

Integrated Coverage and Connectivity Configuration in Wireless Sensor Networks

Xiaorui Wang, Guoliang Xing, Yuanfang Zhang*, Chenyang Lu, Robert Pless, Christopher Gill

Department of Computer Science and Engineering

Washington University in St. Louis

St. Louis, MO 63130-4899

{wang, xing, yfzhang, lu, pless, cdgill}@cse.wustl.edu

ABSTRACT

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 degrees of coverage and connectivity in order to support different applications and environments with diverse requirements. This paper 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 in a unified framework: this is in sharp contrast to several existing approaches that address the two problems in isolation. 3) Finally, we integrate CCP with SPAN to provide both coverage and connectivity 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SenSys'03, November 5–7, 2003, Los Angeles, California, USA.

Copyright 2003 ACM 1-58113-707-9/03/0011...\$5.00.

General Terms

Algorithms, Design, Experimentation

Keywords

Sensor Network, Wireless Ad Hoc Network, Coverage, Connectivity, Energy Conservation, Topology Maintenance, Network Geometry

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 are critical requirements 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 [15] requires every location be monitored by multiple nodes, and distributed tracking and classification [9] requires even higher degrees of coverage. The coverage requirement 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 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.

* The first three authors are listed in alphabetic order

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 -connected, then for *any* possible $K-1$ 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, *i.e.*, 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 the inactive nodes can allow a sensor network to heal after a fault that reduces its connectivity, by activating particular inactive nodes. Alternatively, transient communication disruption can be avoided by maintaining greater connectivity among active nodes. Greater 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 close-to-optimal solution based on global information. Both [2] and [12] apply linear programming techniques to select the minimal set of active nodes for maintaining coverage. More sophisticated coverage model is used to address exposure-based coverage problems in [10][11]. The *maximal breach path* and *maximal support path* in a sensor network are computed using Voronoi diagram and Delaunay Triangulation techniques in [10]. The problem of finding the *minimal exposure path* is addressed in [11]. In [5], sensor deployment strategies were investigated to provide sufficient coverage for distributed detection. Provided scalability and fault-tolerance, localized algorithms are more suitable and robust for large-scale wireless sensor network that operate in dynamic environments. The protocol proposed in [14] depends on local geometric calculation of sponsored sectors to preserve sensing coverage. None of the above coverage maintenance protocols addresses the problem of maintaining network connectivity. On the other hand, several other protocols (*e.g.*, ASCENT [1], SPAN [3], AFECA [16], and GAF [17]) aim to maintain network connectivity, but do not guarantee sensing coverage. Unfortunately, satisfying only coverage or connectivity alone is not sufficient for a sensor network to provide sufficient service. Without sufficient 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 multi-hop 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 for the global algorithm in [2]) is that they can only provide a fixed degree of coverage. They cannot dynamically reconfigure to meet different coverage requirements of applications. Finally, while the PEAS [18] 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 paper are as follows. We provide a geometric analysis of the fundamental relationship between coverage and connectivity. 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 in isolation. We present a Coverage Configuration Protocol (CCP) that can dynamically configure the network to provide *different* feasible degrees of coverage 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. We integrate CCP with a representative connectivity maintenance protocol (SPAN [3]) to provide both coverage and connectivity guarantees.

In the rest of this paper, we first formulate the problem of coverage and connectivity in Section 2. The relationship between coverage and connectivity is analyzed in Section 3. We then present the design and analysis of CCP in Section 4 and propose a simple solution to configure both coverage and connectivity based on CCP in Section 5. Extensive simulation results are presented in Section 6. We offer conclusions in Section 7.

2. PROBLEM FORMULATION

Several coverage models [10][11][12] have been proposed for different application scenarios. In this paper, we assume a point p is *covered* (monitored) by a node v if their Euclidian distance is less than the sensing range of v , R_s , *i.e.*, $|pv| < R_s$. We define the *sensing circle* $C(v)$ of node v as the boundary of v 's coverage region. We assume that any point p on the sensing circle $C(v)$ (*i.e.*, $|pv| = R_s$) is not covered by v . Although this definition has an insignificant practical impact, it simplifies our geometric analysis in following sections. Based on the above coverage model, we define a convex region A (that contains at least one sensing circle) as having a *coverage degree* of K (*i.e.*, being *K-covered*) if every location inside A is covered by at least K nodes. Practically speaking, a network with a higher degree of coverage can achieve higher sensing accuracy and be more robust against sensing failures. The coverage configuration problem can be formulated as follows. Given a convex coverage region A , and a coverage degree K specified by the application (either before or after deployment), we must maximize the number of sleeping nodes under the constraint that the remaining nodes must guarantee A is K -covered.

Despite its simplicity, this coverage model is applicable in a number of applications. For example, it fits well with the decision fusion approach to distributed detection (*e.g.*, Bayesian Detection and Neyman-Pearson Test) [15], where each sensor sends 1 to a fusion node if it detects a target and sends 0 otherwise. The fused detection decision is based on the binary decisions of multiple sensors. The goal of distributed detection problem is to maximize the detection probability under the constraint that the false alarm

rate is less than a required threshold P_F . The solution to this constrained optimization problem at each sensor follows the form of Likelihood Ratio Test (LRT), i.e., if the ratio of the conditional probabilities of the sensor reading is larger than a LRT threshold λ decided by P_F , the sensor will report 1, otherwise report 0 to fusion center. Under some assumptions on the signal decay model and statistical distributions of event signals and background noise, the sensing range of a sensor can be derived from the LRT threshold λ , i.e., the detection probability of the events within the sensing range is maximized while the false alarm probability is less than the required threshold P_F . It should be noted that the term “sensing range” does *not* imply a hard boundary between the area where an event is *always* detected and the area where an event is *never* detected. Instead, the sensing range is often defined by the threshold of false alarm rate (probability). Therefore, the statistical nature of sensor network applications and the environments can be incorporated in the definition of sensing range. Exploring more sophisticated coverage models and corresponding applications is left as our future work.

In addition, we assume that any two nodes u and v can directly communicate with each other if their Euclidian distance is less than a communication range R_c , i.e., $|uv| < R_c$. Given a coverage region A and a sensor coverage degree K_s , the goal of an integrated coverage and connectivity configuration is maximizing 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.

3. RELATIONSHIP BETWEEN COVERAGE AND CONNECTIVITY

The first part of our investigation focuses on understanding the relationship between coverage and connectivity. Does coverage imply connectivity or vice versa so that a sensor network only needs to be configured to satisfy the stronger of the two requirements? In this section, we first derive a sufficient condition when coverage implies connectivity in a network. We then quantify the relationship between the degree of coverage and connectivity. The analysis presented in this section will serve as the foundation for an integrated solution to the problem of integrated coverage and connectivity configuration.

3.1 Sufficient Condition for 1-Coverage to Imply Connectivity

In this subsection, 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 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 easily seen 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 all locations of active nodes are connected. Hence we focus on analyzing the condition for a *covered network* to guarantee connectivity in the rest of this section.

Define the graph $G(V,E)$ to be the communication graph of a set of sensors, where each sensor in the set is represented by a node in V , and for any node x and y in V , the edge $(x,y) \in E$ if and only if the Euclidean distance between x and y , $|xy| < R_c$. Node v and u are connected in $G(V,E)$ if and only if a *network path* consisting of consecutive edges in E exists between node u and v .

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

Proof: For any two nodes u and v in region A , let P_{uv} be the line segment joining them. Since region A is convex, P_{uv} remains entirely within A . Hence any point on P_{uv} is at least 1-covered. Each point p on P_{uv} has a set of one or more closest sensors equidistant from p . A finite sequence $S_{uv} = s_1..s_n$ of closest sensor sets can be constructed for contiguous segments 1..n of P_{uv} , where a segment is defined by all points within it having the same set of closest sensors. S_{uv} starts with $s_1 = \{u\}$ and ends with $s_n = \{v\}$, with intervening sets possibly containing other sensors.

The distance from each point on the line segment P_{uv} to its closest sensor(s) is always less than R_s , as otherwise the path would go through regions that are not sensor-covered. Suppose x and y are two sensors in any two consecutive sets s_j and s_{j+1} in S_{uv} , $x \in s_j$ and $y \in s_{j+1}$. We now prove $|xy| < 2R_s$ by contradiction. Suppose $|xy| \geq 2R_s$, we will show that there exists at least one point on P_{uv} that is not covered. Since x and y are in two consecutive sets s_j and s_{j+1} , there must exist a point p on P_{uv} that is equally closest to x and y among all sensors. Clearly p is the intersection of P_{uv} and the perpendicular bisector of line segment xy . Since $|xy| \geq 2R_s$ and p is equidistant from x and y , from the triangle inequality, p is at least R_s from both x and y and not covered by x or y from the definition of sensing circle in Section 2. This contradicts the assumption that every point on P_{uv} is covered. Therefore, the distance between every pair of sensors in consecutive sets in S_{uv} is less than $2R_s$, and is thus less than R_c , so an edge exists between them in the communication graph. Because each set in S_{uv} contains at least one sensor, we can thus construct a communication path from u to v through each combination of node choices in the sets in S_{uv} . i.e., the communication graph of sensors in region A is connected. \square

Therefore, Theorem 1 establishes a sufficient condition for a 1-covered network to guarantee one-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.

3.2 Relationship between the Degree of Coverage and Connectivity

The previous section argues that if a region is sensor covered, then the sensors covering that region are connected as long as their communication range is no less than twice the sensing range. 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 degrees of coverage or connectivity greater than one.

We define *boundary* sensor as a sensor whose sensing circle intersects with the boundary of the convex sensor deployment region A . Clearly all *boundary* sensors are located within R_s distance to the boundary of A . All the other sensors in region A are *interior* sensors.

Lemma 1: For a K_s -covered convex region A , it is possible to disconnect a *boundary* node from the rest of the nodes in the communication graph by removing K_s sensors if $R_c \geq 2R_s$.

Proof: Consider the scenario illustrated by Figure 1: a sensor u is located at a corner (point q) of the rectangular sensor deployment region A that is K_s -covered. Suppose point p is on the sensing circle of sensor u such that pq has a 45° angle with the horizontal boundary of A .




Figure 1. Removing K_s nodes disconnects a covered network

Suppose K_s coinciding sensors are located at point p . Clearly, these K_s sensors can K_s -cover the quarter circle of sensor u . And we assume there are no other sensors whose sensing circles intersect with sensing circle of u . Then removing these K_s coinciding sensors will create an uncovered region (*i.e.*, a sensing void) surrounding sensor u . Furthermore, when R_c is equal to $2R_s$, there is no sensor within the communication range of sensor u after the removal of these K_s sensors. *i.e.*, the communication graph is disconnected. \square

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: Disconnecting the communication graph G of a set of sensors 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 in G . By Theorem 1, if it is possible to draw a continuous path between two nodes so that every point on the path is sensor-covered, then there exists a communication path between those two nodes. Therefore, to disconnect the graph it is necessary to create a sensing void, so that it is impossible to draw a continuous covered path connecting a node in V_1 to a node in V_2 . That is, 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 .




Figure 2. A disconnected network

To create a sensing void in an originally K_s -covered region A , it is clearly necessary to remove at least K_s sensors. Thus the network connectivity is at least K_s . By Lemma 1, removing K_s sensors could disconnect the communication graph. So the tight lower bound on the connectivity of communication graph is K_s . \square

Intuitively, the connectivity of the *boundary* sensors dominates the overall connectivity of the communication graph. However, in a large-scale sensor network, the *interior* sensors normally route more traffic and higher connectivity is needed for *interior* sensor to maintain the required throughput. We define *interior connectivity* as the number of sensors (either interior or boundary) that must be removed to disconnect any two *interior* sensors in the communication graph of the sensors.

Theorem 3: For a set of sensors 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 sensor deployment region, or the void merges with the boundary of the region.

Case 1: As illustrated in Figure 3, the void does not merge with the boundary. We will prove one must remove at least $2K_s+1$ sensors in this case to create such a void. We prove 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 sensor node j in W . Let's define line vj to be the direction we refer to as ‘vertical’. Now, there are at most $2K_s-1$ remaining sensors (except sensor 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 . Let's define that to be the left side. Draw the line straight left from v until it intersects the void region, and call this point p (note that p is covered by zero sensors.) Point p is at least R_s from node v , and is 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 $2K_s+1$.




Figure 3. Case 1: The void does not merge with boundary




Figure 4. Case 2: The void merges with boundary

Case 2: The void merges with the boundary of region A , as illustrated in Figure 4. 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 sensor 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 sensor can simultaneously cover both points. This implies that at least $2K_s$ sensors were removed in the K_s -covered region A to create the void. We prove this bound is tight by the following example. 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 at perpendicular bisector of A_1A_2 and have distance $(R_s+r)/2 < R_s$ with A_1A_2 and A_3A_4 respectively, as shown in Figure 4. Suppose there are K_s sensors (shown as dotted circles) located at point x and y respectively. W is composed of these $2K_s$ sensors. We assume the sensors (not shown in the figure) whose sensing circles intersect the $2K_s$ sensors 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 left side from the nodes on right side in communication graph.

From the proof of case 1 and case 2, for a set of sensors that K_s -cover a convex region, we have shown that the tight lower bound on the *interior connectivity* is $2K_s$. \square

We should note that the *interior connectivity* defined in this section is different from the connectivity of the communication sub-graph composed of solo *interior* nodes. This is because an *interior* node could connect to another *interior* node via *boundary* nodes and the communication sub-graph composed of solo *interior* nodes could be disconnected if all *boundary* nodes are removed, as illustrated by Figure 4.

From the 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 $2K_s$.

4. 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 has proven that CCP can provide guaranteed degrees of coverage.

4.1 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 5 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.

- The *sensing region* of node v is the region inside its sensing circle, *i.e.*, a point p is in v 's sensing region if and only if $|pv| < R_s$.
- A point $p \in A$ is called an *intersection point* between nodes u and v , *i.e.*, $p \in u \cap v$, 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 , *i.e.*, $p \in v \cap A$ if $|pv|=R_s$.




Figure 5. An example of K_s -eligibility

Theorem 4: A convex region A is K_s -covered by a set of sensors S if 1) there exist in region A intersection points between sensors or between sensors and A 's boundary; 2) all intersection points between any sensors are at least K_s -covered; and 3) all intersection points between any sensor and A 's boundary are at least K_s -covered.

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 is 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 $C(u)$) serving as the boundary of S , crossing this arc (*i.e.* leaving the coverage region of sensor u) would reach an area that is lower covered than point p . This contradicts with the assumption that point p has the lowest coverage degree in region A . Now we consider the following two cases:




Figure 6. A coverage patch bounded by arcs of sensing circles

- 1) The point p lies in a coverage region S whose boundary is only composed of exterior arcs of a collection of sensing circles (as Figure 6 illustrates). Furthermore, since the sensing circles themselves are outside the sensing range of the nodes that define them, the entire boundary of this coverage 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 .
- 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 **Error! Reference source not found.**, point p is in a region bounded by the exterior arcs of sensor 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 sensors u , v , w , x and intersection points between sensors w , x and 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 .




Figure 7. A coverage patch bounded by arcs of sensing circles and boundary of a coverage region

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

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 sensor 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 sensor v needs to consider all the sensors in its sensing neighbor set, $SN(v)$. $SN(v)$ includes all the active nodes that are within a distance of twice of the sensing range to v , i.e., $SN(v) = \{\text{active node } u \mid |uv| \leq 2R_s \text{ and } u \neq v\}$. If there is no intersection point inside the sensing circle of sensor v , v is ineligible when there are K_s or more sensors that are located at sensor v 's position.

The resulting coverage eligibility algorithm is shown in Figure 8. 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 requires the information about 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.

```

int is_eligible (integer  $K_s$ )
begin
  find all intersection points inside  $C(v)$ :
   $SI = \{p \mid (p \in u \cap w \text{ OR } p \in u \cap A) \text{ AND}$ 
     $u, w \in SN(v) \text{ AND } |pv| < R_s\}$ ;
  Find all coinciding sensors:
   $SC = \{u \mid |uv|=0\}$ ;
  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*/
     $sd(p) = |\{u \mid u \in SN(v) \text{ AND } |pu| < R_s\}|$ ;
    if ( $sd(p) < K_s$ ) return ELIGIBLE;
  end
  return INELIGIBLE;
end

```

Figure 8. The K_s -Coverage Eligibility Algorithm

There are two approaches to make each node aware of its multi-hop 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 multi-hop neighbors in its HELLO messages. Specifically, each node may broadcast the locations and status 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 (this approach is also used by SPAN to maintain two-hop neighborhood tables). We should note that, in a network with random topology, such HELLO messages still can't guarantee the discovery of all nodes within a distance of $2R_s$. Since including multi-hop neighbors in the HELLO messages introduce much higher communication overhead compared to a one-hop approach in a dense network, there is a tradeoff between the beacon overhead and the number of active nodes maintained by CCP. We investigate this trade-off through experiments in Section 6.2.

We note that a special case (when coverage degree $K_s = 1$) of Theorem 4 was stated in [8], but it did not provide any proof. 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 paper.

4.2 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, the node sleeps to conserve energy. In the ACTIVE state, the node actively senses the environment and communicates with other sensors. Each node periodically enters the LISTEN state to collect HELLO messages from its neighbors and reevaluates its eligibility to determine its

new state. When a network is deployed, all nodes are initially in the ACTIVE state. 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 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 regain the desired degree of coverage. In CCP a node changes its state independently based on local information. The state transition in CCP is similar to SPAN [3] and several other protocols [14][17]. We now describe the specific rules used in CCP:

- **In SLEEP:** When the sleep timer T_s expires, a node in the sleep state turns the radio on, 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 in the listen state evaluates its eligibility (see Figure 8). If it is eligible, it starts a join timer T_j , otherwise it returns to the SLEEP state. If it becomes ineligible after the join timer is started (e.g., due to the JOIN beacon from a neighbor), it cancels the join timer. If the join timer expires, the node broadcasts a JOIN beacon and enters the ACTIVE state. If the listen timer expires, it starts a sleep timer T_s and returns to the SLEEP node.
- **In ACTIVE:** When a node receives a HELLO message, it updates its sensing neighbor table and executes the coverage eligibility algorithm (see Figure 8) to determine its eligibility to remain active. If it is ineligible, it starts a withdraw timer T_w . If it becomes eligible (due to the reception of a WITHDRAW or HELLO message from a communication neighbor) before the withdraw timer expires, it cancels the withdraw timer. If T_w expires, it broadcasts a WITHDRAW message, starts a sleep timer T_s , and enters the SLEEP node.

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 variations in coverage. Both timers are also related to the density of nodes in the network. 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 should point out 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. For example, intuitively a node that will cover more uncovered area should have a shorter join timer when competing with other competing nodes. The proper ranking heuristics are left as our future work. In this paper, all nodes are deemed to share the same rank.

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

As described in Section 3, 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 inte-

grating CCP with an existing connectivity maintenance protocol, SPAN [3], to provide both sensing coverage and communication connectivity. SPAN [3] 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 states. Each node running SPAN maintains a neighborhood table that includes the location 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 [3]).

The main difference between CCP and SPAN lies in their eligibility rules. In SPAN, a non-coordinator 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 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 a 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 coordinators in 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 4.1, 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.

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 life time of CCP+SPAN protocol.

6.1 Experiment I: Coverage Configuration

Experiment I is performed on the Coverage Simulator (CS) provided by the authors of [14]. Although CS is a simple simulation environment that assumes perfect wireless communication and doesn’t account for communication overhead, this light-weight 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 University of Ottawa [14].

Experiment I compares the performance of CCP to the Ottawa protocol described in [14]. 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 [14] also demonstrated that this protocol can provide better coverage than the PEAS protocol [18], 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 I is 50m×50m if not specified otherwise, and the sensing range is 10m for all sensor nodes.

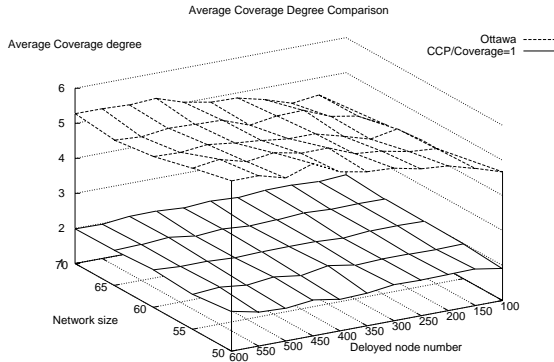


Figure 9. Average Coverage Degree

6.1.1 The Efficiency of CCP

To measure coverage, we divide the entire sensing region into 1m×1m patches. The coverage degree of a patch is approximated by measuring the number of active nodes that cover the center of the patch. Figure 9 compares the average coverage degree of all patches for CCP and the Ottawa protocol. The requested cover-

age 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 10 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 significantly smaller when CCP is used. For example, while CCP results in only 1% of nodes being 4-covered, over 80% of the patches are at least 4-covered with the Ottawa protocol. Figure 11 shows the number of active nodes under the Ottawa protocol and CCP (with different requested coverage degrees).

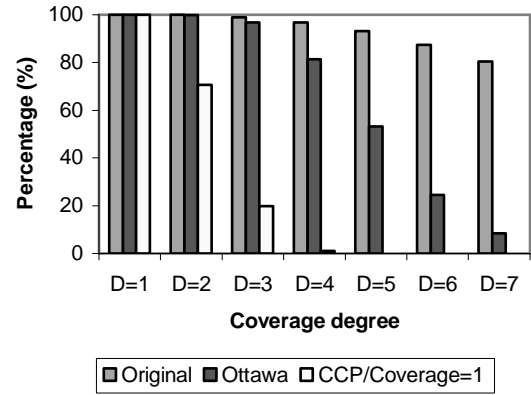


Figure 10. Distribution of Coverage Degree

The number of active nodes used 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 active nodes used 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 subsection, we evaluate CCP’s ability to configure the network to achieve requested coverage degrees. In Figure 11, we plot 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 mini-

imum resulting coverage degree among all patches for different requested coverage degrees.




Figure 11. Comparison of Active Node Number

The minimum coverage degree remains close to the requested coverage degree. This result demonstrates that CCP can guarantee requested degrees of coverage without introducing unnecessary redundancy. Figure 12 also shows that the ratio of average coverage degree to the minimum coverage degree *decreases* as the requested coverage degree increases. Finally, as shown in Figure 12, 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.




Figure 12. Coverage Degree vs. Required Coverage Degree

6.2 Experiment II: Coverage and Communication Performance

Experiment I has shown that CCP can provide configurable coverage by keeping a small number of nodes active. In this subsection, 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 5.

- CCP-2Hop: implemented by adding the locations of a node's neighboring coordinators in its HELLO message (as described in Section 4.1).
- SPAN+CCP-2Hop: SPAN+CCP with extended HELLO messages as in CCP-2Hop.

We simulated all protocols in NS-2 with the CMU wireless extensions [4]. All protocols were run on top of the 802.11 MAC layer with power saving support and improvements from [3]. In a $400 \times 400 \text{m}^2$ coverage region, 160 nodes are randomly distributed in the field initially and remain stationary once deployed. Similar to [3], to ensure a data packet must go through multiple hops before reaching the destination, ten sources and ten sinks are randomly placed in opposite sides of the region. Each of these nodes sends a CBR flow to destination node located on the other side of the region, and each CBR flow sends 128 byte packets with 3Kbps rate. The routing protocol we used is the greedy geographic forwarding algorithm described in [3]. Nodes in our simulations use radios with a 2 Mbps bandwidth and a sensing range of 50 m. We used 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 receiving 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 $K_c = 1$ in all the experiments in this section.

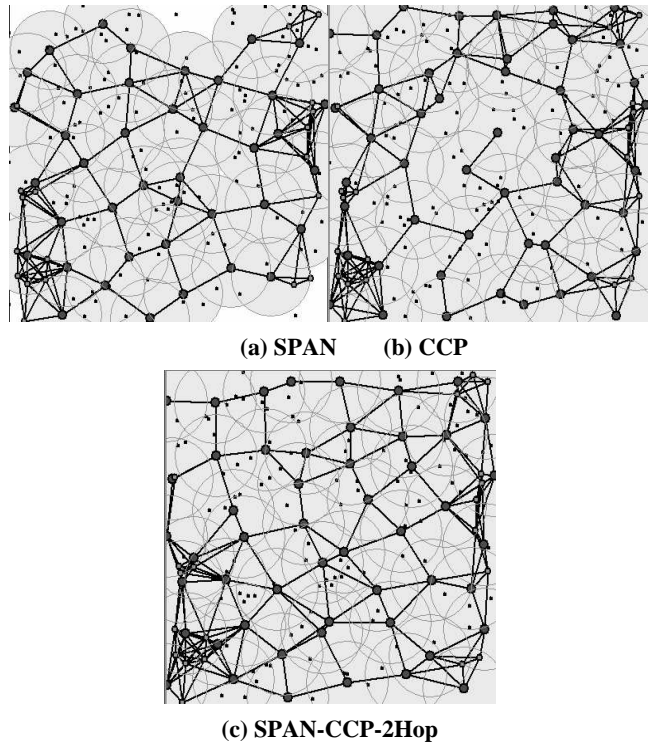


Figure 13. Network Topology and Coverage in a Typical Run ($R_c/R_s = 1.5$)

Figure 13(a-c) show 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. 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. As expected, SPAN leave some areas (close to the boundary) of the region uncovered, even though it maintains network connectivity. Although CCP maintains both connectivity and coverage², its topology has large voids in the network causing low communication throughput. In contrast, SPAN-CCP-2Hop maintains both coverage and satisfactory topology. This example illustrates the need for integrating CCP and SPAN when $R_c/R_s < 2$.

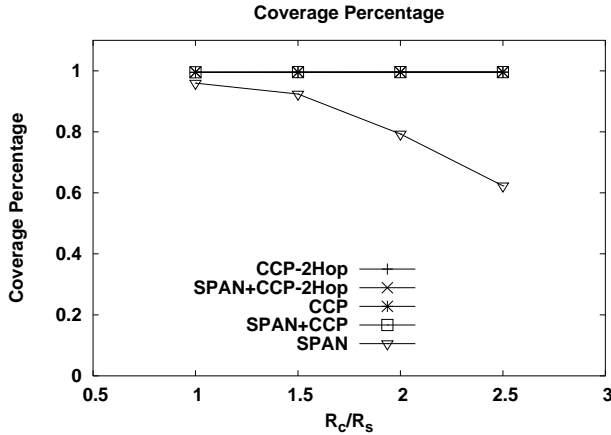


Figure 14. Coverage Degree vs. R_c/R_s

We now present detailed performance results. The goal of our protocols is to maintain both connectivity and coverage while reducing the number of active nodes. Figure 14 shows the coverage percentage of five protocols on a sensor network. The sensing range is 50m and the communication range varies from 50m to 125m. Similar to Experiment I, we divide the field into $1m \times 1m$ patches. A patch is covered if the center of the patch is inside the sensing circle of an active node.

The percentage of coverage is computed as the ratio of the number of covered patches to the total number of patches 300 seconds after the simulation starts. From Figure 14, we can see that CCP, CCP-2Hop, SPAN+CCP, SPAN+CCP-2Hop can maintain coverage percentage 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 randomly distributed sensors of the original network don't provide 100% coverage to the deployment region. The overall results show that CCP can effectively maintain coverage. The coverage percentage 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 selects 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 percentage drops quickly. This result shows that topol-

ogy maintenance protocols alone are not able to maintain coverage.

Figure 15 shows the packet delivery ratios of all protocols over 300 seconds of simulation time. When R_c/R_s increases, all protocols deliver more packets, and 100% of the packets are delivered when $R_c/R_s > 2$. This is because when the communication range increases, the network becomes effectively denser and achieves higher connectivity. When $R_c/R_s < 2$, CCP-2Hop shows the worst delivery ratio since it only considers the coverage requirement, which does not guarantee connectivity under these conditions. CCP performs slightly better than CCP-2Hop since it produces more active nodes and thus higher connectivity due to the lack of location information about two-hop neighboring coordinators. All three remaining protocols perform similarly since SPAN provides better communication connectivity by activating more nodes. As illustrated in Figure 16, in order to provide capacity for both coverage and communication, SPAN+CCP-2Hop produces more active nodes than CCP-2Hop. In addition, although SPAN+CCP-2Hop introduces the overhead of sending location coordinates in HELLO messages, it performs as well as the original SPAN. When R_c is decreased to 50m, the network capacity becomes extremely low and no protocols (including the original SPAN) can deliver more than 50% of the packets. Exactly as predicted by our geometric analysis, CCP provides a 100% delivery ratio when $R_c/R_s \geq 2$ even though it does not explicitly maintain the network topology.

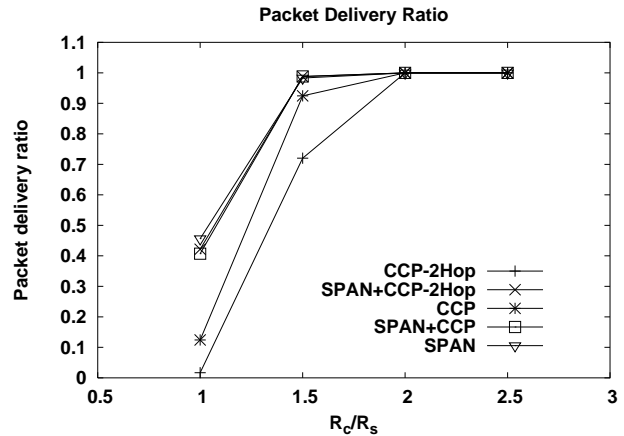


Figure 15. Packet Delivery Ratio vs. R_c/R_s

Figure 16 shows the number of active nodes of five protocols. When R_c/R_s increases, the effective network density increases accordingly, and all protocols activate fewer nodes. SPAN results in the least active nodes since it only maintains connectivity. 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 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 keep more nodes active than CCP-2Hop because the active nodes selected by CCP eligibility rule might not communicate via one hop and SPAN thus activates extra nodes to provide better connectivity. Note Figure 15 shows that the extra nodes activated by SPAN+CCP-2Hop are necessary in order to maintain network connectivity.

² Note that this result does not conflict with Theorem 1 which gives a sufficient but unnecessary condition for connectivity.




Figure 16. Number of Active Nodes vs. R_c/R_s

When R_c/R_s exceeds 2, all protocols except SPAN perform similarly. As we have proven in Section 3.1, the active nodes selected by CCP can guarantee connectivity and SPAN does not take effect any more. In addition, when $R_c > 2R_s$, 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.

6.3 Experiment III: System Life Time

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 lifetime* 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 its specified threshold. For the experiments in this section we define both thresholds to be 90%. Figure 17 and Figure 18 show the *system coverage* and *communication lifetime* of SPAN+CCP and original network with all nodes on, 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 3Kbps rate. Three node deployment densities, 200, 250 and 300 are used for the remaining nodes in the experiments. With each density, the nodes are randomly distributed in a $400 \times 400 m^2$ network field and each of them starts with an initial energy selected randomly within the range from 200 J to 300 J. The ratio of communication and sensing range is 2.5 in all experiments. We sampled the network coverage and delivery ratio from the simulation every 10 seconds. We follow the energy model of Cabletron Roamabout 802.11 DS High Rate network card operating at 2Mbps in base station mode, measured in [3]. The power consumption of Tx (transmit), Rx (receive), Idle and Sleeping modes are 1400mW, 1000mW, 830mW, 130mW respectively [3].

We can see from the Figure 17 that in the original network with all nodes on, the system coverage percentages drop below 90% at 270s with node density 200 and at 280s with densities 250 and 300, and keep dropping sharply thereafter because of a majority of nodes have run out of energy. Figure 18 illustrates similar results. The system delivery ratio drops below 90% after around 330 seconds, which is slightly longer than the system coverage lifetime.

On the other hand, as illustrated in Figure 17, SPAN+CCP keeps the coverage above 90% until 470s with node density 200, 530s with node density 250 and 560 seconds with node density 300.

We can see that the death of active nodes can cause slight fluctuations of the coverage percentage curves. However, the nodes failures do not affect the coverage percentage of original network until a majority of the nodes dies. This is because in original network with all nodes on, a large portion of the field has coverage degrees higher than 1. The system delivery ratios of SPAN+CCP drop below 90% at 650s with node density 200, at 740s with node density 250 and 730 seconds with node density 300 respectively.

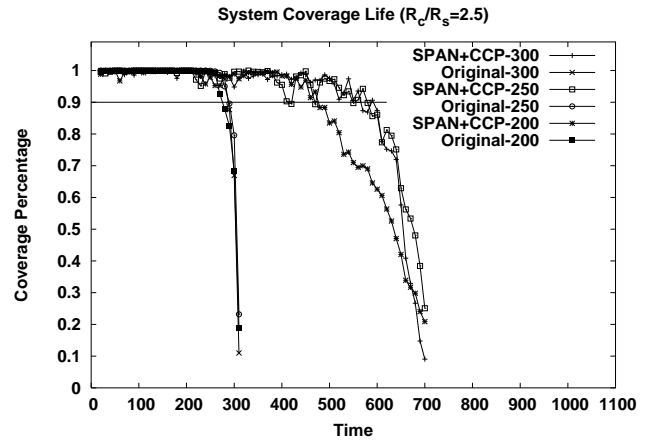


Figure 17. System Coverage Life Time

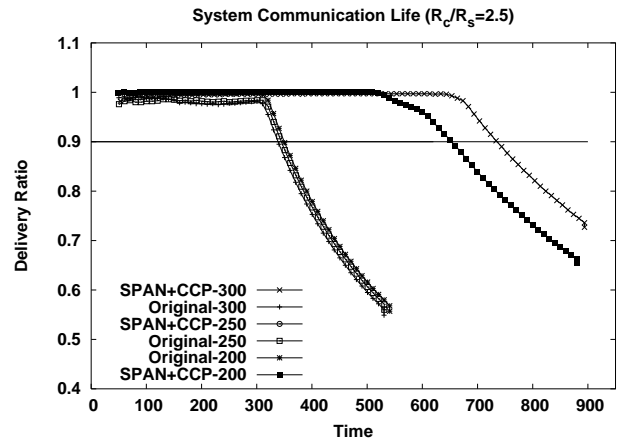


Figure 18. System Communication Life Time

We can see from the figures that the system coverage lifetime dominates the overall system lifetime since maintaining a high coverage percentage requires more active nodes than maintaining a communication backbone. As illustrated in both Figure 17 and Figure 18, the system lifetime doesn't increase much when the node density increases. Similar results are also reported in [3]. This is because the sleep nodes in 802.11 Power Saving Mode must wake up to listen to 802.11 beacons and SPAN HELLO messages periodically and consume considerable energy [3].

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 paper 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) quantify 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 the SPAN to provide both coverage and connectivity guarantees when the sensing range is higher than half of the communication range. 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.

8. ACKNOWLEDGEMENTS

We thank Nicolas D. Georganas and Dian Tian at University of Ottawa for providing the source code of Coverage Simulator, and Benjie Chen for making the SPAN simulation code available on the Web. We would also like to thank our shepherd Loren Clare and anonymous reviewers for their valuable suggestions.

9. REFERENCES

- [1] A. Cerpa and D. Estrin, "ASCENT: Adaptive Self-Configuring Sensor Networks Topologies," INFOCOM, June 2002.
- [2] K. Chakrabarty, S. S. Iyengar, H. Qi, E. Cho, "Grid coverage for surveillance and target location in distributed sensor networks," IEEE Transactions on Computers, 51(12):1448-1453, December 2002.
- [3] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, "Span: An Energy-Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks," ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2001), Rome, Italy, July 16-21, 2001.
- [4] CMU Monarch Extensions to ns. <http://www.monarch.cs.cmu.edu/>.
- [5] T. Couqueur, V. Phipatanasuphorn, P. Ramanathan and K. K. Saluja, "Sensor Deployment Strategy for Target Detection," Proceeding of The First ACM International Workshop on Wireless Sensor Networks and Applications, Sep. 2002.
- [6] T. Clouqueur, P. Ramanathan, K.K. Saluja, and K.-C. Wang. "Value-fusion versus decision-fusion for fault-tolerance in collaborative target detection in sensor networks." In Proceedings of Fourth International Conference on Information Fusion, Aug. 2001.
- [7] A. D'Costa and A. Sayeed, "Collaborative Signal Processing for Distributed Classification in Sensor Networks," The 2nd International Workshop on Information Processing in Sensor Networks (IPSN 2003), April 22-23, 2003, Palo Alto, CA.
- [8] P. Hall, Introduction to the Theory of Coverage Processes. John Wiley & Sons Inc., New York, 1998
- [9] D. Li, K. Wong, Y.H. Hu, A. Sayeed. "Detection, Classification and Tracking of Targets in Distributed Sensor Networks", IEEE Signal Processing Magazine, Volume: 19 Issue: 2, Mar 2002.
- [10] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, "Coverage Problems in Wireless Ad-Hoc Sensor Networks." INFOCOM'01, Vol 3, pp. 1380-1387, April 2001.
- [11] S. Meguerdichian, F. Koushanfar, G. Qu, and M. Potkonjak, "Exposure in Wireless Ad Hoc Sensor Networks." Procs. of 7th Annual International Conference on Mobile Computing and Networking (MobiCom'01), pp. 139-150, July 2001.
- [12] S. Meguerdichian and M. Potkonjak. "Low Power 0/1 Coverage and Scheduling Techniques in Sensor Networks." UCLA Technical Reports 030001. January 2003.
- [13] S. Patten, S. Poduri, and B. Krishnamachari, "Energy-Quality Tradeoffs for Target Tracking in Wireless Sensor Networks," The 2nd Workshop on Information Processing in Sensor Networks (IPSN 2003), April 2003.
- [14] D. Tian and N.D. Georganas, "A Coverage-preserved Node Scheduling scheme for Large Wireless Sensor Networks," Proceedings of First International Workshop on Wireless Sensor Networks and Applications (WSNA'02), Atlanta, USA, September 2002.
- [15] P. Varshney. Distributed Detection and Data Fusion. Springer-Verlag, New York, 1996.
- [16] Y. Xu, J. Heidemann, and D. Estrin, "Adaptive Energy-Conserving Routing for Multihop Ad Hoc Networks," Research Report 527, USC/Information Sciences Institute, October 2000.
- [17] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed Energy Conservation for Ad Hoc Routing," ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2001), Rome, Italy, July 16-21, 2001.
- [18] F. Ye, G. Zhong, S. Lu, and L. Zhang, "PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks". The 23rd International Conference on Distributed Computing Systems (ICDCS'03), May 2003.