Towards Optimal Sensor Placement with Hypercube Cutting Planes

(Invited Paper)

Siddhartha Chaudhuri
Dept. of Computer Science and Engg.
Indian Institute of Technology, Kanpur
sidc@cse.iitk.ac.in

Ratan K. Ghosh
Dept. of Computer Science and Engg.
Indian Institute of Technology, Kanpur
rkg@cse.iitk.ac.in

Sajal K. Das
Dept. of Computer Science and Engg.
University of Texas at Arlington
das@cse.uta.edu

***Abstract*— We study the problem of locating sensors to detect the failure of any set of radiation sources in a system. For computing illumination data, we suggest the use of radiosity methods. We then consider the problem of optimising sensor placement to unambiguously identify any inactive sources. We show that the problem can be transformed from a numerical to a geometrical domain, relate it to set covering, and then attempt to transform it into the domain of graphs. We present some results on hypercube cutting planes that help us progress towards the latter transformation by characterizing its combinatorial structure. Also, we outline an approach to estimate the size of the input space.**

I. Introduction

Optimal sensor placement is a challenging problem, not least because optimality is variously defined in this context. For instance, we may say a layout is optimal if it covers every part of a region with the fewest possible sensors. In computational geometry, this aspect has been extensively studied as the Art Gallery Problem [1]. Alternatively, a layout may be considered optimal if the sensor outputs are highly sensitive to changes in specific parts of the environment.

We will consider optimality in the light of *error detection*. Briefly, we are interested in finding which elements of a set of sources have failed. We assume we cannot examine the sources directly, so we must resort to observing the readings from a set of strategically placed sensors. We also require that we must be able to unambiguously distinguish between different sets of failing sources.

We will show that our model has strong links with studies of linear separability of point sets, common in neural network literature [9]. In particular, our work examines, among other things, the conditions under which a set of m-bit binary strings (which are represented as vertices of the unit hypercube in m-dimensional space) may be separated from all other m-bit binary strings by a hyperplane. Probabilistic estimates of the linear separability of a set of points in general position in space have been derived in [4]. Later work addressed the case when the points are not in general position, specifically when the input consists of the vertices of a hypercube [2]. We try to formulate graph-theoretic conditions for a linear separation of binary strings: this has the advantage of highlighting the combinatorial aspect of our sensor placement problem over the geometric one. Although we do not yet have the general conditions, we provide a subsidiary result for vertex pairs which we hope to be able to extend to the general case.

Our final goal is to prove or disprove NP-completeness of the problem, and use its combinatorial structure to design an efficient algorithm that gives optimal or near-optimal sensor layouts. In this paper, we briefly describe the problem, suggest a method for the generation of input data, examine the problem geometrically, establish a link with set covering, obtain some simple bounds on the size of the input space, present results that help us progress towards a graph formulation and finally outline an approach to refine our estimate of the size of the input space.

II. The Problem

We will consider a region with m radiation sources. We are given a set of sensors that respond to the amount of incident radiation as follows:

Each sensor has k *threshold levels* τ_1, τ_2,... τ_k, each greater than 0, and $k+1$ distinct *output values* v_1, v_2,... v_{k+1}. Let the radiant energy incident on the sensor be e. The sensor output is defined as follows:

$$out(e) = \begin{cases} v_1 & \text{if } e < \tau_1 \\ v_i & \text{if } \tau_{i-1} \leq e < \tau_i,\ 2 \leq i \leq k \\ v_{k+1} & \text{if } e \geq \tau_k \end{cases}$$

Some of the sources may have been disabled or are malfunctioning, and hence do not contribute to the overall illumination. We assume that an active source always radiates with the same (non-zero) strength and the same directionality, and an inactive source has zero strength. What is the minimum number of sensors that can always tell us exactly which of the m sources are inactive, and how do we place them?

We assume that the radiation incident on a sensor is the linear sum of contributions from all sources. Specifically, if $e_{ij} \geq 0$ is the energy incident on the ith sensor due to the jth source (when it is active), then e_i, the total energy incident on the ith sensor, can be written as

$$e_i = \mathbf{E}_i.\mathbf{X}$$

where $\mathbf{E}_i = [e_{i1}\ e_{i2} \ldots e_{im}]$ and the *source vector* $\mathbf{X} =$

$[X_1\ X_2 \ldots X_m]$, where

$$X_i = \begin{cases} 1 & \text{if source } i \text{ is active} \\ 0 & \text{if source } i \text{ is inactive} \end{cases}$$

We observe that our sensor model is essentially identical to the perceptron model of [9], with binary inputs, weights equal to the contribution of each source to the incident energy, and a thresholded output function out.

III. Data Generation with Radiosity

The first step in attacking the problem is to gather illumination data from different parts of the environment. As we have seen above, we need to know the intensity distribution due to each source separately. This can be done by switching on the sources one at a time and measuring each single-source illumination pattern. Such an approach may be tedious and impractical. A reasonable alternative is to use a radiosity method [6] to simulate the illumination model.

The surfaces in the scene are divided into a number of patches, each patch small enough to be considered homogenous. A standard radiosity computation will give the incident radiation at each patch. Let us perform m such computations, keeping exactly one (a different one) of the m sources active each time. This will give the illumination at each patch due to each source individually.

At a patch P, let the incident energy due to the jth source be $\varepsilon_j^P \geq 0$ when the source is active. Let us select n patches P_1, P_2,... P_n for placing n sensors. Then, for the ith sensor, we have

$$e_{ij} = \varepsilon_j^{P_i}, \quad 1 \leq j \leq m$$

We recall that e_i is the total energy incident on the ith sensor. We can write the *sensor input vector* $\mathbf{S}_{in} \equiv [e_1\ e_2 \ldots e_n]$ as:

$$\mathbf{S}_{in} = \mathbf{E}\mathbf{X}^T$$

where

$$\mathbf{E} = [e_{ij}]_{n \times m}$$

The *sensor output vector* $\mathbf{S}_{out}(\mathbf{S}_{in})$ is defined as $[out(e_1)\ out(e_2) \ldots out(e_n)]$. $\mathbf{S}_{out}$ is the observable quantity in our system. Our task is to choose sensor locations such that $\mathbf{S}_{out}$ is unique for each possible $\mathbf{X}$, so that we have a bijective mapping from sensor outputs to sets of inactive sources.

IV. The Geometrical Picture

The problem may be expressed geometrically. Let the input to the ith sensor be e_i for source vector $\mathbf{X}$ and e'_i for source vector $\mathbf{X}'$. We say the sensor *distinguishes* between $\mathbf{X}$ and $\mathbf{X}'$ if $out(e_i) \neq out(e'_i)$, i.e. if there is some threshold level between e_i and e'_i.

Let us now consider the m-dimensional space $\mathbb{R}^m$. We represent a point in this space as $\mathbf{x} = [x_1\ x_2 \ldots x_m]$. Each source vector $\mathbf{X}$ represents one of the 2^m vertices of the unit hypercube Q_m in this space. Sensor i distinguishes between $\mathbf{X}$ and $\mathbf{X}'$ if and only if there is some threshold level τ_h such that $\mathbf{X}$ and $\mathbf{X}'$ lie on opposite sides of the plane

$$\mathbf{E}_j.\mathbf{x} = \tau_h$$

For each sensor, there are k parallel planes, one for each threshold value. For a set of n sensors, there are a total of nk planes. If the sensor output vector is unique for each source vector, then there must be at least one sensor that distinguishes between each possible pair of source vectors. In other words, the linear subdivision of $\mathbb{R}^m$ induced by these planes must contain each vertex of Q_m in an unique cell (maximal connected m-dimensional region). We will require that no plane contains a vertex.

If we connect each pair of distinct vertices of Q_m, we obtain Q_m^{clique}, the *hypercube clique*, with $\binom{2^m}{2} = O(2^{2m})$ edges. To distinguish between all possible source vectors, we must place sensors at selected patches so that the resulting set of planes intersects the interior of every edge of this graph.

At this point, let us define the following two concepts:

Definition 1: A plane $\mathbf{A}.\mathbf{x} = b$ is "non-negative" if all the coordinates of $\mathbf{A}$ are non-negative and b is strictly positive. We observe that each plane generated by illumination data for a sensor is non-negative.

Definition 2: A set of edges of Q_m^{clique} is "valid" if there is some non-negative plane that intersects the interior of each edge in the set. Each valid set directly corresponds to an unique linear separation of the vertices of the hypercube.

V. Set Covering

We note that the problem is essentially a set covering problem. With each patch we may associate a set of k planes and hence a set of intersected edges of Q_m^{clique}. We must select patches so that the entire set of edges in Q_m^{clique} is covered.

Unconstrained set covering is known to be an NP-complete problem [8]. A simple greedy algorithm provides a solution within a factor $\alpha(\eta)$ of the optimum, where η is the size of the ground set and $\alpha(\eta) \equiv \ln \eta - \ln \ln \eta + \Theta(1)$, with the last term in $[-0.31, 0.78]$ [10]. Using the greedy algorithm, we can easily obtain a fair approximation to the optimal sensor locations in polynomial time.

However, it is not established whether our particular problem is NP-complete or not, since it is a constrained version of set covering. Certain subsets do not occur in the input: all the edges in a such a subset cannot be simultaneously intersected by a set of k parallel non-negative planes. For example, in the 3-dimensional cube, it is easy to check that no single non-negative plane ($k = 1$) can intersect both the edges $([0\ 0\ 0], [1\ 0\ 0])$ and $([1\ 1\ 0], [1\ 1\ 1])$ (see Fig. 1).

Let N_k b the number of subsets that *can* occur. To prove (or disprove) NP-completeness, it would be helpful to have some idea of the relative size of the input space, i.e. the ratio $N_k/|\Phi|$, where $\Phi = 2^{EDGES(Q_m^{clique})}$ is the power set of the edges of the hypercube clique. We note that $|\Phi| = 2^{|EDGES(Q_m^{clique})|} = 2^{\binom{2^m}{2}} = 2^{(4^m - 2^m)/2}$. Also, N_1 is simply the number of different linear separations of the vertices of the hypercube.

VI. Some Simple Bounds

Theorem 1: A plane must intersect at least $2^m - 1$ edges of Q_m^{clique}, if it intersects any edge and does not contain a vertex, and at most 2^{2m-2} edges.

Proof: Let plane $\mathbf{A}.\mathbf{x} = b$ intersect at least one edge of Q_m^{clique}. Of the 2^m vertices, say p vertices are on the positive side of the plane ($\mathbf{A}.\mathbf{x} > b$) and $q = 2^m - p$ vertices on the negative side. A pair of vertices on opposite sides of the plane corresponds to an intersected edge. This is a bijective mapping, since no intersected edge can join two vertices on the same side of the plane. The graph defined by the vertices and intersected edges of Q_m^{clique} is thus isomorphic to the complete bipartite graph $K_{p,q}$ (Fig. 1). The number of edges in such a graph is simply $pq = p(2^m - p)$. The minimum of this expression is $2^m - 1$, and it is obtained when p or q is 1. The maximum is 2^{2m-2}, and it is obtained when $p = q = 2^{m-1}$. ■

Theorem 2: There is at least one non-negative plane that intersects exactly $2^m - 1$ edges of Q_m^{clique} and at least one non-negative plane that intersects exactly 2^{2m-2} edges.

Proof: To prove the first part, we must show that some non-negative plane $\mathbf{A}.\mathbf{x} = b$ can isolate a single vertex. Choose any positive b, and make each coordinate of $\mathbf{A}$ greater than b. The resulting plane has the origin on its negative side and all other vertices of Q_m^{clique} on its positive side. The number of intersected edges is $2^m - 1$.

To prove the second part, we must construct a non-negative plane that has half the vertices on its positive side and half on its negative side. Let us make the first coordinate of $\mathbf{A}$ greater than b and all other coordinates 0. Since exactly half the vertices of the unit hypercube have their first coordinate 1 and the rest have 0, this plane evenly splits the vertices. The number of intersected edges is 2^{2m-2}. (We note that the edges and vertices on each side of this plane define graphs isomorphic to Q_{m-1}^{clique}.) ■

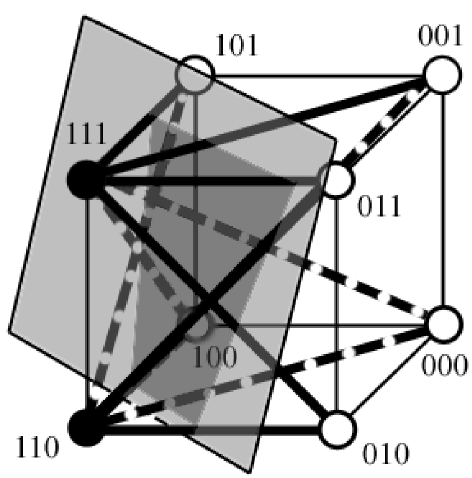

Fig. 1. Bipartite graph $K_{2,6}$ formed by intersection of a plane (grey) with the 3-cube clique. The black vertices are on one side of the plane and the white vertices on the other side. Thick lines denote intersected edges (some non-intersected edges have been omitted for clarity).

These results show that the size of each valid set is tightly bounded above and below by 2^{2m-2} and $2^m - 1$ respectively.

Also, the number of different possible cuts of Q_m^{clique} is bounded by 2^{2^m}, the number of ways of separating the vertices into two sets. Each valid set corresponds to exactly one cut. Therefore the number of valid sets is less than 2^{2^m}, and

$$\frac{N_1}{|\Phi|} \;<\; \frac{2^{2^m}}{2^{\binom{2^m}{2}}} \;=\; 2^{(3\times 2^m - 4^m)/2} \;\rightarrow\; 0$$

We note in passing that the maximum number of cells in a linear subdivision induced by H hyperplanes in $\mathbb{R}^m$ [7] is

$$\sum_{i=0}^{m} \binom{H}{m-i}$$

This tells us that for isolating each of 2^m points, we need at least m hyperplanes. A set of m hyperplanes aligned with the coordinate planes (and with suitable intercepts) is an obvious example that does the job. Interpreting these results for our sensor layout, we find that at least $\lceil m/k \rceil$ sensors are required, and if we are lucky we may be able to make do with just this many.

VII. Towards a Pure Graph Problem

We would like to remove the geometrical component and transform the problem into the graph domain, so that we can try to use the large body of results in graph theory that have been associated with studies of NP-completeness and optimisation. As a first step, we shall construct a directed graph that has some properties equivalent to those related to the intersection of the hypercube clique with a plane.

We will use boldface to denote the position vector of a point, i.e. $\mathbf{p}$ is the position vector of p.

Consider a vertex u of Q_m. Its position vector $[u_1\ u_2 \ldots u_m]$ is a bit vector, i.e. it contains only 0's and 1's as elements. Let us define $ONES(u)$ as the index set of the 1's in $\mathbf{u}$, that is,

$$ONES(u) = \{i \mid u_i = 1\}$$

We observe that for any plane $\mathbf{A}.\mathbf{x} = b$, where $\mathbf{A} = [a_1\ a_2 \ldots a_m]$,

$$\mathbf{A}.\mathbf{u} = \sum_{i=1}^{m} a_i u_i = \sum_{i \in ONES(u)} a_i$$

Consider two vertices u and v such that $ONES(u) \subset ONES(v)$, where $\subset$ denotes the *proper subset* relation. Then for any non-negative plane $\mathbf{A}.\mathbf{x} = b$ we have $\mathbf{A}.\mathbf{u} \leq \mathbf{A}.\mathbf{v}$. Consider the *directed hypercube graph* $\vec{Q}_m$ on the vertices of Q_m. Its edges coincide with the edges of the hypercube (those that join vertices differing in exactly one coordinate) and are directed from $[0\ 0 \ldots 0]$ to $[1\ 1 \ldots 1]$, i.e. edge (u, v) is present iff $\mathbf{v}$ has a 1 where $\mathbf{u}$ has a 0 and they agree in all other coordinates.

Lemma 1: There is a path of non-zero length from u to v in $\vec{Q}_m$ if and only if $ONES(u) \subset ONES(v)$.

Proof: **If part**: Successively change each 0 in $\mathbf{u}$ to 1, if $\mathbf{v}$ has a 1 in that coordinate. After a finite number of steps, we will obtain $\mathbf{v}$, since $ONES(u) \subset ONES(v)$. Each step corresponds to an edge of $\vec{Q}_m$, since the initial and final values differ by exactly one 1. We see by induction that there is a path $u \rightsquigarrow v$ in $\vec{Q}_m$. Also, since $u \neq v$, we must change a 0 to a 1 at least once, so the path has non-zero length.

Only-if part: Consider any edge (u', v') on the path $u \rightsquigarrow v$ (there must be at least one such edge since the path has non-zero length). By the construction of $\vec{Q}_m$, we know that $\mathbf{v}'$ is the same as $\mathbf{u}'$ but for an extra 1. So $ONES(u') \subset ONES(v')$. Further, $\subset$ is a transitive relation. Applying this inductively to the vertices in the path, starting from u, we obtain $ONES(u) \subset ONES(v)$. ■

Since the subset relationship and path existence in $\vec{Q}_m$ have been shown to be equivalent, we will introduce a common notation for them. We say that $u \to v$ if the following equivalent statements hold:

1) $ONES(u) \subset ONES(v)$ ($\subset$ denotes *proper subset*).
2) There is a path of non-zero length from u to v in $\vec{Q}_m$.

If $u \to v$, then for any non-negative plane $\mathbf{A}.\mathbf{x} = b$, $\mathbf{A}.\mathbf{u} \leq \mathbf{A}.\mathbf{v}$.

We will now present a result about *valid pairs*, i.e. valid sets of 2 edges.

Theorem 3: There is a non-negative plane that intersects the distinct edges (u_1, v_1) and (u_2, v_2) in Q_m^{clique} if and only if both u_1 and v_1 do not lie on paths from the origin $\mathbf{0} = [0\ 0 \ldots 0]$ to u_2 and v_2 in $\vec{Q}_m$, and vice versa.

In other words, there is a non-negative plane that intersects (u_1, v_1) and (u_2, v_2) in Q_m^{clique} iff the following do not hold simultaneously (when u_1, v_1, u_2 and v_2 are all distinct):

$u_1 \to u_2$, $v_1 \to u_2$, $u_1 \to v_2$, $v_1 \to v_2$

and also, the following do not hold simultaneously:

$u_2 \to u_1$, $v_2 \to u_1$, $u_2 \to v_1$, $v_2 \to v_1$

These "forbidden configurations" may be expressed graphically as

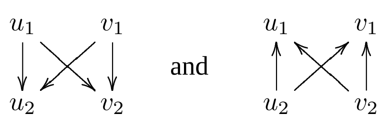

If the edges share a common endpoint, say $u_1 = u_2 = u$, then the forbidden configurations reduce to

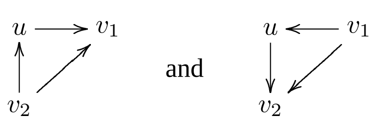

Proof: **Only-if part (by contradiction)**: There is some non-negative plane $\mathbf{A}.\mathbf{x} = b$ that intersects (u_1, v_1) and (u_2, v_2). Let us suppose, without loss of generality, that both u_1 and v_1 lie on paths from the origin to u_2 and v_2 in $\vec{Q}_m$, i.e., $u_1 \to u_2$, $v_1 \to u_2$, $u_1 \to v_2$ and $v_1 \to v_2$ hold simultaneously. This implies that:

$$\begin{aligned} \mathbf{A}.\mathbf{u}_1 &\leq \mathbf{A}.\mathbf{u}_2 \\ \mathbf{A}.\mathbf{v}_1 &\leq \mathbf{A}.\mathbf{u}_2 \\ \mathbf{A}.\mathbf{u}_1 &\leq \mathbf{A}.\mathbf{v}_2 \\ \mathbf{A}.\mathbf{v}_1 &\leq \mathbf{A}.\mathbf{v}_2 \end{aligned}$$

Now if the plane intersects (u_2, v_2), we must have

$$\begin{aligned} &\mathbf{A}.\mathbf{u}_2 < b < \mathbf{A}.\mathbf{v}_2, \quad \text{or} \\ &\mathbf{A}.\mathbf{v}_2 < b < \mathbf{A}.\mathbf{u}_2 \end{aligned}$$

In the first case, we have $\mathbf{A}.\mathbf{u}_1 \leq \mathbf{A}.\mathbf{u}_2 < b$ and $\mathbf{A}.\mathbf{v}_1 \leq \mathbf{A}.\mathbf{u}_2 < b$, so the plane cannot intersect (u_1, v_1). Similarly in the second case, we have $\mathbf{A}.\mathbf{u}_1 \leq \mathbf{A}.\mathbf{v}_2 < b$ and $\mathbf{A}.\mathbf{v}_1 \leq \mathbf{A}.\mathbf{v}_2 < b$, so the plane cannot intersect (u_1, v_1).

This is a contradiction, so our supposition was incorrect: u_1 and v_1 cannot both lie on paths from the origin to u_2 and v_2 in $\vec{Q}_m$. By an identical argument, with subscripts 1 and 2 interchanged, we can show that u_2 and v_2 cannot both lie on paths from the origin to u_1 and v_1.

If part: We know that the configuration is allowed. We will examine each possible configuration of (u_1, v_1) and (u_2, v_2) that does not contain any set of forbidden relationships.

The base configuration space is very large: each pair (p, q) from the set $\{u_1, v_1, u_2, v_2\}$ may be related as $p \to q$, or $q \to p$, or not related at all, a total of 3 possibilities. There are $\binom{4}{2} = 6$ possible pairs, so the total number of possible configurations is $3^6 = 729$. Fortunately, three factors drastically reduce the number of configurations we must examine. These are:

1) There are many forbidden configurations.
2) The subset relationship is transitive, so configurations such as the following are equivalent:

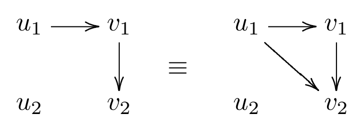

Also, configurations that contain a cycle do not occur in our graph: we cannot have $p \to q \to r \to p$ because this would imply that $p \to p$, which is impossible since a set cannot be a proper subset of itself.
3) Symmetries may be exploited. We can (a) swap the endpoints of either edge or (b) exchange the two edges, without essentially altering the configuration and the associated arguments.

We will enumerate all possible unique allowed configurations (by unique we mean that we will not consider equivalent configurations separately) and show that in each case, we can construct a non-negative plane $\mathbf{A}.\mathbf{x} = b$, where $\mathbf{A} = [a_1\ a_2 \ldots a_m]$, that intersects both the edges.

A. u_1, u_2, v_1 and v_2 are distinct

We may divide the possible configurations into 7 cases, as presented below. For lack of space, we will not prove each of these cases in this paper. The complete proof may be found in [3].

Case 1: Configurations with the following subgraph or its symmetrical equivalents:

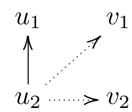

where a dotted line indicates *absence* of the corresponding arrow. If this subgraph is present, then for each $s \in ONES(u_2)$, set $a_s = (b+\epsilon)/|ONES(u_2)|$, $\epsilon > 0$. Set all other coefficients to 0. This gives $\mathbf{A}.\mathbf{u}_1 > b$ and $\mathbf{A}.\mathbf{u}_2 > b$. Since neither $ONES(v_1)$ or $ONES(v_2)$ contains all the elements of $ONES(u_2)$, ϵ can be made small enough so that $\mathbf{A}.\mathbf{v}_1 < b$ and $\mathbf{A}.\mathbf{v}_2 < b$. Hence the plane intersects (u_1, v_1) and (u_2, v_2).

Some of the configurations in this class are:

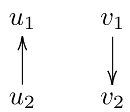 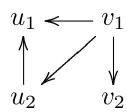 and 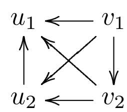

Case 2: Configurations with the following subgraph or its symmetrical equivalents:

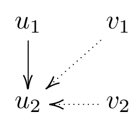

where a dotted line again indicates absence of the corresponding arrow. If this subgraph is present, $ONES(v_1)$ must have some element s not in $ONES(u_2)$, hence not in $ONES(u_1)$ either, and $ONES(v_2)$ must have a similar element t (s and t need not be distinct). Set a_s, $a_t > b$, and all other coefficients to 0. This gives $\mathbf{A}.\mathbf{u}_1 = 0 < b < \mathbf{A}.\mathbf{v}_1 = a_s$ and $\mathbf{A}.\mathbf{u}_2 = 0 < b < \mathbf{A}.\mathbf{v}_2 = a_t$. Hence the plane intersects (u_1, v_1) and (u_2, v_2).

Some of the configurations in this class are:

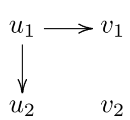 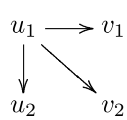 and 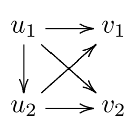

Cases 3 - 7: Individual configurations not covered by cases 1 and 2 above:

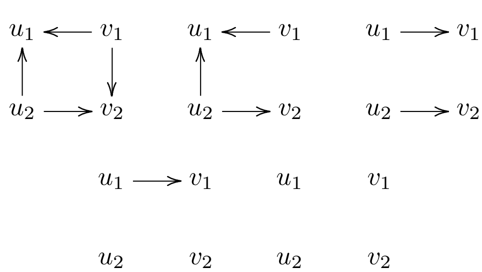

For brevity, let us only prove Case 7, with the configuration

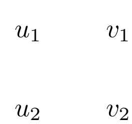

We observe that none of u_1, v_1, u_2, v_2 can be the origin 0, since $0 \to p$ for any other vertex p of $\vec{Q}^m$. Therefore $|ONES(p)| > 0$ for each $p \in \{u_1, v_1, u_2, v_2\}$.

Let us consider the case when $ONES(u_1)$ has an element t not in $ONES(v_1)$ or $ONES(v_2)$. We set $a_t > b$. Also, we set $a_s = (b+\epsilon)/|ONES(u_2)|$ for each $s \in ONES(u_2)$, $s \neq t$, $\epsilon > 0$. All other a_i's are set to 0. This gives us $\mathbf{A}.\mathbf{u}_1 > b$ and $\mathbf{A}.\mathbf{u}_2 > b$.

$ONES(v_1)$ does not have t, nor does it have all the elements of $ONES(u_2)$ (else $u_2 \to v_1$). Similarly for $ONES(v_2)$. Therefore for small enough ϵ, $\mathbf{A}.\mathbf{v}_1 < b$ and $\mathbf{A}.\mathbf{v}_2 < b$.

$\therefore$ This plane intersects (u_1, v_1) and (u_2, v_2).

In general, we can construct a plane in similar fashion whenever the endpoint of an edge has a 1 where the other endpoint of the same edge and at least one endpoint of the other edge have 0's.

If there is no such 1, then the following sets are empty:

$$\begin{array}{l}
ONES(u_1) \cap \overline{ONES(v_1)} \cap \overline{ONES(u_2)} \\
ONES(u_1) \cap \overline{ONES(v_1)} \cap \overline{ONES(v_2)} \\
ONES(u_2) \cap \overline{ONES(v_2)} \cap \overline{ONES(u_1)} \\
ONES(u_2) \cap \overline{ONES(v_2)} \cap \overline{ONES(v_1)} \\
ONES(v_1) \cap \overline{ONES(u_1)} \cap \overline{ONES(u_2)} \\
ONES(v_1) \cap \overline{ONES(u_1)} \cap \overline{ONES(v_2)} \\
ONES(v_2) \cap \overline{ONES(u_2)} \cap \overline{ONES(u_1)} \\
ONES(v_2) \cap \overline{ONES(u_2)} \cap \overline{ONES(v_1)}
\end{array}$$

These correspond to the greyed-out regions in the Venn diagram in Fig. 2.

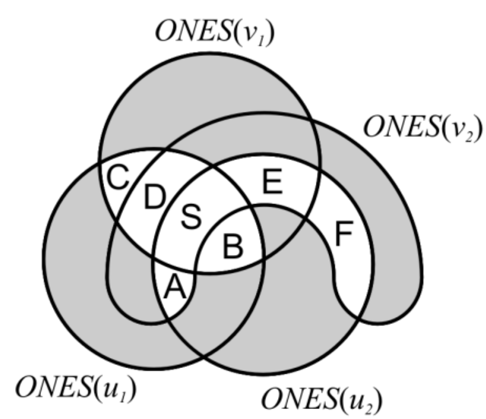

Fig. 2. Venn Diagram for Case 7 (all vertices distinct)

So we have

$$\begin{aligned}
ONES(u_1) &= S \cup A \cup B \cup C \cup D \\
ONES(u_2) &= S \cup A \cup B \cup E \cup F \\
ONES(v_1) &= S \cup B \cup C \cup D \cup E \\
ONES(v_2) &= S \cup A \cup D \cup E \cup F
\end{aligned}$$

where S, A, B, C, D, E and F, as shown in Fig. 2, are all disjoint.

A and E cannot be empty, since then $ONES(u_1) \subseteq ONES(v_1)$ or vice versa. Similarly B and D cannot be empty.

For each $s \in A \cup B$, we set $a_s = (b + \epsilon)/|A \cup B|$, $\epsilon > 0$. We also set all other a_i's to 0. This gives $\mathbf{A}.\mathbf{u}_1 > b$ and $\mathbf{A}.\mathbf{u}_2 > b$, since both $ONES(u_1)$ and $ONES(u_2)$ are supersets of $A \cup B$. Also, $ONES(v_1)$ does not contain A and $ONES(v_2)$ does not contain B. Since A and B are disjoint and non-empty, $|A| < |A \cup B|$ and $|B| < |A \cup B|$. So for small enough ϵ, $\mathbf{A}.\mathbf{v}_1 < b$ and $\mathbf{A}.\mathbf{v}_2 < b$.

$\therefore$ This plane intersects (u_1, v_1) and (u_2, v_2).

It is simple, though tedious, to check that the cases listed above cover all valid configurations when u_1, u_2, v_1 and v_2 are distinct.

B. The edges share a common endpoint, say $u_1 = u_2 = u$

Again, we may divide our work into cases. These are:

Case 1: Configurations with the following subgraph or its symmetrical equivalents:

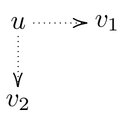

Case 2: Configurations with the following subgraph or its symmetrical equivalents:

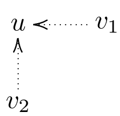

As usual, a dotted line indicates absence of the corresponding arrow.

Again, for lack of space, we will not give the proofs for the above cases in this paper. Suffice it to say that the proofs correspond very closely to the proofs of Cases 1 and 2 when all the vertices are distinct. Details are available in [3]. ■

VIII. Generalization

If we could extend Theorem 3 to sets of edges of any size, we would have a purely combinatorial way to test if a hyperplane can perform a given linear separation. This would completely characterize valid sets in graph-theoretic terms, and a combinatorial enumeration could be possible.

In the absence of such a result at present, we outline a different approach to estimate N_1, the number of valid sets. We observe that the partition induced by a non-negative plane on the vertices of a hypercube has a pleasing structure. Firstly, as we have mentioned in Sec. VI, the intersected edges define a complete bipartite subgraph of Q_m^{clique}. Secondly, from Sec. VII, if a vertex u lies on the positive side of the plane, then all vertices v such that $u \to v$ also lie on the positive side (a similar result holds for the negative side).

Let us consider the *basis set* $\mathcal{B}^+$, which comprises all the vertices v on the positive side for which there are no other vertices u also on the positive side such that $u \to v$. It is easy to check that $\mathcal{B}^+$ uniquely defines the complete set of vertices on the positive side. A similar basis set $\mathcal{B}^-$ may be obtained for the negative side. Further, $\mathcal{B}^+$ and $\mathcal{B}^-$ are complementary, so specifying either one of them completely and uniquely defines the partition, and hence the corresponding valid set of edges. Our comments on $\mathcal{B}^+$ in the next few paragraphs apply equivalently to $\mathcal{B}^-$.

Let Γ^+ be the set of all possible basis sets $\mathcal{B}^+$. We define Λ as the set of all subsets V of the vertices of Q_m such that no $u, v \in V$ have $u \to v$, i.e. it is not possible to go from u to v via the edges of $\vec{Q}_m$. Evidently, $\Gamma^+ \subseteq \Lambda$. So

$$N_1 = |\Gamma^+| \leq |\Lambda|$$

We may write an expression for $|\Lambda|$ using the inclusion-exclusion principle and attempt to bound the sum. This is a work under progress, so we will only state the straightforward result that the number of non-zero length paths in $\vec{Q}_m$ is $3^m - 2^m$ (see [3] for a derivation).

IX. Conclusion

We have described a problem in sensor placement and outlined an approach towards its analysis and solution. The results obtained in this paper are a first step towards characterizing the problem combinatorially. We have shown that the problem is a highly constrained version of set covering, with a considerably smaller input space. We are working on establishing whether the problem is NP-complete or not, using the results in this paper and those of other authors. Our immediate goal is a tight estimate of the exact size of the input space. With a generalized version of Theorem 3, we envisage an optimized covering algorithm for sensor placement that yields better results in less time.

References

[1] M. de Berg, M. van Kreveld, M. Overmars and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*, 2nd rev. ed., Springer-Verlag, Berlin, 2000.

[2] M. Budinich, "On linear separability of random subsets of hypercube vertices", *J. Phys. A: Math. Gen.* 24, pp. L211-L213, 1991.

[3] S. Chaudhuri and R. K. Ghosh, "Optimal sensor placement", Technical Report, IIT Kanpur, 2004.

[4] T. M. Cover, "Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition", *IEEE Trans. Electron. Comput.*, EC-14, pp. 326-34, 1965.

[5] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, New York, 1979.

[6] C. M. Goral, K. E. Torrance, D. P. Greenberg and B. Battaile, "Modeling the interaction of light between diffuse surfaces", *Computer Graphics (SIGGRAPH 84 Proceedings)*, 18(3):213222, July 1984.

[7] D. Halperin, "Arrangements", *Handbook of Discrete and Computational Geometry*, ed. J. E. Goodman and J. O'Rourke, ch. 21, CRC Press LLC, Boca Raton, FL, 1997.

[8] R. M. Karp, "Reducibility among combinatorial problems", *Complexity of Computer Computations*, ed. R. Miller and J. Thatcher, pp. 85-103, Plenum Press, 1972.

[9] M. Minsky and S. Papert, *Perceptrons*, MIT Press, Cambridge, MA, 1969.

[10] P. Slavík, "A tight analysis of the greedy algorithm for set cover", Technical Report, SUNY, Buffalo, 1995. URL: `http://www.cse.buffalo.edu/tech-reports/95-54.ps`.