

Integrated Coverage and Connectivity Configuration for Energy Conservation in Sensor Networks

GUOLIANG XING, XIAORUI WANG, YUANFANG ZHANG, CHENYANG LU,
ROBERT PLESS, and CHRISTOPHER GILL

Washington University in St. Louis

An effective approach for energy conservation in wireless sensor networks is scheduling sleep intervals for extraneous nodes while the remaining nodes stay active to provide continuous service. For the sensor network to operate successfully, the active nodes must maintain both sensing coverage and network connectivity. Furthermore, the network must be able to configure itself to any feasible degree of coverage and connectivity in order to support different applications and environments with diverse requirements. This article presents the design and analysis of novel protocols that can dynamically configure a network to achieve guaranteed degrees of coverage and connectivity. This work differs from existing connectivity or coverage maintenance protocols in several key ways. (1) We present a Coverage Configuration Protocol (CCP) that can provide different degrees of coverage requested by applications. This flexibility allows the network to self-configure for a wide range of applications and (possibly dynamic) environments. (2) We provide a geometric analysis of the relationship between coverage and connectivity. This analysis yields key insights for treating coverage and connectivity within a unified framework; in sharp contrast to several existing approaches that address the two problems in isolation. (3) We integrate CCP with SPAN to provide both coverage and connectivity guarantees. (4) We propose a probabilistic coverage model and extend CCP to provide probabilistic coverage guarantees. We demonstrate the capability of our protocols to provide guaranteed coverage and connectivity configurations through both geometric analysis and extensive simulations.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols—*Applications*; C.3 [Special-Purpose and Application-Based Systems]:—*Real-time and embedded systems*

General Terms: Algorithms, Design, Experimentation

Additional Key Words and Phrases: Sensor network, wireless ad hoc network, coverage, connectivity, energy conservation, topology maintenance, network geometry

This work is funded, in part, by the NSF under an ITR Grant CCR-0325529.

Authors' address: G. Xing, X. Wang, Y. Zhang, C. Lu, R. Pless, and C. Gill, Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130-4899; email: {xing,wang,yfzhang,lu,pless,cdgill}@cse.wustl.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 1550-4859/05/0800-0036 \$5.00

1. INTRODUCTION

Energy is a paramount concern in wireless sensor network applications that need to operate for a long time on battery power. For example, habitat monitoring may require continuous operation for months, and monitoring civil structures (e.g., bridges) requires an operational lifetime of several years. Recent research has found that significant energy savings can be achieved by dynamic management of node duty cycles in sensor networks with high node density. In this approach, some nodes are scheduled to sleep (or enter a power-saving mode) while the remaining active nodes provide continuous service. A fundamental problem is to minimize the number of nodes that remain active while still achieving acceptable quality of service for applications. In particular, maintaining sufficient sensing coverage and network connectivity with the active nodes is a critical requirement in sensor networks.

Sensing coverage characterizes the monitoring quality provided by a sensor network in a designated region. Different applications require different degrees of sensing coverage. While some applications may only require that every location in a region be monitored by one node, other applications require significantly higher degrees of coverage. For example, distributed detection based on data fusion [Varshney 1996] requires that every location be monitored by multiple nodes, and distributed tracking and classification [Li et al. 2002] requires even higher degrees of coverage. The coverage requirement for a sensor network also depends on the number of faults that must be tolerated. A network with a higher degree of coverage can maintain acceptable coverage in face of higher rates of node failures. The coverage requirement may also change after a network has been deployed, for instance, due to changes in application modes or environmental conditions. For example, a surveillance sensor network may initially maintain a low degree of coverage required for distributed detection. After an intruder is detected, however, the region in the vicinity of the intruder must reconfigure itself to achieve a higher degree of coverage required for distributed tracking.

Sensing is only one responsibility of a sensor network. To operate successfully, a sensor network must also provide satisfactory connectivity so that nodes can communicate for data fusion and reporting to base stations. The connectivity of a graph is the minimum number of nodes that must be removed in order to partition the graph into more than one connected component. The active nodes of a sensor network define a graph with links between nodes that can communicate. If this graph is K_c -connected, then for any possible K_c-1 active nodes which fail, the sensor network will remain connected. Connectivity affects the robustness and achievable throughput of communication in a sensor network.

Most sensor networks must remain connected, that is, the active nodes should not be partitioned in any configured schedule of node duty cycles. However, single connectivity is not sufficient for many sensor networks because a single failure could disconnect the network. At a minimum, redundant potential connectivity through inactive nodes can allow a sensor network to heal after a fault that reduces its connectivity, by activating more nodes. Alternatively, even transient communication disruption can be avoided by maintaining higher

connectivity among active nodes. Higher connectivity may also be necessary to maintain good throughput by avoiding communication bottlenecks.

Although achieving energy conservation by scheduling nodes to sleep is not a new approach, none of the existing protocols satisfy the complete set of requirements in sensor networks. First, most existing solutions have treated the problems of sensing coverage and network connectivity separately. The problem of sensing coverage has been investigated extensively. Several algorithms aim to find a close-to-optimal solution based on global information. Both Cerpa and Estrin [2002] and Meguerdichian and Potkonjak [2003] apply linear programming techniques to select the minimal set of active nodes for maintaining coverage. A more sophisticated coverage model is used to address exposure-based coverage problems in Meguerdichian et al. [2001a, 2001b]. The problem of finding the minimal exposure path is addressed in Meguerdichian et al. [2001a]. The maximal breach path and maximal support path in a sensor network are computed using Voronoi diagram and Delaunay Triangulation techniques in Meguerdichian et al. [2001b]. In Couqueur et al. [2002], node deployment strategies were investigated to provide sufficient coverage for distributed detection. Due to requirements for scalability and fault-tolerance, localized algorithms are more suitable and robust for large-scale wireless sensor networks that operate in dynamic environments. The protocol proposed in Tian and Georganas [2002] depends on local geometric calculation of sponsored sectors to preserve sensing coverage. The differentiated surveillance protocol proposed in Yan et al. [2003] was designed to achieve different degrees of coverage by dynamically scheduling nodes' duty cycles based on global clock synchronization. None of the above coverage maintenance protocols addresses the problem of maintaining network connectivity. On the other hand, several other protocols (e.g., ASCENT [Cerpa and Estrin 2002], SPAN [Chen et al. 2002], AFECA [Xu et al. 2002], and GAF [Xu et al. 2001]) aim to maintain network connectivity but do not guarantee sensing coverage. Unfortunately, satisfying coverage or connectivity alone is not sufficient for a sensor network to provide adequate service. Without enough sensing coverage, the network cannot monitor the environment with sufficient accuracy or may even suffer from "sensing voids"—locations where no sensing can occur. Without sufficient connectivity, nodes may not be able to coordinate effectively or transmit data back to base stations. The combination of coverage and connectivity is a special requirement introduced by sensor networks that integrate multihop wireless communication and sensing capabilities into a single platform. In contrast, traditional mobile ad hoc networks comprised of laptops only need to maintain network connectivity.

A second limitation of the aforementioned coverage protocols (except Chakrabarty et al. [2002] and Yan et al. [2003]) is that they can only provide a fixed degree of coverage. They cannot be dynamically reconfigured to meet different coverage requirements of applications. The algorithm proposed in Chakrabarty et al. [2002] requires global knowledge about the network and does not scale well in large-scale networks. Yan et al. [2003] can achieve differentiated degrees of coverage. However, the approach is not based on rigorous geometric analysis. In addition, as mentioned earlier, Yan et al. [2003] does

not address the problem of integrated coverage and connectivity configuration. Finally, while the PEAS [Ye et al. 2003] protocol was designed to address both coverage and connectivity in a configurable fashion it does not provide analytical guarantees on the degree of coverage and connectivity which are required by many critical sensor network applications (e.g., surveillance and structural monitoring).

The main contributions of this article are as follows. We first provide a geometric analysis of the fundamental relationship between coverage and connectivity based on a simple circular communication/sensing model. This analysis gives underlying insights for treating coverage and connectivity in a unified framework. This is in sharp contrast to several existing works that address the two problems separately. The problem of integrated coverage and connectivity configuration is formulated in Section 2 and the detailed analysis on the relationship between coverage and connectivity is presented in Section 2.2. Second, we present a Coverage Configuration Protocol (CCP) that can dynamically configure the network to provide different degrees of coverage as requested by applications. This flexibility allows the network to self-configure for a wide range of applications and environments with diverse or changing coverage requirements. CCP can provide both coverage and connectivity guarantees when the ratio of communication range and sensing range is no lower than 2, according to our analysis in Section 2. The design and analysis of CCP is presented in Section 3. Third, in Section 4, we integrate CCP with a representative connectivity maintenance protocol (SPAN [Chen et al. 2002]) to provide both coverage and connectivity guarantees when the ratio of communication range and sensing range is lower than 2. Fourth, we extend our theoretical analyses and CCP to more realistic communication/sensing models where the coverage can be probabilistic and the communication/sensing ranges irregular (see Section 5). We present simulation results in Section 6, and offer conclusions in Section 7.

2. PROBLEM FORMULATION

We define a convex region A as having a coverage degree of K_s (i.e., being K_s -covered) if every location inside A is covered by at least K_s nodes. Practically speaking, a network that provides a higher degree of coverage can achieve higher sensing accuracy and be more robust against sensing failures. Given a coverage region A and a node coverage degree K_s , the goal of an integrated coverage and connectivity configuration is to maximize the number of nodes that are scheduled to sleep under the constraints that the remaining nodes must guarantee: (1) A is at least K_s -covered, and (2) all active nodes are connected.

We now introduce the following simplifying assumptions that are useful for our initial analysis presented in Sections 2–4. Assumptions A(1)–A(4) describe a simple communication/sensing model where each node has uniform circular communication/sensing ranges. In Section 5, we will discuss how our results are extended when these assumptions are relaxed.

A(1) Every node v has a sensing region $S(v)$. Any point *inside* $S(v)$ is covered by v .

- A(2) The sensing region of every node is circular.
- A(3) The circular sensing region of every node has a same radius R_s . R_s is referred to as the *sensing range*. The circle $C(v, R_s)$ ¹ is called the sensing circle of node v .
- A(4) Any two nodes u and v can directly communicate with each other if their Euclidian distance is less than a communication range R_c , that is, $|uv| \leq R_c$.
- A(5) Every node knows its accurate location (e.g., through GPS or location service [Hightower and Borriello 2001]).

We assume any point on the boundary of a sensing region is *not* covered by the node (i.e., assumption A(1)). Although this assumption has insignificant practical impact, it simplifies our geometric analysis in following sections. In addition, we assume that region A contains at least one sensing circle.

In the rest of this section, we investigate the relationship between sensing coverage and network connectivity. We first show in Section 2.1 that sensing coverage always implies network connectivity if sensing range is at least twice that of communication range. In Section 2.2, we extend our results to a more general case where a network has a sensing coverage K_s ($K_s \geq 1$) and analyze the quantitative relationship between K_s -coverage and network connectivity. These results give insights into the design of our coverage protocol that will be presented in Section 3.

2.1 Sufficient Condition for 1-Coverage to Imply Connectivity

In this section, we analyze the relationship between 1-coverage and connectivity in a network. We note that connectivity only requires that the location of any active node be within the communication range of one or more active nodes such that *all* active nodes can form a connected communication backbone, while coverage requires all locations in the coverage region to be within the sensing range of at least one active node.

Intuitively, the relationship between connectivity and coverage depends on the ratio of the communication range to the sensing range. However, it is easy to see that a connected network may not guarantee its coverage regardless of the ranges. This is because coverage is concerned with whether *any* location is uncovered, while connectivity only requires that the locations of all active nodes are connected. Hence in the rest of this section, we focus on analyzing the sufficient condition for a *covered network* to guarantee connectivity. We have the following theorem:

THEOREM 1. *For a set of nodes that at least 1-cover a convex region A , the communication graph is connected if $R_c \geq 2R_s$.*

PROOF. We prove the statement using the Voronoi diagram of the nodes in the network, as shown in Figure 1. Let $\text{Vor}(u)$ represent the Voronoi cell of node u . We first prove that any two nodes whose Voronoi cells are adjacent

¹ $C(v, R)$ represents the circle that is centered at point v and has a radius R .




Fig. 1. The Voronoi diagram of the nodes that 1-cover a region. The cell of node v is composed of the points whose closest node is v .

can communicate with each other. As illustrated in Figure 1, p is the Voronoi vertex of three adjacent Voronoi cells $\text{Vor}(u)$, $\text{Vor}(v)$, and $\text{Vor}(w)$. According to the definition of a Voronoi diagram, u , v , and w are equally distant from p and are closest to p among all nodes. Hence, p must be covered by u , v , and w , otherwise it will not be covered by any nodes. According to the triangle inequality, we have:

$$|uv| \leq |pu| + |pv| < 2R_s \leq R_c.$$

We now prove that the network is connected by showing that there is a communication path between any two nodes s and t in the network. Suppose line segment st intersects consecutive Voronoi cells $\text{Vor}(s) = \text{Vor}(u_1)$, $\text{Vor}(u_2) \dots \text{Vor}(u_n) = \text{Vor}(t)$. Any two consecutive nodes in the series u_1 to u_n , since their Voronoi cells are adjacent, can communicate with each other according to our earlier discussion. Hence nodes u_1 to u_n constitute a communication path from s to t . The dotted path between s and t in Figure 1 illustrates such a path. \square

Theorem 1 establishes a sufficient condition for a 1-covered network to guarantee 1-connectivity. Under the condition that $R_c \geq 2R_s$, a sensor network only needs to be configured to guarantee coverage in order to satisfy both coverage and connectivity. In the next section, we extend our result to a more general case where a network can have a coverage degree of K_s ($K_s \geq 1$).

2.2 Relationship Between the Degree of Coverage and Connectivity

The previous section argues that if a region is covered, then the nodes covering that region are connected as long as $R_c \geq 2R_s$. If we maintain the condition of $R_c \geq 2R_s$, we can quantify the relationship between the degree of coverage and connectivity. This result is important for applications that require higher degrees of coverage or connectivity.

THEOREM 2. *A set of nodes that K_s -cover a convex region A forms a K_s -connected communication graph if $R_c \geq 2R_s$.*

PROOF. We first show that the lower bound on the connectivity of K_s -covered networks is K_s . We then show the tightness of this bound by a scenario where a node could be disconnected from other nodes by removing K_s nodes from a




Fig. 2. A partitioned network must have an uncovered region that separates two connected subnetworks.




Fig. 3. A scenario in which removing K_s nodes, located at p , disconnects node u from the rest of the network.

K_s -covered network. Disconnecting the communication graph of a set of nodes creates (at least) 3 disjoint sets of nodes, the set of nodes W that is removed, and two sets of nodes V_1 and V_2 , such that there are no edges from any node in V_1 to any node in V_2 . By Theorem 1, if it is possible to draw a continuous path between two nodes so that every point on the path is covered, then there exists a communication path between those two nodes. Therefore, to disconnect the graph, it is necessary to create a sensing void such that it is impossible to draw a continuous covered path connecting a node in V_1 to a node in V_2 . As illustrated in Figure 2, the nodes of V_1 may all lie in region S , the nodes in V_2 may all lie in region Q , and a set of nodes W must be removed to make a region T that is 0-covered. The nodes that are removed may actually lie in the region labeled S or Q , but their removal leaves the 0-covered region labeled as T .

To create a sensing void in an originally K_s -covered region A , it is clearly necessary to remove at least K_s nodes. Thus the network connectivity is at least K_s .

We now prove that K_s is the tight lower bound of the network connectivity by showing a scenario where a node can be disconnected from the rest of the network by removing K_s nodes if $R_c \geq 2R_s$. Consider the scenario illustrated by Figure 3. A node u is located at a corner (point q) of the rectangular node deployment region A that is K_s -covered. Suppose point p is on the sensing circle of node u such that pq has a 45° angle with the horizontal boundary of A . Suppose K_s coinciding nodes are located at point p . Clearly, these K_s nodes can K_s -cover the quarter circle of node u . We assume there are no other nodes whose sensing circles intersect the sensing circle of u . Removing these K_s coinciding nodes will create an uncovered region (i.e., a sensing void) surrounding node u . Furthermore, when $R_c = 2R_s$, there is no node within the communication range




Fig. 4. *Case 1.* The void (represented by the white region) does not intersect the region boundary.

of node u after the removal of these K_s nodes, that is, the communication graph is disconnected.

Hence the tight lower bound on the connectivity of K_s -covered networks is K_s . \square

We define *boundary* node as a node whose sensing circle intersects the boundary of the convex node deployment region A . Clearly all boundary nodes are located within R_s distance to the boundary of A . All the other nodes in region A are referred to as *interior* nodes. Intuitively, the connectivity of the boundary nodes dominates the overall connectivity of the communication graph. However, in a large-scale sensor network the interior nodes normally route more traffic, and higher connectivity is needed for interior nodes to maintain the required throughput. We define *interior connectivity* as the number of nodes (either interior or boundary) that must be removed to disconnect any two interior nodes in the communication graph of the nodes. We have the following theorem regarding the interior connectivity of K_s -covered networks.

THEOREM 3. *For a set of nodes that K_s -cover a convex region A , the interior connectivity is $2K_s$ if $R_c \geq 2R_s$.*

PROOF. Suppose u and v are two interior nodes, and the removal of a set of nodes W disconnects node u and node v . In order for nodes v and u to be disconnected, there must be a void region that separates node v from node u . There are two cases: either this void is completely contained within the node deployment region, or it merges with the boundary of the region.

Case 1. As illustrated in Figure 4, the void does not merge with the boundary. We will prove that one must remove at least $2K_s + 1$ nodes in this case to create such a void. We prove this by contradiction. Suppose $|W| < 2K_s + 1$. In this case, the void must completely surround a set of nodes, including node v . Since node v remains active, the sensing void must be at a distance at least R_s from v . Draw a line from v through a node j in W . Define line vj to be the direction we refer to as vertical. Now, there are at most $2K_s - 1$ remaining nodes (except node j) in W which are either on the line vj or to the left or the right of line vj . By the pigeonhole principle, there must be one side that has less than K_s nodes from the set W ; define that to be the left side. Draw the line




Fig. 5. Case 2. The void (represented by the white region) intersects the region boundary.

perpendicular to vj at v , to the left until it intersects the void region, and call this point p (note that p is covered by zero nodes.) Point p is at least R_s from node v and at least R_s from any point on or to the right of the vertical line. However, there are at most $K_s - 1$ nodes in the set W that are to the left of the line. This contradicts the assertion that p was originally K_s covered, and the removal of the nodes of W leaves it 0-covered. Thus $|W|$ is at least $2K_s + 1$.

Case 2. The void merges with the boundary of region A , as illustrated in Figure 5. In this case, the removal of a set of nodes W creates a void which separates the nodes v and u , and this void merges with the boundary of the region A that is being sensed. Since v is an interior node, all the points within a radius R_s from v are inside region A , and the same holds true for u . Furthermore, since the region A is convex, the line connecting any point v' within R_s from v and any point u' within R_s from u are inside the region A and must be intersected by the void, otherwise there will exist a continuous path ($vv'u'u$) from v to u which remains entirely within the node covered region and defines a network path in the communication graph (from Theorem 1). Thus the minimum width of the void that separates u from v is at least $2R_s$. Consider any two points in the void that are a distance of $2R_s$ apart. No node can simultaneously cover both points. This implies that at least $2K_s$ nodes were removed in the K_s -covered region A to create the void. We prove this bound is tight by the following case. Suppose the K_s -covered region A is a rectangle $A_1A_2A_3A_4$ with width $2R_s + r$ ($0 < r < R_s$). Two points x and y are located on the perpendicular bisector of A_1A_2 and A_3A_4 , respectively, as shown in Figure 5. Suppose there are K_s nodes (shown as dotted circles) located at point x and y , respectively. W is composed of these $2K_s$ nodes. We assume the nodes (not shown in the figure) whose sensing circles intersect the $2K_s$ nodes in W are far enough from point x and y such that the void created by the removal of W intersects both A_1A_2 and A_3A_4 . It is clear that the void disconnects the nodes on the left side from the nodes on the right side in communication graph. We have thus shown from the proof of Case 1 and Case 2, for a set of nodes that K_s -cover a convex region that the tight lower bound on the interior connectivity is $2K_s$. \square

We note that the interior connectivity defined in this section is different from the connectivity of the communication subgraph composed of solo interior nodes. This is because an interior node could connect to another interior node via boundary nodes, and the communication subgraph composed of solo interior nodes could be disconnected if all boundary nodes are removed as illustrated by Figure 5.




Fig. 6. An example of 1-coverage eligibility. The node with the bold sensing circle is ineligible since every point in its sensing range is covered by other nodes.

From Theorems 2 and 3, we can draw the conclusion that the boundary nodes that are located within R_s distance to the boundary of the coverage region are K_s connected. To the rest of the network, the interior connectivity is $2K_s$.

3. COVERAGE AND CONNECTIVITY CONFIGURATION WHEN $R_c \geq 2R_s$

Based on Theorems 1, 2, and 3, the integrated coverage and connectivity configuration problem can be handled by a coverage configuration protocol if $R_c \geq 2R_s$. In this section, we present a new coverage configuration protocol called CCP that uses this principle. CCP has several key benefits. (1) CCP can configure a network to the specific coverage degree requested by the application. (2) It is a decentralized protocol that only depends on local states of sensing neighbors. This allows CCP to scale effectively in large sensor networks in which nodes can fail at run time. It also allows applications to change its coverage degree at run time without incurring high communication overhead. (3) Our geometric analysis shows that CCP can provide guaranteed degrees of coverage.

3.1 K_s -Coverage Eligibility Algorithm

Each node executes an eligibility algorithm to determine whether it is necessary to become active. Given a requested coverage degree K_s , a node v is ineligible if every location within its coverage range is already K_s -covered by other active nodes in its neighborhood. For example, assume the nodes covering the shaded circles in Figure 6 are active, the node with the bold sensing circle is ineligible for $K_s = 1$, but eligible for $K_s > 1$. Before presenting the eligibility algorithm, we define the following notation.

- A point $p \in$ coverage region A is called an intersection point between nodes u and v , that is, $p \in u \cap v$, if p is an intersection point of the sensing circles of u and v .
- A point p on the boundary of the coverage region A is called an intersection point between node v and A , that is, $p \in v \cap A$ if $|pv| = R_s$.

THEOREM 4. *A convex region A is K_s -covered by a set of nodes if (1) there exist in region A intersection points between nodes or between nodes and A 's boundary; (2) all intersection points between any nodes are at least K_s -covered; and (3) all intersections points between any node and A 's boundary are at least K_s -covered.*




Fig. 7. A coverage patch is bounded by the arcs of five sensing circles. All points in the patch, including the boundary points, share the same coverage degree.

PROOF. We prove by contradiction. Let p be the point that has the lowest coverage degree k in region A , and $k < K_s$. Furthermore, suppose there is no intersection point in A which is covered to a degree less than K_s . The set of sensing circles partition A into a collection of coverage patches, each of them bounded by arcs of sensing circles and/or the boundary of A , and all points in each coverage patch have the same coverage degree. Suppose point p is located in coverage patch S . First we prove that the interior arc of any sensing circle cannot serve as the boundary of S . We prove by contradiction. Assume there exists an interior arc (of sensing circle $C(u, R_s)$) serving as the boundary of S , and crossing this arc (i.e., leaving the coverage region of node u) would reach an area that is lower covered than point p . This contradicts the assumption that point p has the lowest coverage degree in region A . Now we consider the following two cases.

Case 1. The point p lies in a region S whose boundary is only composed of exterior arcs of a collection of sensing circles (as Figure 7 illustrates). Since the sensing circles themselves are outside the sensing range of the nodes defining them, the entire boundary of this patch, including the intersection points of the sensing circles defining the boundary, has the same coverage degree as point p . This contradicts the assertion that p is covered to a degree less than K_s and all intersection points have coverage degree at least K_s .

Case 2. The point p lies in a coverage region S that is bounded by the exterior arcs of a collection of sensing circles and the boundary of A . As shown in Figure 8, point p is in a region bounded by the exterior arcs of node u , v , w , x and the boundary of region A . Similarly as Case 1, the entire boundary of this coverage patch, including the intersection points of nodes u , v , w , x and intersection points between nodes w , x and the boundary of A , has the same coverage degree as point p . This contradicts the assertion that p is covered to a degree less than K_s and all intersection points have coverage degree at least K_s .

Clearly the point p cannot lie in a coverage patch that is bounded solely by the boundary of region A . Otherwise A has the same coverage degree as point p . This contradicts the assumption that the region A is K_s covered. From the above discussion, the point p with lower coverage degree than K_s does not exist. Thus the region A is K_s covered. \square




Fig. 8. A coverage patch bounded by the arcs of four sensing circles and the region boundary. All points in the patch including those on the patch boundary share the same coverage degree.

```

int is_eligible (integer  $K_s$ )
begin
  find all intersection points inside  $C(v, R_s)$ :
  SI = {p | (p ∈ C(u, R_s) ∩ C(w, R_s) OR p ∈ C(u, R_s) ∩ A)
        AND u, w ∈ SN(v) AND |pv| < R_s};
  Find all coinciding nodes:
  SC = {u | |uv|=0};
  if (|SI|=0) {
    if (|SC| ≥  $K_s$ ) return INELIGIBLE;
    else return ELIGIBLE;
  }
  for (each point p ∈ SI)
  begin
    /*compute p's coverage degree*/
    sd(p) = |{u | u ∈ SN(v) AND |pu| < R_s}|;
    if (sd(p) <  $K_s$ ) return ELIGIBLE;
  end
  return INELIGIBLE;
end

```

Fig. 9. The K_s -coverage eligibility algorithm.

Theorem 4 allows us to transform the problem of determining the coverage degree of a region to the simpler problem of determining the coverage degrees of all the intersection points in the same region. A node is ineligible for turning active if all the intersection points inside its sensing circle are at least K_s -covered. To find all the intersection points inside its sensing circle, a node v needs to consider all the nodes in its sensing neighbor set, $SN(v)$. $SN(v)$ includes all the active nodes whose sensing circles intersect the sensing circle of v , that is, $SN(v) = \{\text{active node } u \mid |uv| < 2R_s \text{ and } u \neq v\}$. If there is no intersection point inside the sensing circle of node v , v is ineligible when there are K_s or more nodes that are located at node v 's position.

The resulting coverage eligibility algorithm is shown in Figure 9. The computational complexity for the eligibility algorithm is $O(N^3)$, where N is the number of nodes in the sensing neighbor set. The eligibility algorithm only requires information about the locations of all sensing neighbors. CCP maintains a table of known sensing neighbors based on the beacons (HELLO messages) that it receives from its communication neighbors. When $R_c \geq 2R_s$, the HELLO message from each node only needs to include its own location. When $R_c < 2R_s$, however, a node may not be aware of all sensing neighbors through

such HELLO messages. Since some sensing neighbors may be hidden from a node, it might activate itself to cover a perceived sensing void that is actually covered by its hidden sensing neighbors. Thus the number of active nodes would be higher than necessary in this case. To address this limitation, there must be some mechanism for a node to advertise its existence to the neighborhood of $2R_s$ range.

There are two approaches to make each node aware of its multihop neighbors. One is to broadcast HELLO messages in multiple hops by setting the TTL of each HELLO message. The other is to let each node include the locations of all known multihop neighbors in its HELLO messages. Specifically, each node may broadcast the locations and states of all active nodes within $\lceil 2R_s/R_c \rceil$ hops. The second approach reduces the number of broadcasts and is adopted by CCP (it is also used by SPAN [Chen et al. 2002] to maintain two-hop neighborhood tables). We should note that, in a network with random topology, such HELLO messages still cannot guarantee the discovery of all nodes within a distance of $2R_s$. Since including multihop neighbors in the HELLO messages introduces much higher communication overhead compared to a one-hop approach in a dense network, there is a trade-off between the beacon overhead and the number of active nodes maintained by CCP. We investigate this trade-off through experiments in Section 6.

We note that a special case (when coverage degree $K_s = 1$) of Theorem 4 was stated in Hall [1998] in which no proof is provided. Moreover, Theorem 4 presents a more general case that applies to any degree of coverage. This general case is important because flexible coverage configuration is a focus of this article.

3.2 The State Transition of CCP

In CCP, each node determines its eligibility using the K_s -coverage eligibility algorithm based on the information about its sensing neighbors and may switch state dynamically when its eligibility changes. A node can be in one of three states: SLEEP, ACTIVE, and LISTEN. In the SLEEP state, a node turns its radio off to conserve energy. Each sleeping node periodically turns its radio on and enters the LISTEN state to receive HELLO messages and reevaluate its eligibility. When a network is deployed, all nodes are initially in the ACTIVE state. In the ACTIVE state, a node actively senses the environment and communicates with other nodes. If an area exceeds the required degree of coverage due to high density, redundant nodes will find themselves ineligible and switch to the SLEEP state until no more nodes can be turned off without causing an insufficient degree of coverage. Over time, an active node may run out of energy which may cause the degree of coverage to decrease below the desired level. In this case, some nodes originally in the SLEEP state will find themselves becoming eligible and enter the ACTIVE state so that the network regains the desired degree of coverage. Since each node determines its state independently based on local information, there could be conflicting state transitions in the neighborhood. For example, when an active node dies and creates a void, several of its neighbors in LISTEN states may become active to cover the void

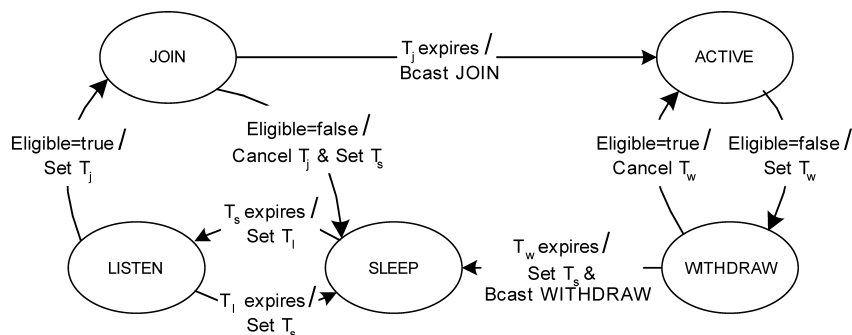


Fig. 10. State diagram of CCP.

simultaneously resulting in unnecessarily high coverage. We use two transient states, JOIN and WITHDRAW, to reduce the contention among neighbors in the transition from LISTEN to ACTIVE and the transition from ACTIVE to SLEEP, respectively. The state transition in CCP is similar to SPAN [Chen et al. 2002] and several other protocols [Tian and Georganas 2002; Xu et al. 2002]. We now describe the specific rules used in CCP (see Figure 10).

- In *SLEEP*. When the sleep timer T_s expires, a node turns on the radio, starts a listen timer T_l , and enters the LISTEN state.
- In *LISTEN*. When a beacon (HELLO, WITHDRAW, or JOIN message) is received, a node evaluates its eligibility. If it is eligible, it starts a join timer T_j and enters the JOIN state. Otherwise, it sets a sleep timer T_s and returns to the SLEEP state when T_l expires.
- In *JOIN*. If a node becomes ineligible before T_j expires (e.g., due to the reception of a JOIN message), it cancels T_j , starts a sleep timer T_s , and returns to the SLEEP state. If T_j expires, it broadcasts a JOIN message and enters the ACTIVE state.
- In *ACTIVE*. When a node receives a HELLO message, it executes the coverage eligibility algorithm to determine its eligibility to remain active. If it is ineligible, it starts a withdraw timer T_w and enters the WITHDRAW state.
- In *WITHDRAW*. If a node becomes eligible (due to the reception of a WITHDRAW or HELLO message from a neighbor) before the T_w expires, it cancels the T_w and returns to the ACTIVE state. If T_w expires, it broadcasts a WITHDRAW message, starts a sleep timer T_s , and enters the SLEEP mode.

Both the join and withdraw timers are randomized to avoid collisions among multiple nodes that decide to join or withdraw. The values of T_j and T_w affect the responsiveness of CCP. Shorter timers lead to quicker response to the variations in coverage. Both timers should be set appropriately according to the network density. For example, for a denser network where a node has more neighbors, both timers should be increased to give a node enough time to collect the JOIN or WITHDRAW messages from its neighbors. In addition, we note that ranking the expiration time of join or withdraw timers according to the utility of the

node may result in a better coverage topology and fewer active coverage nodes. Intuitively, a node that will cover more uncovered area should have a shorter join timer when competing against other nodes. The proper ranking heuristics are left to our future work. In this article, all nodes are deemed to share the same rank.

4. COVERAGE AND CONNECTIVITY CONFIGURATION WHEN $R_c < 2R_s$

As described in Section 2.1, CCP does not guarantee connectivity when the ratio of the communication range to the sensing range is less than 2. In this section, we present a simple approach for integrating CCP with an existing connectivity maintenance protocol, SPAN [Chen et al. 2002], to provide both sensing coverage and communication connectivity. SPAN is a decentralized coordination protocol that conserves energy by turning off unnecessary nodes while maintaining a communication backbone composed of active nodes. The communication backbone maintains the topology of the network such that all active nodes are connected through the backbone and all inactive nodes are directly connected to at least one active node. Although SPAN is not designed to configure the network into different connectivity, its eligibility algorithm results in a communication backbone that is capable of maintaining comparable network capacity and communication delay as the original network with all nodes active.

Integrating CCP with SPAN is simplified by the fact that they share a similar structure and similar states. Each node running SPAN maintains a neighborhood table that includes the locations of its one-hop neighbors as well as the IDs of their active neighbors and makes local decisions on whether to sleep or to stay awake as a coordinator and participate in the communication backbone (the details of SPAN are presented in Chen et al. [2002]).

The main difference between CCP and SPAN lies in their eligibility rules. In SPAN, a noncoordinator will become eligible to serve as a coordinator whenever it finds it satisfies the connectivity eligibility rule: at least one pair of its neighbors cannot reach each other either directly or via one or two active nodes. A coordinator will withdraw if it becomes ineligible. A straightforward way to provide both coverage and connectivity is to combine the eligibility according to both SPAN and CCP when a node makes a decision to join or withdraw. The resulting eligibility algorithm for providing both coverage and connectivity is as follows:

- *Eligibility rule for inactive nodes.* An inactive node will be eligible to become active if it is eligible according to the eligibility rule of SPAN *or* CCP.
- *Eligibility rule for active nodes.* An active node will withdraw if it satisfies the eligibility rule of *neither* SPAN *nor* CCP.

When $R_c/R_s < 2$, the active nodes picked by the CCP eligibility rule guarantee that the region is covered to the required degree. However, these active nodes might not communicate with each other. In this case, the eligibility rule SPAN will activate extra nodes so that every node can reach an active node within its communication range.

In SPAN, a HELLO message includes the node's location coordinates and the IDs of neighboring coordinators. Thus a node can know the existences of the coordinators in a two-hop neighborhood. We modified the structure of the SPAN HELLO message to include the coordinates of each neighboring coordinator. Thus, a node can maintain a neighborhood table that includes the locations of all two-hop neighboring coordinators from the HELLO messages. As discussed in Section 3.1, the information about the locations of two-hop active neighbors can reduce the number of active nodes under CCP when $R_c/R_s < 2$. We examine the effect of using 2-hop information in Section 6.

5. RELAXATION OF ASSUMPTIONS

The theoretical results and the CCP protocol presented so far are based on the assumptions made in Section 2. In this section, we extend our results to more realistic cases by relaxing some of those assumptions.

5.1 Relationship Between Coverage and Connectivity

In previous sections, we assumed that all nodes in a sensor network have uniform and circular communication/sensing regions (i.e., assumptions A(2)-(4) in Section 2). However, these assumptions may not be strictly accurate in real-world sensor network platforms. For example, empirical studies have found that the communication range of Mica Motes [Crossbow 2003] is highly dependent on the environments [Zhao and Govindan 2003]. The sensing range of a node depends on the node modality and is affected by the background noise in environments. In this section, we assume that nodes may have nonuniform and irregular (i.e., possibly noncircular) communication and sensing regions. The analysis of the relationship between coverage and connectivity presented in Section 2 therefore needs to be reexamined under these more realistic assumptions. We define the following concepts for the convenience of our discussion.

- The *minimum communication range (MCR)* of node v , $R_{cmin}(v)$ is the minimum distance between node v and the boundary of its communication region, that is, the region in which all the nodes can be reached by v .
- The *maximum sensing range (MSR)* of node v , $R_{smax}(v)$ is the maximal distance between node v and the boundary of its sensing region.
- The set of sensing neighbors of node v , $SN(v)$ includes all the active nodes whose sensing regions intersect v 's sensing region, that is, $SN(v) = \{\text{active node } u \mid S(u) \cap S(v) \neq \emptyset \text{ and } u \neq v\}$.
- The *minimum communication range of a sensor network*, R_{cmin} is defined as the minimum MCR of all nodes in the network.
- The *maximum sensing range of a sensor network*, R_{smax} is defined as the maximum MSR of all nodes in the network.

We then have the following theorem.

THEOREM 5. *For a set of nodes that K_s -cover ($K_s \geq 1$) a convex region A , Theorems 1, 2 and 3 still hold if R_{cmin} and R_{smax} are substituted for R_c and R_s , respectively.*

PROOF. Since region A is K_s -covered by the nodes and the actual sensing range of every node is upper-bounded by R_{smax} , A is K_s -covered by the circles that are centered at the nodes and have a radius R_{smax} . Hence Theorems 1–3 hold if the communication range of every node is R_{cmin} . From the definition of R_{cmin} , the actual communication range of every node is lower-bounded by R_{cmin} . Hence the results on the network connectivity proved in Theorems 1–3 still hold. \square

Theorem 5 depends on the knowledge of two global network properties, R_{smax} and R_{cmin} , which may not be easily available in a large-scale sensor network. Furthermore, from Theorem 5, the sufficient condition to guarantee the network connectivity becomes $R_{smax} \geq 2R_{cmin}$ which may be too conservative for heterogeneous sensor networks where nodes may have different types of network interfaces and/or node modalities. The proof of sufficient condition for network connectivity in Theorem 1 (see Section 2.1) depends on the fact that, when $R_c \geq 2R_s$, any two sensing neighbors can communicate directly. This observation allows us to extend Theorem 1 to the case where nodes have different communication and sensing ranges.

THEOREM 6. *For a set of nodes that K_s -cover ($K_s \geq 1$) a convex region A , Theorem 1 still holds if the following property holds for any node u in the network.*

$$\forall v \in SN(u), R_{cmin}(u) \geq R_{smax}(v) + R_{smax}(u). \quad (1)$$

PROOF. Let node v be a sensing neighbor of node u . Since the sensing regions of u and v are contained by the circles $C(u, R_{smax}(u))$ and $C(v, R_{smax}(v))$, respectively, $C(u, R_{smax}(u))$ and $C(v, R_{smax}(v))$ intersect. Hence $|uv| < R_{smax}(u) + R_{smax}(v)$. From (1), $R_{cmin}(u) > |uv|$, that is, node v is within the communication range of node u . Similarly, it can be shown that node u is within the communication range of node v . That is, any two sensing neighbors are connected in the communication graph. For any two nodes i and j , similar to the proof of Theorem 1, it can be shown that a communication path can be constructed along the line segment joining i and j , since any two sensing neighbors whose sensing regions are intersected by line ij can communicate with each other. \square

For a sensing-covered network, Theorem 6 gives a sufficient condition for connectivity based on the communication and sensing ranges of sensing neighbors. This condition is less pessimistic than Theorem 5 in heterogeneous network platforms. It also allows a sensing-covered network to determine whether it needs an explicit connectivity configuration based on local states.

5.2 Eligibility Algorithm

In this section, we extend CCP to more realistic cases. Similar to Section 5.1, we also assume that nodes may have nonuniform and irregular (i.e., possibly noncircular) communication and sensing regions. In addition, we assume that the sensing region of every node is convex. Under these relaxed assumptions, we find that the proof of Theorem 4 in Section 3.1 is still valid after we substitute sensing circles with arbitrary convex shapes. Specifically, we have the following theorem (the proof is similar to the proof of Theorem 4 and is omitted here).

```

int is_eligible (integer  $K_s$ )
begin
  find all intersection points inside  $C(v)$ :
   $SI = \{p \mid (p \in C(u) \cap C(w) \text{ OR } p \in C(u) \cap A) \text{ AND } u, w \in SN(v) \text{ AND } p \in S(v)\}$ ;
  Find all nodes whose sensing regions
  contain  $C(v)$ :
   $SC = \{u \mid C(v) \subseteq C(u)\}$ ;
  if ( $|SI|=0$ ) {
    if ( $|SC| \geq K_s$ ) return INELIGIBLE;
    else return ELIGIBLE;
  }
  for (each point  $p \in SI$ )
  begin
    /*compute p's coverage degree*/
     $sd(p) = |\{u \mid u \in SN(v) \text{ AND } p \in S(u)\}|$ ;
    if ( $sd(p) < K_s$ ) return ELIGIBLE;
  end
  return INELIGIBLE;

```

Fig. 11. The K_s -coverage eligibility algorithm for convex sensing regions.

THEOREM 7. *For a set of nodes that K_s -cover ($K_s \geq 1$) a convex region A , Theorem 4 still holds as long as the sensing region of every node is convex.*

To accommodate the above extension to the sensing region in the eligibility algorithm of CCP (see Figure 9), the procedure to compute the intersection points of sensing circles needs to be extended to a more general algorithm that can compute the intersection points of arbitrary convex shapes. In addition, we need to consider the case where node v 's sensing region lies entirely in the sensing regions of other nodes. In such a case, v is ineligible to be active even if there are no intersection points of other sensing regions within the sensing region of node v . The modified eligibility algorithm of CCP is shown in Figure 11, where $S(v)$ represents the sensing region of node v , and $C(v)$ represents the *boundary* of $S(v)$.

5.3 Probabilistic Coverage Model

In Section 5.1 and Section 5.2, we have discussed how our results can be extended when assumptions A(2) and A(3) are relaxed. In this section, we extend CCP to a probabilistic coverage model by further relaxing assumption A(1). In Section 2, we assumed that a point inside the sensing region of node v is guaranteed to be covered by v (i.e., assumption A(1)). However, this deterministic coverage model does not capture the stochastic nature of many realistic sensing tasks in sensor networks. For example, in distributed detection applications, the probability that an event can be detected by an acoustic node depends on the distance between the event and node [Duarte and Hu 2003].

Similar to Section 5.1, we assume that nodes may have nonuniform and irregular (i.e., possibly noncircular) communication and sensing regions. Let $S(v)$ represent the sensing region of node v . We further assume that the probability that any point within $S(v)$ is *sensed* by node v is lower-bounded by P ($0 < P < 1$). P and $S(v)$ are known parameters, and the relationship between them depends

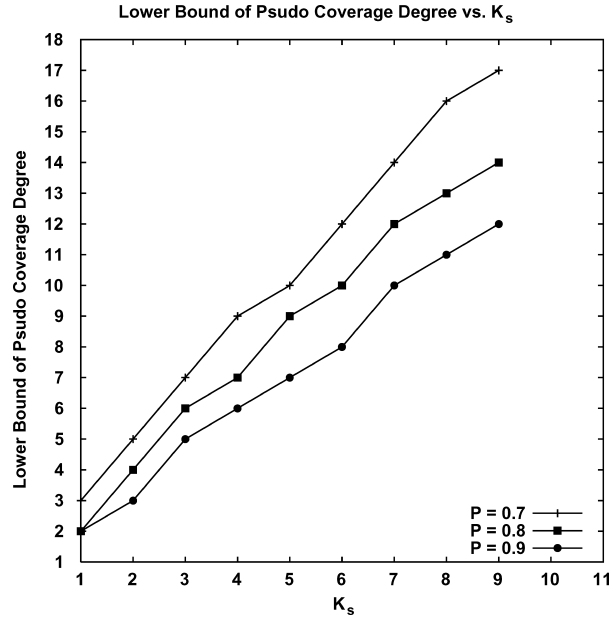


Fig. 12. The lower bound of pseudo coverage degree increases roughly linearly with K_s .

on the signal propagation properties and the characteristics of node v . Based on this probabilistic coverage model, the coverage configuration problem can be reformulated as follows. Given a convex coverage region A , and parameters K_s ($K_s \geq 1$) and β ($0 < \beta \leq 1$) specified by the application, we must maximize the number of sleeping nodes under the constraint that the remaining nodes must guarantee that the probability at which any point in A is sensed by at least K_s nodes is no lower than β . We refer to this probabilistic coverage model as (K_s, β) -coverage.

We now show how to use CCP to provide probabilistic coverage. The central idea is to map a (K_s, β) -coverage requirement to a *pseudo coverage degree*, K'_s , as the input parameter to the original CCP algorithm shown in Figure 11. Suppose each node can sense every point within its sensing region with probability P , and CCP is executed with the input parameter K'_s to provide the coverage, the probability that a point is sensed by at least K_s nodes must be no lower than β .

$$1 - \sum_{i=0}^{K_s-1} \binom{K'_s}{i} P^i (1-P)^{K'_s-i} \geq \beta. \quad (2)$$

When P , K_s , and β are known, the lower bound of pseudo coverage degree K'_s can be derived from (2) which is then used as input to CCP to achieve the probabilistic sensing coverage over convex deployment region A .

Figure 12 shows the lower bound on the pseudo coverage degree computed from (2) for different K_s when $\beta = 0.95$ and P varies from 0.7 to 0.9. We can see that the pseudo coverage degree increases roughly linearly as a function of K_s . This result indicates that CCP can effectively support the (K_s, β) coverage model for applications that require high degrees of probabilistic coverage.

5.4 Applying A Probabilistic Coverage Model to Distributed Sensing Applications

The (K_s, β) -coverage model defined in Section 5.3 is applicable to a number of real-world sensing applications. As an example, we discuss in this section how to apply the (K_s, β) -coverage model to a distributed target detection application based on the Constant False Alarm Rate (CFAR) detector [Varshney 1996]. In CFAR detector, each node sends 1 to a *fusion node* if its sensor reading exceeds a decision threshold λ and sends 0 otherwise. The overall decision at the fusion node is obtained from fusing the binary decisions of multiple nodes using a *fusion rule*. A *false alarm* occurs when the fusion node decides on 1 while no target is present. The goal of the application is to choose the minimum number of active nodes in a geographic region such that any point in the region has a detection probability higher than a threshold β , while the overall false alarm rate is below a threshold α .

The (K_s, β) -coverage model can be directly mapped to the CFAR detector whose overall detection probability required by the application is β , and the fusion rule is K_s out of K'_s , that is, the fusion node decides on 1 when there are at least K_s nodes out of the total K'_s nodes reporting 1 to the fusion node. Hence (2) describes the relation between the overall detection probability β and individual nodes' detection probability P using the K_s out of K'_s rule. Similarly, the relation between the overall false alarm rate α and the decision threshold λ can be derived.² Then λ can be used by each node to perform target detection.

In order to solve K'_s from (2), the detection probability P of a node must be known. In addition, the sensing range of nodes must be obtained before running CCP to achieve K'_s coverage and hence the desired detection probability β . The sensing range of a sensor depends on the sensor modality, sensor design, and the environment. The sensing range has a significant impact on the performance of a sensing application and is usually determined empirically to satisfy the desirable Signal to Noise Ratio (SNR) or other requirement of the application (e.g., the target detection probability P in our example). Given the decision threshold λ , the characteristics of a node's detection performance versus distance can be measured through experiments.

As an example of how to measure the sensing range in reality, we now briefly discuss a real-world experiment based on sGate [Sensoria 2003], a sensor platform from Sensoria Corp., performed by Duarte and Hu [2003]. In the experiment, various military vehicles drove through the node deployment region, and the types of the vehicles were identified based on their acoustic measurements. The experimental results showed that the probability of correct vehicle classification decreases quickly with the sensor-target distance, and drops below 50% when the sensor-target distance exceeds 100 m. From such empirical results, appropriate detection probability P and sensing range R_s can be chosen. Note that there is a fundamental trade-off between the detection performance and cost. Although choosing a conservative sensing range always leads to a higher detection probability at each node, more nodes are needed to cover the region of

²The detailed discussion is beyond the scope of this article. Similar derivation, based on the majority fusion rule, is presented in Xing et al. [2003].




Fig. 13. The procedure of applying the (K_s, β) probabilistic coverage model to a distributed target detection application based on CFAR detector. The application requires the detection probability of every point in the region to be at least β . A target is detected if at least K_s nodes out of K'_s nodes detect the target and report 1 to the fusion node.

interest. When the sensing ranges of nodes are irregular, the approximations discussed in Section 5.1 can be applied.

In summary, Figure 13 illustrates the procedure of applying the (K_s, β) -coverage model to the target detection application. We note that a more complex fusion rule than K_s out of K'_s , for example, the distance-based fusion rule proposed by Duarte and Hu [2003], can be used to achieve better detection performance and hence fewer active nodes. Further discussion on this topic is beyond the scope of this article.

5.5 The Effect of Location Accuracy

We have assumed so far that each node knows its accurate location (i.e., assumption A(5)). CCP can also be extended to tolerate bounded location errors. In this section, we assume that the location error (defined as the distance between the actual location of a node and its estimated location) is upper-bounded by δ . In addition, we assume nodes may have different circular sensing ranges. We then have the following theorem.

THEOREM 8. *If the location error is upper-bounded by δ , and the set of nodes activated by CCP can K_s -cover a convex region A assuming all nodes locate on their estimated locations and each node uses a sensing range δ shorter than its actual sensing range, the same set of nodes K_s -covers A with their actual sensing ranges when they are at their actual locations.*

PROOF. We prove by contradiction. Suppose the nodes activated by CCP cannot K_s -cover A with actual sensing ranges when there is no location error. There must exist a point p in A that is covered by less than K_s nodes. On the other hand, p must be covered by at least K_s active nodes u_1, u_2, \dots, u_{K_s} , which have estimated locations and sensing ranges δ shorter than the actual sensing ranges. We have (let a and a' represent the actual and estimated locations of point a , respectively):

$$|u'_i p| < R_s(u_i) - \delta, 1 \leq i \leq K_s$$

where $R_s(u_i)$ represents the sensing range of u_i . Since $|u_i u'_i| \leq \delta$, ($1 \leq i \leq K_s$), from triangle inequality we have:

$$|u_i p| \leq |u'_i p| + |u_i u'_i| < R_s(u_i), 1 \leq i \leq K_s.$$

Hence p is covered by u_1 to u_{K_s} that have the actual locations and sensing ranges which contradicts our assumption that the coverage degree of p is smaller than K_s . \square

Based on Theorem 8, a network can achieve desired coverage by executing CCP with conservative sensing ranges when there are bounded estimated location errors.

6. EXPERIMENTATION

In this section, we present the results of three sets of simulation experiments. Experiment I tests CCP's capability to provide different degrees of coverage. Experiment II evaluates CCP and CCP+SPAN in terms of both coverage and connectivity on NS-2. Experiment III tests the system lifetime of CCP+SPAN protocol.

6.1 Experiment 1: Coverage Configuration

Experiment 1 is performed on the Coverage Simulator (CS) provided by the authors of Tian and Georganas [2002]. Although CS is a simple simulation environment that assumes perfect wireless communication and does not account for communication overhead, this lightweight simulator allows us to evaluate CCP's eligibility algorithm over a wide range of network settings. It has also been shown to provide similar coverage performance results to NS-2 when evaluating the coverage preservation protocol developed by the University of Ottawa [Tian and Georganas 2002].

Experiment 1 compares the performance of CCP to the Ottawa protocol described in Tian and Georganas [2002]. Similar to CCP, the Ottawa protocol is a decentralized protocol designed to preserve coverage while turning off redundant nodes to conserve energy in a sensor network. Simulation results reported in Tian and Georganas [2002] also demonstrated that this protocol can provide better coverage than the PEAS protocol [Ye et al. 2003] which is designed to control density rather than coverage. The Ottawa protocol and CCP utilize different eligibility rules. The main advantage of CCP over the Ottawa protocol lies in its ability to configure the network to the specific coverage degree requested by an application, while the Ottawa protocol does not support different coverage configurations. In addition, our experimental results show that even when only 1-coverage is required, CCP results in a smaller number of active nodes and hence leads to more energy conservation than the Ottawa protocol. All the results in this section are based on five runs with different random network topologies. The region used for testing in Experiment 1 is $50 \text{ m} \times 50 \text{ m}$ if not specified otherwise, and the sensing range is 10m for all nodes.

6.1.1 The Efficiency of CCP. To measure coverage, we divide the entire sensing region into $1 \text{ m} \times 1 \text{ m}$ patches. The coverage degree of a patch is

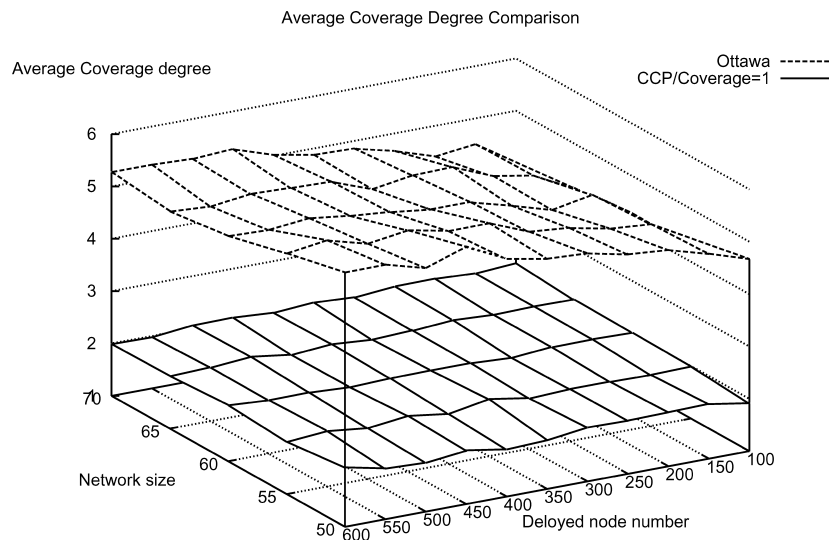


Fig. 14. Average coverage degree of all patches under CCP and the Ottawa protocol when the requested coverage degree is 1. CCP maintains an average coverage degree around 2, while the average coverage degree of Ottawa protocol is between 4 and 6 and increases with the number of nodes.

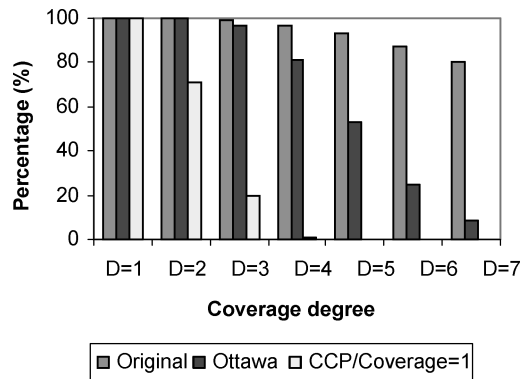


Fig. 15. Distributions of coverage degrees of all patches under CCP and the Ottawa protocol.

approximated by measuring the number of active nodes that cover the center of the patch. Figure 14 compares the average coverage degree of all patches for CCP and the Ottawa protocol. The requested coverage degree is $K_s = 1$ for CCP. The average coverage degree of CCP remains around 2 in all combinations of network size and numbers of nodes. In contrast, the Ottawa protocol results in an average coverage degree between 4 and 6 and increases with the number of nodes. Figure 15 shows the distribution of coverage degrees with 100 nodes. Each data point represents the percentage of patches with a coverage degree no lower than that specific level. The data set “Original” represents the coverage percentage of the original network. While both protocols achieve full coverage as required, the number of nodes that has unnecessarily high coverage degrees is

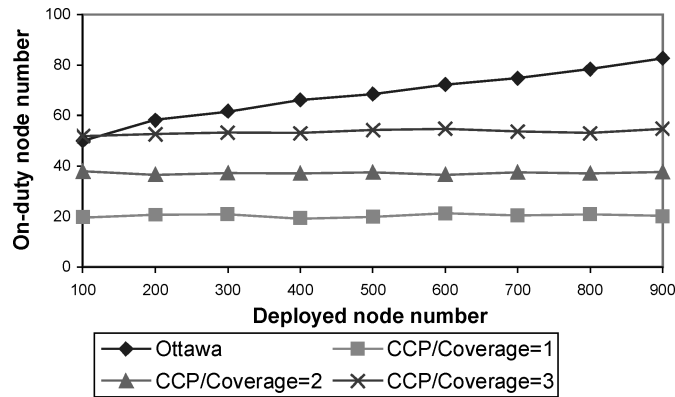


Fig. 16. The number of active nodes of CCP and the Ottawa protocol with different total number of nodes and requested coverage degrees.

significantly smaller when CCP is used. For example, while CCP results in only 1% of nodes that are 4-covered, over 80% of the patches are at least 4-covered with the Ottawa protocol. Figure 16 shows the number of active nodes under the Ottawa protocol and CCP (with different requested coverage degrees).

The number of nodes activated by CCP (when $K_s = 1$) is less than half of the number of nodes activated by the Ottawa protocol when the number of deployed nodes is 100. When the number of deployed nodes reaches 900, the number of active nodes for CCP is less than 25% of that for the Ottawa protocol. The number of nodes activated by the Ottawa protocol increases when the number of deployed nodes increases, while CCP maintains the same number of active nodes. This is because the eligibility rule in CCP makes decisions based on knowledge about the nodes within twice the sensing range, while the eligibility algorithm in the Ottawa protocol can only utilize the information nodes within the sensing range. In addition, the Ottawa protocol requires that all nodes close to the boundary of the region remain active which can lead to a large number of additional active nodes when a large number of nodes are deployed. In contrast, CCP is able to turn off redundant nodes close to the network boundary. In summary, the above experiments show that our eligibility rule can preserve coverage with fewer active nodes. That, in turn, will consume less power, and thus extend the lifetime of the network.

6.1.2 The Configurability of CCP. In this section, we evaluate CCP's ability to configure the network to achieve requested coverage degrees. In Figure 17, we plot the resulting coverage degrees under different requested coverage degrees and different numbers of deployed nodes (500, 700, and 900). The line labeled "Min-500, 700, 900" represents the minimum resulting coverage degree among all patches for different requested coverage degrees.

We can see that the minimum coverage degree is always equal to the required coverage degree and remains close to the the average coverage degree. This result demonstrates that CCP can guarantee requested degrees of coverage without introducing unnecessary redundancy. Figure 17 also shows that the




Fig. 17. Average coverage degree vs. required coverage degree under different requested coverage degrees and numbers of deployed nodes.

ratio of the average coverage degree to the minimum coverage degree *decreases* as the requested coverage degree increases. Finally, as shown in Figure 17, the number of active nodes of CCP is proportional to the degree of coverage. This allows CCP to scale to any feasible degree of coverage requested by the application.

6.1.3 Probabilistic Coverage Performance. In this section, we examine the effectiveness of the (K_s, β) probabilistic coverage model discussed in Section 5.3. In this model, when the sensing probability (P in Section 5.3) associated with the sensing range of each node is known, CCP can be run with the pseudo coverage degree K'_s computed by (2), to guarantee that the probability that every point in a region is sensed by at least K_s nodes is no lower than β . We examine in this section the redundancy in the coverage probability produced by our approach. Smaller redundancy usually leads to more energy savings (e.g., by activating fewer nodes).

We first discuss our experimental methodology. For a pair of required K_s and β , we first solve a K'_s from (2) and then run CCP to achieve K'_s coverage over a region. Then for each point in the region, we can measure an actual coverage degree. By replacing K'_s with the actual coverage degree, we can calculate a K_s from (2). We define the K_s calculated above as K_s^* which represents the actual number of nodes needed to achieve the required sensing probability β at a point under the actual coverage degree. Since the actual coverage degree is never lower than K'_s (enforced by CCP), K_s^* is no lower than K_s accordingly. The difference between K_s^* and K_s indicates the level of unnecessary redundancy in the coverage probability produced by the (K_s, β) model. In the experiment, 1000 nodes are deployed in a 400×400 m² region. CCP is run in the Coverage Simulator from the University of Ottawa [Tian and Georganas 2002]. Similar to Experiment 1, we divide the region into $1 \text{ m} \times 1 \text{ m}$ patches. The K_s^* of the center of each patch is then calculated. The result of this section is an average of five runs.




Fig. 18. K_s^* vs. K_s . Minimum K_s^* remain the same with K_s , and the average K_s^* increases roughly linearly with K_s and remain close to K_s .

Figure 18 shows the average and minimum K_s^* of all patches in the region when K_s varies from 1 to 6, and β varies from 0.8 to 0.95. We can see that all minimum K_s^* coincide with K_s which indicates that the (K_s, β) model can effectively achieve the required probabilistic coverage. The average K_s^* increases with K_s and remains close to K_s all the time. The overall result shows that the (K_s, β) model can achieve the required probabilistic coverage with reasonable redundancy.

6.2 Experiment 2: Coverage and Communication Performance

Experiment 1 has shown that CCP can provide configurable coverage by keeping a small number of nodes active. In this section, we evaluate the capability of several protocols in terms of providing integrated coverage and connectivity configuration in NS-2. The following protocols are compared:

- SPAN, obtained from MIT (<http://www.pdos.lcs.mit.edu/span/>);
- CCP, implemented by replacing the SPAN's coordinator eligibility rule with CCP's;
- SPAN+CCP, implemented by combining the eligibility rules of SPAN and CCP as described in Section 4;
- CCP-2Hop, implemented by adding the locations of a node's neighboring coordinators in its HELLO message (as described in Section 4); and
- SPAN+CCP-2Hop, SPAN+CCP with extended HELLO messages as in CCP-2Hop.

All protocols were run on top of the 802.11 MAC layer with power saving support and improvements from Chen et al. [2002]. In a $400 \text{ m} \times 400 \text{ m}$ coverage




Fig. 19. Network topology and coverage under different protocols when $R_c/R_s = 1.5$. The medium-sized dots represent source and sink nodes located at two opposite sides of the network; the large dots represent active nodes, and the small dots are inactive nodes. The sensing ranges of active nodes are represented by circles.

region, 160 nodes are randomly distributed in the field initially and remain stationary once deployed. Nodes in our simulations have a sensing range of 50 m. We used the TwoRayGround radio propagation model in all NS-2 simulations. To measure the performance of different protocols under different ratios of communication range/sensing range, we varied the communication range by setting appropriate values of the reception power threshold in the network interface. All experimental results presented in this section are averages of five runs on different randomly chosen scenarios. The requested coverage degree is $K_s = 1$ in all the experiments in this section. The period of broadcasting beacon messages is fixed to 3 seconds for all protocols.³ We present the results on coverage, delivery ratio, the number of active nodes and overhead in Section 6.2.1 to 6.2.4. The goal of our protocols is to maintain both connectivity and coverage, while reducing the number of active nodes.

6.2.1 Coverage Performance. Figure 19(a)–(c) shows the network topology and coverage produced by SPAN, CCP, and SPAN-CCP-2Hop for $R_c/R_s = 1.5$ after 300 seconds of simulation time in 3 typical runs. As expected, SPAN leaves some areas (close to the boundary, as shown in Figure 19(a)) of the region uncovered, even though it maintains network connectivity. Although CCP maintains

³A node may broadcast a beacon before the end of the current period due to state transitions, for example, a new beacon is issued when a node becomes active or withdraws from being active.




Fig. 20. The coverage ratio vs. R_c/R_s . All protocols that integrate CCP successfully maintain full coverage under all settings. The coverage of SPAN decreases as R_c/R_s increases.

both connectivity and coverage,⁴ its topology has large voids in the network causing low communication throughput. In contrast, SPAN-CCP-2Hop maintains both coverage and satisfactory connectivity topology. This example illustrates the need for integrating CCP and SPAN when $R_c/R_s < 2$.

We now present detailed performance results. The sensing range is fixed to 50 m, and the communication range varies from 25 m to 125 m in the experiments. Similar to Experiment 1, we divide the field into $1 \text{ m} \times 1 \text{ m}$ patches. A patch is covered if the center of the patch is inside the sensing circle of an active node. We define *coverage ratio* as the ratio between the number of covered patches and the total number of patches. Figure 20 shows the average coverage ratio of five protocols 300 seconds after the simulation starts. From Figure 20, we can see that CCP, CCP-2Hop, SPAN+CCP, and SPAN+CCP-2Hop can maintain a coverage ratio close to 100% for all R_c/R_s ratios. Specifically, a majority of the coverage numbers is 100%, and all remaining numbers are above 99.99%. After a further investigation, we found this is because in some rounds of experiments, the 160 nodes randomly distributed in the original network do not provide 100% coverage to the deployment region. The overall results show that CCP can effectively maintain required coverage. The coverage ratio provided by SPAN increases when the R_c/R_s ratio drops and reaches about 96% when $R_c/R_s = 1$. This is because when the radio radius drops, network connectivity decreases accordingly, and SPAN activates more communication coordinators to maintain the communication capacity. Since SPAN does not consider coverage requirement at all, it fails to achieve full coverage in any of the tested configurations. When R_c/R_s increases, the coverage ratio of SPAN drops quickly. This result shows that topology maintenance protocols alone are not able to maintain coverage.

⁴Note that this result does not conflict with Theorem 1 which states a sufficient but unnecessary condition for connectivity.

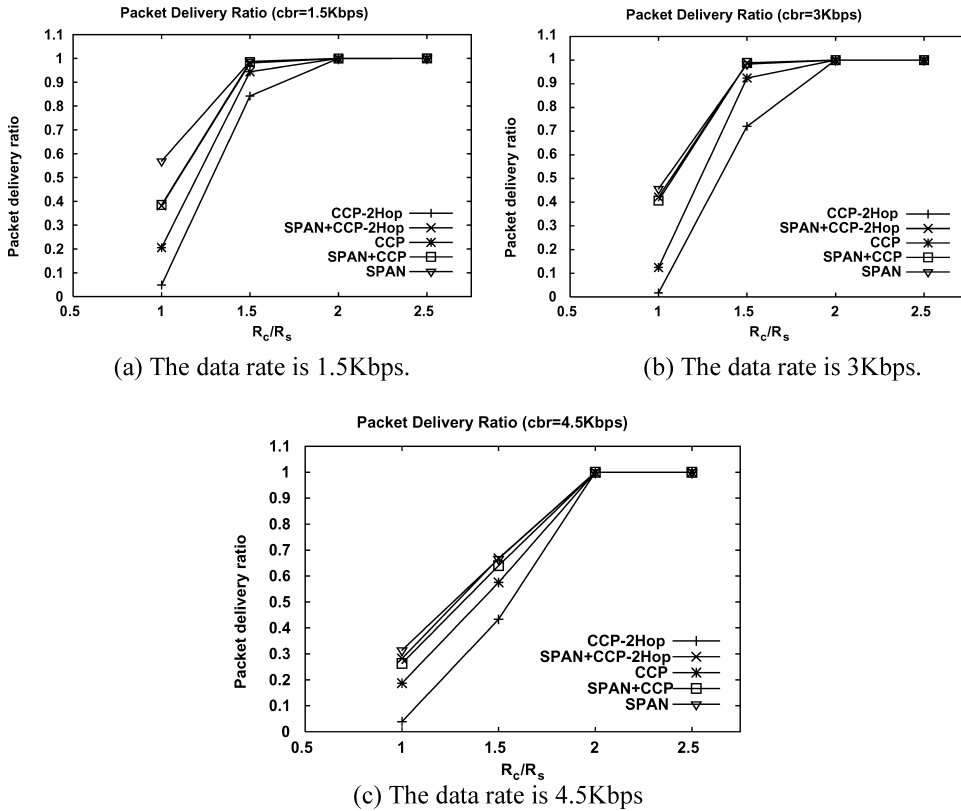


Fig. 21. Packet delivery ratio vs. R_c/R_s under different traffic loads. All protocols deliver fewer packets when R_c/R_s increases and achieve 100% delivery ratio when R_c/R_s is above 2.

6.2.2 Delivery Ratio. To test the network connectivity and communication performance, we measure the delivery ratio of the protocols under different network traffic workloads. Similar to Chen et al. [2002], to ensure that a data packet must go through multiple hops before reaching the destination, ten sources and ten sinks are randomly placed on opposite sides of the region. Each of these nodes sends a constant bit rate (CBR) flow to the destination node located on the other side of the region, and each CBR flow sends 128 byte packets. Three data rates are used in the simulations: 1.5 Kbps, 3 Kbps, and 4.5 Kbps. The routing protocol we used is the greedy geographic forwarding algorithm implemented in SPAN [Chen et al. 2002].

Figure 21(a)–(c) shows the packet delivery ratios of all protocols over 300 seconds of simulation time under 3 different data rates. The network bandwidth is 2 Mbps. First, we focus on Figure 19(a). When R_c/R_s increases, all protocols deliver more packets, and 100% of the packets are delivered when $R_c/R_s \geq 2$. This is because, when the communication range increases, the network becomes effectively denser and achieves a higher connectivity. Although CCP does not explicitly maintain connectivity, it provides good connectivity and achieves a 100% delivery ratio when $R_c/R_s \geq 2$. This result conforms to our geometric

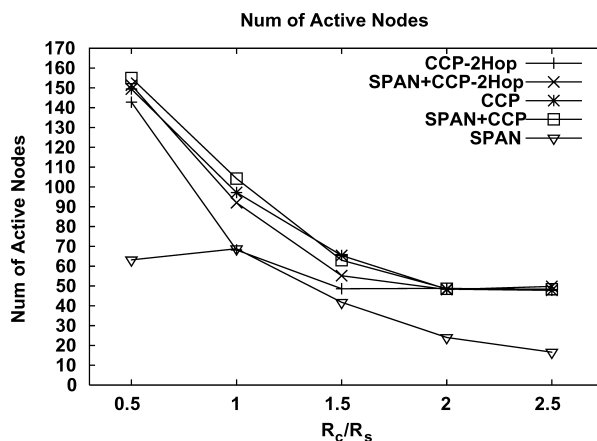


Fig. 22. Number of active nodes vs. R_c/R_s . The protocols that integrate CCP result in more active nodes in order to achieve coverage. Two-hop protocols result in fewer active nodes than their one-hop counterparts due to the knowledge of more active nodes within the sensing neighborhood.

analysis. When $R_c/R_s < 2$, CCP-2Hop has the worst delivery ratio since coverage does not guarantee connectivity in this case. CCP performs slightly better than CCP-2Hop since it produces more active nodes due to the lack of location information about two-hop active neighbors (see Figure 22). All three remaining protocols perform similarly since SPAN provides better network connectivity by activating more nodes. When $R_c/R_s = 1$, the network connectivity becomes extremely low, and none of the protocols (including SPAN) can deliver more than 50% of the packets. We found that most packet drops are due to network holes, that is, local minima of greedy forwarding when a routing node cannot find an active neighbor closer to the destination than itself. This result suggests that more complex routing schemes (e.g., geometric face routing algorithms [Kuhn et al. 2003] designed to handle network holes) are more appropriate when $R_c/R_s < 2$.

As shown in Figure 21(b)–(c), when $R_c/R_s < 2$, all protocols perform worse when the data rate increases because more packets are dropped due to buffer overflows on the routing paths. In this case, the delivery ratios of CCP and CCP-2Hop are consistently lower than those of the protocols based on SPAN. This result shows the need for explicit consideration of both connectivity and coverage in order to achieve both guarantees.

6.2.3 The Number of Active Nodes. Figure 22 shows the number of active nodes for five protocols. When R_c/R_s increases, the effective network density increases accordingly, and all protocols except SPAN activate fewer nodes. SPAN results in the least active nodes since it only maintains connectivity. When R_c/R_s decreases from 2.5 to 1, SPAN activates more nodes to maintain network connectivity. When R_c/R_s is 0.5, however, the number of active nodes for SPAN does not increase because many nodes are disconnected and hence are turned off by SPAN. SPAN+CCP and CCP perform similarly and result in the most active nodes. The 2-hop protocols outperform one-hop protocols when $R_c/R_s < 2$. This

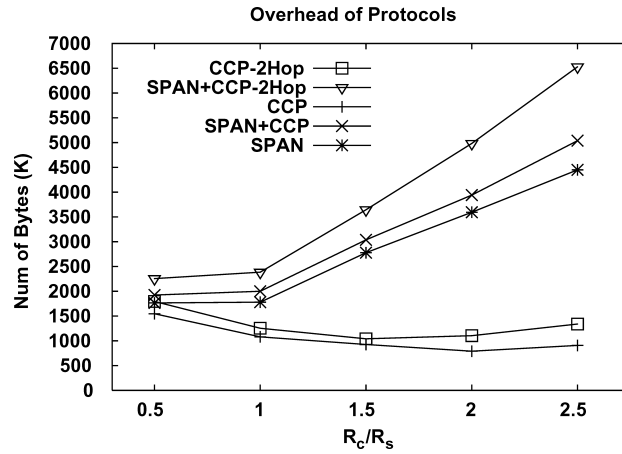


Fig. 23. The total number of bytes in the beacons of several protocols over 300 seconds. CCP and CCP-2Hop incur much lower overheads than other protocols. The difference between the overheads of the two-hop protocols and their one-hop counterparts increases with R_c/R_s .

matches our expectation since in 2-hop protocols each node bases its decision on the knowledge of more active nodes in its sensing neighborhood. Also in this region, SPAN+CCP-2Hop keeps more nodes active than CCP-2Hop because the active nodes selected by the CCP eligibility rule might not communicate via one hop, and SPAN thus activates extra nodes to provide better connectivity. Note Figure 22 shows that the extra nodes activated by SPAN+CCP-2Hop are necessary in order to maintain network connectivity.

When R_c/R_s exceeds 2, all protocols except SPAN perform similarly. This is because, as we have proven in Section 2.1, the active nodes selected by CCP can guarantee connectivity and SPAN does not take effect any more. In addition, when $R_c \geq 2R_s$, nodes can reach all coordinators in a $2R_s$ neighborhood through direct communication, and thus the 2-hop extension no longer reduces the number of active nodes.

6.2.4 Overhead. In this section, we compare the overhead of different protocols. The metric we adopted is the total number of bytes in the beacons broadcast by each protocol over 300 seconds in simulations.⁵ As in MIT's implementation of greedy geographic routing [Chen et al. 2002], a node location is represented by two 16-bit integers for all the protocols in our simulations. The overhead of a protocol depends on the number of beaconing nodes, the beacon period, and the size of each beacon. The beacon period is fixed at 3 seconds for all protocols. To better understand the results shown in Figure 23, we first compare the beacon mechanisms of SPAN and CCP. In SPAN, as required by the eligibility rule [Chen et al. 2002], each node maintains a neighbor table consisting of the IDs of its two-hop (active and sleeping) neighbors. Hence each node needs to include its ID and the IDs of its one-hop (active and sleeping)

⁵The result for SPAN only includes the SPAN-specific overhead in the beacon messages. For example, the location information in beacon messages is only used by greedy forwarding and not counted.

neighbors in the beacon messages. In contrast, a node in CCP only needs to know the locations of its *active* neighbors since the coverage of the network is solely due to active nodes. Hence only active nodes in CCP broadcast beacon messages. Each beacon message includes the location of itself (CCP) or the locations of its one-hop active neighbors (in CCP-2Hop). In the following, we first compare the overheads of CCP and SPAN and then study the impact of 2-hop neighborhood on different protocols.

CCP vs. SPAN. As shown in Figure 23, CCP and CCP-2Hop incur much smaller overheads than other protocols when R_c/R_s is larger than 1 due to the small number of active nodes. CCP and CCP-2Hop have similar overheads to other protocol when R_c/R_s is 0.5 since almost all nodes in the network become active (see Figure 22) and broadcast beacon messages as in SPAN.


All protocols that integrate SPAN (SPAN, SPAN+CCP, and SPAN+CCP-2Hop) have higher overheads when the communication range increases since the network becomes denser and each node has more neighbors resulting in more bytes in each beacon message. In contrast, the overheads of CCP and CCP-2Hop become smaller when R_c/R_s increases from 0.5 to 1.5. This is because the number of active nodes drops quickly when R_c increases (see Figure 22), resulting in fewer beacon messages. On the other hand, the overheads of CCP and CCP-2Hop increase slightly when R_c/R_s increases from 2 to 2.5 since each active node has more neighbors, resulting in more bytes in each beacon message while the total number of active nodes remains similar (see Figure 22).

Impact of 2-Hop Beacons. We now discuss the impact of 2-hop beacons on different protocols. For the protocols that integrate SPAN (SPAN+CCP, SPAN+CCP-2Hop), the difference between the overhead of the two and one-hop implementations increases with R_c/R_s since the number of two-hop neighbors of a node grows quicker than the number of one-hop neighbors when R_c/R_s increases. In contrast, the difference between CCP and CCP-2Hop remains small until R_c/R_s reaches 2. This is because, although each node may have more active neighbors in CCP-2Hop and hence larger beacon messages, the total number of active nodes of CCP-2Hop is smaller than its one-hop counterpart (see Figure 22). When R_c/R_s is larger than 2, CCP-2hop produces a similar number of active nodes to CCP and hence has a considerably higher overhead due to larger beacon messages.

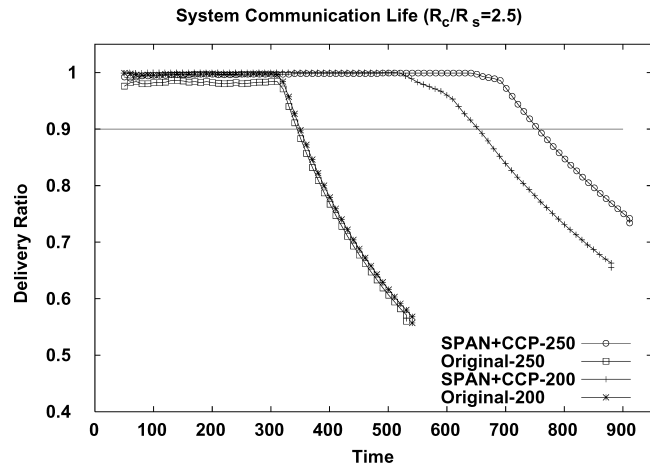
In summary, the key results in this section show that (1) CCP and CCP-2Hop have much lower overheads than other protocols, and (2) the difference between the overheads of CCP and CCP-2Hop grows slowly with R_c/R_s .

6.3 Experiment 3: System Lifetime

This section shows that SPAN+CCP can extend the system lifetime significantly while maintaining both coverage and communication capacity. The metrics used in evaluating system lifetime are the coverage and the communication lifetime. The overall system lifetime is the continuous operational time of the system before either the coverage or delivery ratio drops below specified



(a) System coverage lifetime.

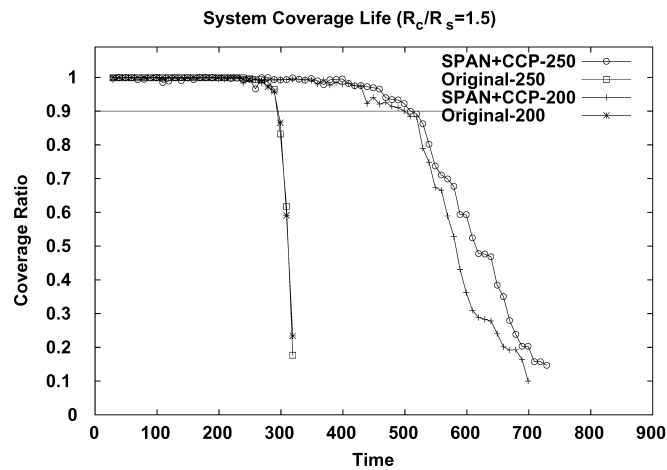


(b) System communication lifetime.

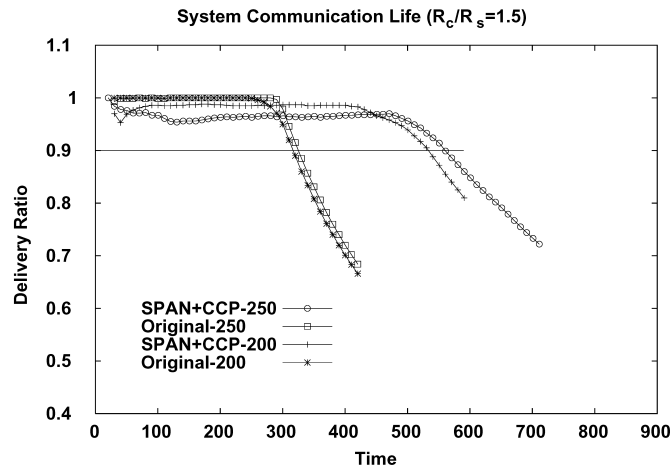
Fig. 24. System lifetime of original network and the network with SPAN+CCP when $R_c/R_s = 2.5$. SPAN+CCP can significantly improve both system coverage and communication lifetimes.

thresholds. In this section, we define both the coverage ratio threshold and the delivery ratio threshold to be 90%.

Figure 24 and Figure 25 show the system coverage and communication lifetimes of SPAN+CCP and the original network where all nodes are active when R_c/R_s is 2.5 and 1.5, respectively. In these experiments, each of 20 source and sink nodes starts with 5000 Joules of energy. Each source node sends a CBR traffic with 3 Kbps rate. Two node deployment densities, 200 and 250, are used for the remaining nodes in the experiments. With each density, the nodes are randomly distributed in a $400 \times 400 \text{ m}^2$ network field and each of them starts with an initial energy selected randomly within the range from 200J to 300J. The coverage ratio and delivery ratio were sampled from the simulations every 10 seconds. We used the energy model of Cabletron Roamabout 802.11 DS



(a) System coverage lifetime.



(b) System communication lifetime

Fig. 25. System lifetime of the original network and the network with SPAN+CCP when $R_c/R_s = 1.5$. The coverage lifetime remains similar to the case of $R_c/R_s = 1.5$, while the communication lifetime becomes shorter.

High Rate network card operating at 2 Mbps in base station mode, measured in Chen et al. [2002]. The power consumption of T_x (transmit), R_x (receive), Idle and Sleeping modes are 1400 mW, 1000 mW, 830 mW, and 130 mW, respectively [Chen et al. 2002].

Figure 24(a) and (b) show the system coverage and communication lifetimes when R_c/R_s is 2.5. First, we look at the results for 200 nodes. As shown by Figure 24(a), the coverage ratio of the original networks drops below 90% at 270s and keeps dropping quickly thereafter because a majority of nodes have run out of energy. In comparison, SPAN+CCP keeps the coverage ratio above 90% until 470s (the slight fluctuation in coverage ratio under CCP is due to the transient effect when an active node runs out of energy). As shown in Figure 24(b),

the delivery ratio of the original networks drops below 90% at 330s which is slightly longer than the system coverage lifetime. In comparison, the delivery ratio of SPAN+CCP drops below 90% at 650s with node density 200.

Overall, SPAN+CCP improves the coverage and communication lifetimes by 74% and 97%, respectively.

As expected, SPAN+CCP achieves longer lifetimes when the number of nodes increase to 250. However, the increase in system lifetime is not proportional to the increase in node density. A similar result is also reported for SPAN [Chen et al. 2002]. This is because the sleeping nodes operating in 802.11 Power Saving Mode must wake up to listen to beacons periodically and consume considerable energy.

Figure 25(a) and (b) show the system coverage and communication lifetimes when R_c/R_s is 1.5. SPAN+CCP again achieves significant improvement in coverage and communication lifetimes. Compared to the results when $R_c/R_s = 2.5$, the system coverage lifetime of SPAN+CCP remains similar, while the communication lifetime becomes shorter. This result is expected because more active nodes are needed to maintain the network connectivity when R_c/R_s falls below 2.

6.4 Summary of Simulation Results

In summary, the key results of our experiments are as follows.

- Coverage Efficiency.* CCP can provide one-coverage while keeping a significantly smaller number of active nodes than the Ottawa protocol. The number of active nodes remains steady with respect to network density for the same requested coverage degree.
- Coverage Configuration.* The CCP eligibility algorithm can effectively enforce different coverage degrees specified by the application. The number of active nodes remains proportional to the requested coverage degree.
- Integrated Coverage and Connectivity Configuration.* When $R_c/R_s \geq 2$, all protocols that employ CCP perform well in terms of packet delivery ratio, coverage, and the number of active nodes. When $R_c/R_s < 2$, CCP+SPAN-2Hop is the most effective protocol that provides both sufficient coverage and communication. SPAN cannot guarantee coverage under all tested conditions. These empirical results match our geometric analysis.

7. CONCLUSIONS AND FUTURE WORK

This article explores the problem of energy conservation while maintaining both desired coverage and connectivity in wireless sensor networks. We provided a geometric analysis that (1) proves sensing coverage implies network connectivity when the sensing range is no more than half of the communication range; and (2) quantifies the relationship between the degree of coverage and connectivity. We developed the Coverage Configuration Protocol (CCP) that can achieve different degrees of coverage requested by applications. This flexibility allows the network to self-configure for a wide range of applications and (possibly dynamic) environments. We also integrate CCP with SPAN to provide

both coverage and connectivity guarantees when the sensing range is larger than half of the communication range. We further extend the analysis and CCP to handle probabilistic sensing and communication models. Simulation results demonstrate that CCP and CCP+SPAN+2Hop can effectively configure the network to achieve both requested coverage degrees and satisfactory communication capacity under different ratios of sensing/communication ranges as predicted by our geometric analysis. In the future, we will extend our solution to handle more sophisticated coverage models and connectivity configuration and develop adaptive coverage reconfiguration for energy-efficient distributed detection and tracking techniques.

ACKNOWLEDGMENTS

We thank Nicolas D. Georganas and Di Tian at University of Ottawa for providing the source code of Coverage Simulator, Benjie Chen at MIT for making the SPAN simulation code available on the Web, and Loren Clare for his feedbacks on this article.

REFERENCES

- BAHL, P. AND PADMANABHAN, V. N. 2000. Radar: An in building RF-based user location and tracking system. In *Proceedings of the IEEE 19th Annual Joint Conference (INFOCOM'00)* (April), Tel Aviv, Israel, 775–784.
- CERPA, A. AND ESTRIN, D. 2002. ASCENT: Adaptive self-configuring sensor networks topologies. In *Proceedings of the IEEE 21st Annual Joint Conference (INFOCOM'02)* New York, NY.
- CHAKRABARTY, K., IYENGAR, S. S., QI, H., AND CHO, E. 2002. Grid coverage for surveillance and target location in distributed sensor networks. *IEEE Trans. Comput.* 51, 12, 1448–1453.
- CHEN, B., JAMIESON, K., BALAKRISHNAN, H., AND MORRIS, R. 2002. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *ACM Wirel. Netw. J.* 8, 5, 481–494.
- COUQUEUR, T., PHIPATANASUPHORN, V., RAMANATHAN, P., AND SALUJA, K. K. 2002. Node deployment strategy for target detection. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor networks and Applications* (Sept.), Atlanta, GA.
- CLOUQUEUR, T., RAMANATHAN, P., SALUJA, K. K., AND WANG, K.-C. 2001. Value-fusion versus decision-fusion for fault-tolerance in collaborative target detection in sensor networks. In *Proceedings of the 4th International Conference on Information Fusion* (Aug.), Annapolis, MD.
- D’COSTA, A. AND SAYEED, A. 2003. Collaborative signal processing for distributed classification in sensor networks. *The 2nd International Workshop on Information Processing in Sensor networks (IPSN’03)* (April), Palo Alto, CA.
- DUARTE, M. AND HU, Y. 2003. Distance based decision fusion in a distributed wireless sensor network. In *Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks (IPSN’03)* (April), Palo Alto, CA. 392–404.
- HALL, P. 1998. *Introduction to the Theory of Coverage Processes*. John Wiley & Sons Inc., New York, NY.
- HIGHTOWER, J. AND BORRIELLO, G. 2001. Location systems for ubiquitous computing. *IEEE Comput.* 34, 8, 57–66.
- KUHN, F., WATTENHOFER, R., AND ZOLLINGER, A. 2003. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *Proceedings of the 4th ACM International Symposium on Mobile Ad-Hoc Networking and Computing (MobiHoc)* Annapolis, MD.
- LI, D., WONG, K., HU, Y. H., AND SAYEED, A. 2002. Detection, classification and tracking of targets in distributed sensor networks. *IEEE Signal Process. Mag.* 19, 2, 17–29.
- MEGUERDICHIAN, S., KOUSHANFAR, F., QU, G., AND POTKONJAK, M. 2001a. Exposure in wireless ad hoc sensor networks. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom’01)* (July), Rome, Italy, 139–150.

- MEGUERDICHIAN, S., KOUSHANFAR, F., POTKONJAK, M., AND SRIVASTAVA, M. 2001b. Coverage problems in wireless ad-hoc sensor networks. In *Proceedings of the IEEE 20th Annual Joint Conference (INFOCOM'01)* (April) Anchorage, AK. 1380–1387.
- MEGUERDICHIAN, S. AND POTKONJAK, M. 2003. Low power 0/1 coverage and scheduling techniques in sensor networks. UCLA Tech. Rep. 030001 (Jan.).
- CROSSBOW. 2003. Mica wireless measurement system datasheet. Available at <http://www.xbow.com>.
- PATTEM, S., PODURI, S., AND KRISHNAMACHARI, B. 2003. Energy-quality tradeoffs for target tracking in wireless sensor networks. *The 2nd Workshop on Information Processing in Sensor networks (IPSN'03)* (April), Palo Alto, CA.
- SAVARESE, C., RABAAY, J. M., AND BEUTEL, J. 2001. Locationing in distributed ad-hoc wireless sensor networks. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP'01)* Salt Lake City, UT.
- NODEIA 2003. Sgate datasheet.
- SILVENTOINEN, M. J. AND RANTALAINEN, T. 1995. Mobile station locating in GSM. In *IEEE Wireless Communication System Symposium* (Nov.), Long Island, NY.
- TIAN, D. AND GEORGANAS, N. D. 2002. A coverage-preserved node scheduling scheme for large wireless sensor networks. In *Proceedings of the 1st International Workshop on Wireless Sensor networks and Applications (WSNA'02)* (Sept.), Atlanta, GA.
- VARSHNEY, P. 1996. *Distributed Detection and Data Fusion*. Springer-Verlag, New York, NY.
- WANG, X., XING, G., ZHANG, Y., LU, C., PLESS, R., AND GILL, C. D. 2003. Integrated coverage and connectivity configuration in wireless sensor networks. In *Proceedings of the 1st ACM Conference on Embedded Networked Node Systems (SenSys'03)* (Nov.), Los Angeles, CA.
- WYLIE, M. AND HOLTZMAN, J. 1996. The non-line of sight problem in mobile location estimation. In *Proceedings of the 5th International Conference on Universal Personal Communications*, Cambridge, MA. 827–831.
- XING, G., LU, C., PLESS R., AND O'SULLIVAN, J. A. 2004. Co-Grid: An efficient coverage maintenance protocol for distributed sensor networks. In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN'04)* (April), Berkeley, CA.
- XU, Y., HEIDEMANN, J., AND ESTRIN, D. 2002. Adaptive energy-conserving routing for multihop ad hoc networks. Res. Rep. 527, USC/Information Sciences Institute (Oct.).
- XU, Y., HEIDEMANN, J., AND ESTRIN, D. 2001. Geography-informed energy conservation for ad hoc routing. *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'01)* (July), Rome, Italy.
- YAN, T., HE, T., AND STANKOVIC, J. A. 2003. Differentiated surveillance service for sensor networks. *1st ACM Conference on Embedded Networked Node Systems (SenSys'03)* (Nov.), Los Angeles, CA.
- YE, F., ZHONG, G., LU, S., AND ZHANG, L. 2003. PEAS: A robust energy conserving protocol for long-lived sensor networks. *The 23rd International Conference on Distributed Computing Systems (ICDCS'03)* (May), Providence, RI.
- ZHAO, J. AND GOVINDAN, R. 2003. Understanding packet delivery performance in dense wireless node networks. In *Proceedings of 1st ACM Conference on Embedded Networked Node Systems* (Nov.), Los Angeles, CA.

Received January 2004; revised September 2004; accepted December 2004