent landmarks of A has not replied (i.e., it is faulty), Figure 4:4–6.

7.3. Analysis

We evaluate our protocol with respect to its time complexity and communication cost, and analyze the RQ quorum system using our metrics defined in Section 6. More precisely, we bound its access cost and load under the best and worst failure configuration. Note that the performance of our protocol and consequently of the RQ system depends not only on the number of faulty nodes but also on their topology. As shown in Lemma 2 under no failures the message is diffused from node (i, j) across R after l steps, where l is the maximum number of nodes between the sender and the border of R. We call l diameter of R with respect to sender (i, j). Figures 5 and 6 show how the message is diffused under two different failure configurations in case R contains 11^2 nodes and the sender's location is at the center of R. Figure 5 shows the landmark points in case of faulty components of diameter at most 2. In that case the protocol terminates after 7 steps, in absence of failures it completes after 5 steps. Figure 6 shows a worst-case failure configuration in which the message has to be forwarded by each correct nodes in R because its landmark points and adjacent nodes are all faulty except one. In this case our data diffusion protocol does not show any improvement over the basic protocol of regional diffusion across R using clusters. The following Lemma 1 shows the correctness of our protocol under any failure configuration, and Lemma 2 computes its time complexity and its communication cost under the best and worst failure configuration.

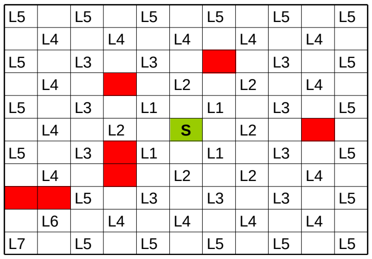

L5		L5		L5		L5		L5		L5
	L4		L4		L4		L4		L4	
L5		L3		L3		🟥		L3		L5
	L4		🟥		L2		L2		L4	
L5		L3		L1		L1		L3		L5
	L4		L2		S		L2		🟥	
L5		L3	🟥	L1		L1		L3		L5
	L4		🟥		L2		L2		L4	
🟥	🟥	L5		L3		L3		L3		L5
	L6		L4		L4		L4		L4	
L7		L5		L5		L5		L5		L5

Figure 5. Diffusion protocol with few failures.

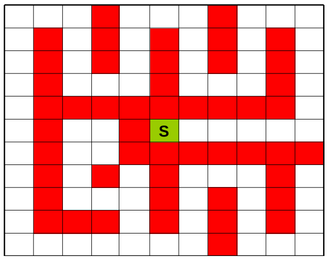

			🟥				🟥			
	🟥		🟥		🟥		🟥		🟥	
	🟥		🟥		🟥		🟥		🟥	
	🟥				🟥				🟥	
	🟥	🟥	🟥	🟥	🟥	🟥	🟥	🟥	🟥	
	🟥			🟥	S					
	🟥			🟥	🟥	🟥	🟥	🟥	🟥	🟥
	🟥		🟥		🟥				🟥	
	🟥				🟥		🟥		🟥	
	🟥	🟥	🟥		🟥		🟥		🟥	
							🟥			

Figure 6. Diffusion protocol under a worst failure configuration.

Lemma 1 *Our data diffusion protocol succeeds in diffusing a message to at least $q - f$ correct nodes in the geographic area R computed by the sender provided the shortest path of any correct node in R is contained in R.*

Proof: The protocol works in steps. For each step i we denote by R_i the area containing the correct nodes in R that have received the message, and by $\mathcal{L}_i$ the set of landmark points that broadcast a request at step i. Note that the nodes in $R_i \setminus R_{i-1}$ receives a message broadcast by nodes in $\mathcal{L}_{i-1}$. We show here that at any step i, if $|R_i| < q$, there exists a correct node in R_{i+1} that receives the request. A node $N \in R_i \setminus R_{i-1}$ that receives a request from $A \in \mathcal{L}_{i-1}$ does not forward the request under one of these conditions:

- if $N \notin R$ there is no need for N to broadcast the message;
- if both of the landmark points L_1 and L_2 of A that are adjacent to N are correct. Let us denote by S_x the subregion covered by the radio broadcast of node x. Then, $S_N \subseteq S_{L_1} \cup S_{L_2}$, and there is no need for N to broadcast since its broadcast area is covered by the broadcast area of L_1 and L_2.
- one landmark of L adjacent to N is alive (say L_1) and the other is faulty. In this case if the node C contained in $(S_{L_1} \cup S_{L_2}) \setminus R_i$ sends a broadcast, N does not need to transmit, since $S_N \setminus R_{i+1} \subset S_C$.

The correctness of our protocol follows from our assumption that any two correct nodes in R are connected and their shortest path is contained in R. □

Lemma 2 *In absence of failures at least q nodes in the regional area R computed by the sender will have received the message after l steps where* $\lceil \frac{\sqrt{q}-1}{2} \rceil \leq l \leq \lceil \sqrt{q} \rceil$. *The communication cost of diffusing data across R is at most* $\lceil \frac{q}{2} \rceil$ *in absence of failures and at most* $r^2 - f$ *under the worst failure configuration, where* $|R| = r^2$ *and* $r^2 \geq q$.

Proof: The diameter l of R with respect to the sender varies according to the location of the sender since R is contained in the geographic system region. Since R is the smallest square region containing q, then $\lceil \frac{\sqrt{q}-1}{2} \rceil \leq l \leq \lceil \sqrt{q} \rceil$. Note that $l = \frac{\sqrt{q}-1}{2}$ if the sender is at the center of region R and $q = (2l+1)^2$. Therefore, the protocol terminates after l steps since l is the farthest distance from the sender to the border of R. Under no failures the communication cost is equal to the number of landmark points in R, which is at most $\lceil \frac{q}{2} \rceil$. Under the worst failure configuration each node acts as a landmark node and the communication cost is equal to the number of nodes in R minus f. □

The following lemma provides an upper bound for the access cost of the RQ system under the best and worst failure configuration, and the load balance. It follows from the previous lemmas.

Lemma 3 $c(\mathcal{Q}_r) \leq \lceil \frac{q}{2} \rceil$ *and* $c_f(\mathcal{Q}_r) \leq r^2 - f$. *In addition,* $l(\mathcal{Q}_r) \approx \frac{1}{3}$.

Note that the RQ system represents an improvement over high-resilient quorums for sensor networks in case of no failures or sparse failures, because it accounts intermediate nodes as part of the quorum set, its construction does not require coordination as [4, 6], and it uses topology information. Moreover, it cuts the communication cost of accessing a quorum in half by using properties of the radio broadcast.

7.4. Enhancing Efficiency

Note that the data diffusion protocol we described in the previous section is not optimized. In fact, it requires acknowledgement from $q - f$ focal points in R in order to guarantee that the message has been received by a sufficient large number of focal points, if the communication path between the sender and $q - f$ active focal points in R is not necessary contained in R. Moreover, as shown in Lemma 2, in the worst failure scenario each correct focal point in R has to forward the message, which is inefficient (see Figure 6). We can enhance the efficiency of our data diffusion protocol by allowing the intermediate nodes to dynamically adapt region R based on their knowledge of faulty focal points.

Similarly to the previous protocol, the sender computes the region R and broadcasts its request along with an integer h that is equal to the diameter of R and represents the number of times the request has to be forwarded in absence of failures. A landmark node decreases h by one before forwarding the request, a node forwards the request until h becomes equal to zero. If the intermediate node detects that some of its adjacent landmark points are faulty it can increase the diameter h to overcome the missed transmissions. The diameter h is increased in order to guarantee that the message will be forwarded to a sufficient large number of focal points (i.e., it can be increased by the maximum distance between an active point and its closest faulty landmark points). This variation of the data diffusion protocol reduces the amount of transmissions in the worst failure scenario sketched in Figure 6.

8. Conclusions

In this paper we have shown how quorum system techniques can be applied to sensor networks to reduce their energy consumption, if properly designed. We have shown the unsuitability of previous quorums and quorum metrics proposed for wired networks, and redefined quorum metrics such as access cost, load balance, and capacity, taking into account the limitations and characteristics of sensor networks. In addition, we have proposed a family of energy-based quorum systems, and a specific quorum construction that reduces the quorum access cost. Moreover, we have illustrated an efficient data diffusion protocol that reduces transmissions by a factor of 2 with respect to the quorum size, and reduces collisions. Our work suggests an unexplored direction for reducing energy in sensor networks, and as a result, it leaves many open questions that we think are worthy of investigation, such as the study of other useful sensor properties and relations among metrics that can easily be used by applications to save energy.

References

[1] J. Jiang, Y. Tseng, T. Lai *Quorum-based asynchronous power-saving protocols for IEEE 802.11 ad-hoc networks.* Mobile Networks and Applications, 10, pp. 169–181, 2005.

[2] H. Lee, J. Welch, N. Vaidya *Location tracking using quorums in mobile ad-hoc networks.* Ad Hoc Networks 1(4): 371–381 (2003)

[3] S. Dolev, S. Gilbert, N. Lynch, A. Shvartsman, J. Welch. *GeoQuorums: Implementing Atomic Memory in Mobile ad-hoc Networks.* In Proc. 17th Intl. Conf. of Distributed Computing, pp. 306–320.

[4] P. Carmi, S. Dolev, S. Har-Peled, M. Katz, M. Segal. *Geographic Quorum System Approximations.* Algorithmica 41(4): 233–244.

[5] S. Bhattacharya. *Mobile wireless networks: Randomized location service in mobile ad hoc networks.* In Proc. 6th Intl. Workshop on Modeling analysis and simulation of wireless and mobile systems, pp. 66–73.

[6] M. Cheng, D. Du, D. Du. *Location management in mobile ad hoc wireless networks using quorums and clusters.* Wireless Communications and Mobile Computing, 5(7), pp. 793–803.

[7] X. Li, K. Moaveninejad, O. Frieder. *Regional gossip routing for wireless ad hoc networks.* Mobile Networks and Applications, (10), pp. 61–77.

[8] J. Bredin, E. Demaine, M. Hajiaghayi, D. Rus. *Deploying Sensor Networks with Guaranteed Capacity and Fault Tolerance.* In Proc. 6th Intl. Symp. on Mobile Ad Hoc Networking and Computing, pp. 309–319, May 2005.

[9] D. Tulone. *Ensuring data consistency in highly mobile networks via quorum systems.* To appear in Ad Hoc Networks.

[10] Z. Haas, B. Liang *Ad-hoc mobility management with uniform quorum systems.* Trans. on Networking, 7, pp. 228–240, 1999.

[11] R. Prakash, Z. Haas, M. Singhal *Load-balanced location management for cellular mobile systems using quorums and dynamic hashing.* Wireless Networks, 7(5), pp. 497–512.

[12] I. Stojmenovic *A scalable quorum based location update scheme for routing in ad-hoc wireless networks.* SITE, University of Ottawa, TR-99-09, Sept 1999.

[13] J. Luo, J-P. Hubaux, P. Eugster *PAN: providing reliable storage in mobile ad hoc networks with probabilistic quorum systems.* In Proc. 4th Intl. Symp. on Mobile ad-hoc networking and computing, pp. 1–12, 2003.

[14] M. Naor, U. Wieder. *Scalable and dynamic quorum systems.* Distributed Computing 17(4): 311–322 (2005).

[15] D. Malkhi, M. Naor, D. Ratajczak. *Viceroy: a scalable and dynamic emulation of the batterfly.* In Proc. Principle of Distributed Computing, 2002, pp. 183–192.

[16] U. Nadav, M. Naor. *The Dynamic And-Or Quorum System.* In Proc. Intl. Conf. of Distributed Computing, pp. 472–486.

[17] I. Abraham, D. Dolev, D. Malkhi *Ad-hoc networks: LLS: a locality aware location service for mobile ad-hoc networks.* In Proc. of Discrete Algorithms and Methods.

[18] I. Abraham, D. Malkhi *Probabilistic Quorums for Dynamic Systems.* In Proc. 17th Intl Conf. on Distributed Computing, pp. 60–74, 2003.

[19] D. Peleg, A. Wool *Crumbling walls: a class of practical and efficient quorum systems.* Distributed Computing, 10, pp. 69–83, 1997.

[20] M. Naor, A. Wool *The load, capacity and availability of quorum systems.* SIAM Journal of Computing, 27, pp. 423–447, 1998.

[21] M. Naor, A. Wool *Access control and signature via quorum secret sharing.* In Proc. 3rd Conf. Comm. Security, pp. 157–168, 1996.

[22] M. Maekawa A $\sqrt{n}$ algorithm for mutual exclusion in decentralized systems. ACM Transactions on Computer Systems, 3, pp. 145–159, 1985.

[23] K. Lorincz, D. Malan, T. Fulford-Jones, A. Nawoj, A. Clavel, V. Shnayder, G. Mainland, S. Moulton, M. Welsh. *Sensor Networks for Emergency Response: Challenges and Opportunities* In IEEE Pervasive Computing, Special Issue on Pervasive Computing for First Response, Oct-Dec 2004.

[24] N. Ramanathan and T. Schoellhammer and D. Estrin and M. Hansen and T. Harmon and E. Kohler and M. Srivastava. *The Final Frontier: Embedding Networked Sensors in the Soil.* CENS, UCLA, Tech. Report 68, Nov 2006.

[25] M. Herlihy *A quorum-consensus replication method for abstract types.* ACM Transactions on Computer Systems, 4, pp. 32–53, 1986.

[26] R. Bazzi *Planar quorums.* In Proc. 10th Intl. Workshop on Distributed Algorithms, pp. 582–592, 1996.

[27] D. Gifford *Weighted voting for replicated data.* In Proc. 7th Symp. of Operating Systems Principles, pp. 150–159, 1979.

[28] Dahlia Malkhi, Michael K. Reiter. *Byzantine Quorum Systems.* Distributed Computing, 11(4), pp. 203–213 (1998).

[29] Dahlia Malkhi, Michael K. Reiter, Avishai Wool. *The Load and Availability of Byzantine Quorum Systems.* SIAM J. Comput. 29(6), pp. 1889–1906 (2000).

[30] D. Malkhi, M. K. Reiter, A. Wool, R. Wright. *Probabilistic quorum systems.* Information and Computation 170(2): 184–206, 1 November 2001.

[31] M. Castro, B. Liskov. *Practical Byzantine Fault Tolerance and Proactive Recovery.* ACM Trans. on Computer Systems, 20(4), pp. 398–461, November 2002.

[32] D. Malkhi, M. K. Reiter, D. Tulone, E. Ziskind. *Persistent objects in the Fleet system.* In Proc. 2nd DARPA Information Survivability Conference and Exposition, June 2001.

[33] L. Zhou, F. Schneider, R. Van Renesse. *COCA: A secure distributed online certification authority.* ACM Trans. on Computer Systems, 20(4), pp. 329–368, November 2002.

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

[35] S. Madden. *The Design and Evaluation of a Query Processing Architecture for Sensor Networks.* Ph.D. Thesis. UC Berkeley. Fall, 2003.