APPROXIMATING THE MINIMAL SENSOR SELECTION FOR SUPERVISORY CONTROL

Samir Khuller * **Guy Kortsarz** ** **Kurt R. Rohloff** ***,[1]

* *Department of Computer Science, University of Maryland College Park, MD 20742, USA samir@cs.umd.edu, www.cs.umd.edu/users/samir*
** *Computer Science Department, Business and Science Building, Rutgers University, Camden, NJ, USA guyk@camden.rutgers.edu*
*** *Department of EECS, The University of Michigan 1301 Beal Avenue, Ann Arbor, MI 48109-2122, USA krohloff@eecs.umich.edu, www.eecs.umich.edu/umdes*

Abstract: This paper discusses the problem of selecting a set of sensors of minimum cardinality that can be used for the synthesis of a supervisory controller. We show how this sensor selection problem is related to a type of directed graph *st*-cut problem that has not been previously discussed in the literature. Approximation algorithms to solve the sensor selection problem can be used to solve the graph cut problem and vice-versa. Polynomial time algorithms to find good approximate solutions to either problem most likely do not exist (under certain complexity assumptions), but a time efficient approximation algorithm is shown that solves a special case of these problems.

Keywords: Discrete-Event Systems, Supervisory Control, Graph Theory, Approximation Algorithms, Observers, Computer-Aided Control System Design.

1. INTRODUCTION

When a controller operates on a system to match a specification, the controller may not need information from all sensors at its disposal. Therefore, for reasons of economy or simplicity, the control designer may desire the controller to use as few sensors as possible. There may be many sensor selections sufficient for the controller to match the specification, and the optimal selection may not be obvious. Unfortunately, the problem of finding this smallest cardinality sensor selection set is NP-complete (Yoo and Lafortune (2002)). Effective polynomial time approximation algorithms exist for many real-world NP-complete optimization problems (Ausiello et al. (1999)), and an approximation of this minimal solution may be acceptable. We explore the problem of approximating solutions to the sensor minimization problem.

In the next section the problem statement is formulated in the supervisory control framework of Lin and Wonham (1988) and necessary background information from computer science is given. We relate the sensor selection problem to a

[1] This paper reports on work commenced while the third author was under the direction of Prof. Jan H. van Schuppen at CWI in Amsterdam. Financial support in part for the investigation was made available by the European Commission through the project Control and Computation (IST-2001-33520) of the Information Society Technologies Program and by NSF grant CCR-0082784.

type of directed graph st-cut problem in Section 3 and we show some inapproximability results for our sensor selection and graph cutting problems in Section 4. We show a polynomial time approximation algorithm for a special case of the sensor selection and graph cutting problems in Section 5 and close with a brief discussion in Section 6.

2. THE SENSOR SELECTION PROBLEM AND APPROXIMATION ALGORITHMS

In our framework systems and specifications are modelled as the automata $G = (X^G, x_0^G, \Sigma, \delta^G)$ and $H = (X^H, x_0^H, \Sigma, \delta^H)$, respectively. The behavior generated by G is denoted by $\mathcal{L}(G)$ and the behavior generated by the controller S controlling G is denoted by $\mathcal{L}(S/G)$. The system S/G is said to match the specification H if $\mathcal{L}(S/G) = \mathcal{L}(H)$.

Controllable events ($\Sigma_c \subseteq \Sigma$) and observable events ($\Sigma_o \subseteq \Sigma$) are those events that can be respectively disabled or observed by the controller. Due to the controllability and observability theorem, Lin and Wonham (1988), there exists a controller S such that $\mathcal{L}(S/G) = \mathcal{L}(H)$ if and only if $\mathcal{L}(H)$ is controllable with respect to $\mathcal{L}(G)$ and $\Sigma \setminus \Sigma_c$ and $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c. We assume that $\mathcal{L}(H)$ is always controllable. See Cassandras and Lafortune (1999) for a deeper introduction to supervisory control.

When designing a controller, the set of observable events can be chosen by selecting the appropriate respective sensors for those events. We call $\Sigma_o \subseteq \Sigma$ a *sufficient sensor selection* with respect to G, H and Σ_c if $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c. We call the problem of finding a minimal cardinality set of observable events the *sensor selection problem.*

Problem 1. Sensor Selection: Given G, H and $\Sigma_c \subseteq \Sigma$, find a sufficient sensor selection Σ_o^{min} such that for any other sufficient sensor selection Σ_o, $|\Sigma_o^{min}| \leq |\Sigma_o|$.

For a survey of literature relevant to Problem 1, please consult Rohloff (2004).

Despite the NP-completeness of Problem 1 (Yoo and Lafortune (2002)), we may still be required to find a sufficient sensor selection Σ_o whose cardinality is approximately $|\Sigma_o^{min}|$. Fortunately, some NP-complete minimization problems have fairly accurate polynomial time approximation algorithms. This means sufficient and approximate solutions can be found for many difficult problems. However, not all NP-complete minimization problems are believed to have this property. We further discuss these approximation properties for Problem 1.

To better quantify what we mean by an approximation to Problem 1, suppose P is the set of instances of Problem 1. Let $p \in P$ be a specific problem instance corresponding to G, H and Σ_c. Suppose $\Sigma_o^{min}(p)$ is the solution of this problem instance and A is an algorithm that when given input p, returns $\Sigma_o^A(p)$ such that $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, $\Sigma_o^A(p)$ and Σ_c. We measure the utility of the approximation $\Sigma_o^A(p)$ with the ratio $|\Sigma_o^A(p)|/|\Sigma_o^{min}(p)|$. Problem 1 has an *r-approximation* if $\forall p \in P, |\Sigma_o^A(p)|/|\Sigma_o^{min}(p)| \leq r$. This r-approximation notation also holds for other approximation problems. For a deeper discussion of these topics, see Ausiello et al. (1999).

Other necessary background from computer science is that $DTIME(f(n))$ defines the set of all problems that can be solved by deterministic algorithms with time complexity in $O(f(n))$. Using standard computer science notation, n is the size of the encoding of the problem instance. Of interest is the class $DTIME(n^{\text{polylog } n})$. It is believed that $NP \nsubseteq DTIME(n^{\text{polylog } n})$, but this has not been proved (Arora and Lund (1997)).

3. THE GRAPH CUTTING PROBLEM

We now show how the sensor selection problem is related to a special type of directed graph st-cut problem. Examples of the constructions in this section are given in Rohloff (2004).

Suppose we are given an edge-colored directed graph $D = (V, A, C)$ where V is a set of vertices, $A \subseteq V \times V$ are directed edges, $C = \{c_1, \ldots, c_p\}$ is the set of colors and for $s, t \in V$, there is a path of directed edges from s to t. Each edge is assigned a color in C and let A_i be the edges having color c_i. Given $I \subseteq C$, let $A_I = \cup_{c_i \in I} A_i$. I is a colored st-cut if $(V, A \setminus A_I, C)$ has no path from s to t. This prompts us to define the colored cut problem.

Problem 2. Minimal Colored Cut: Given an edge colored directed graph $D = (V, A, C)$ and two vertices, $s, t \in V$, find a colored st-cut $I^{min} \subseteq C$ such that for any other colored st-cut $I \subseteq C$, $|I^{min}| \leq |I|$.

We now convert an instance of a colored cut problem to an instance of a sensor selection problem. Suppose we are given the edge colored directed graph $D = (V, A, C)$ and two vertices s, t. We first construct a system G.

If $C = \{c_1, \ldots, c_p\}$, let there be a set of events $\{\sigma_1, \ldots, \sigma_p\}$ such that for every color c_i, there is a corresponding event σ_i. Let γ be another event and define $\Sigma = \{\sigma_1, \ldots, \sigma_p, \gamma\}$. Define $X^G = V \cup \{s', s'', t'\}$ where s', s'', t' are states not in V. Let $x_0^G = s$. To define the state transition function, let v_1, v_2 be any vertices except s. To define δ^G, if $(v_1, v_2) \in A_i$, then $\delta^G(v_1, \sigma_i) = v_2$. If $(s, v_2) \in A_i$, then $\delta^G(s, \sigma_i) = v_2$ and $\delta^G(s'', \sigma_i) = v_2$. If $(v_1, s) \in A_i$, then $\delta^G(v_1, \sigma_i) = s''$. For simplicity we assume that $(s, s) \notin A$. Also, define $\delta^G(s, \gamma) = s'$ and $\delta^G(t, \gamma) = t'$. Let H be a copy of G except that $\delta^H(t, \gamma)$ is undefined. Let $\Sigma_c = \{\gamma\}$.

In the system above, γ must be enabled at s and be disabled at t. There is a control conflict if there is a path in G from s to t where no event is observed. Therefore, as system behavior progresses, if any event is observed, then γ can be disabled. Hence, a set of colors $I = \{c_a, \ldots, c_z\}$ is a colored cut for D if and only if selecting sensors for $\{\sigma_a, \ldots, \sigma_z\}$ makes the system observable. Therefore any approximation algorithm for the sensor selection problem can also be used with the same absolute effectiveness for the colored cut problem.

We now show the converse by converting an instance of Problem 1 to an instance of Problem 2. We start with a modified $\mathcal{M}_{\Sigma_o}$-machine method for testing observability (Tsitsiklis (1989)). Suppose we are given $H = (X^H, x_0^H, \Sigma, \delta^H)$, $G = (X^G, x_0^G, \Sigma, \delta^G)$, Σ_o and Σ_c such that we wish to test if $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c. Let Σ' be a copy of the event set Σ where for every event $\sigma \in \Sigma$, there is a corresponding event $\sigma' \in \Sigma'$. We can now give the following definition: $\mathcal{M}_{\Sigma_o} = (X^{\mathcal{M}_{\Sigma_o}}, x_0^{\mathcal{M}_{\Sigma_o}}, \Sigma^{\mathcal{M}_{\Sigma_o}}, \delta^{\mathcal{M}_{\Sigma_o}})$ where $X^{\mathcal{M}_{\Sigma_o}} := X^H \times X^H \times X^G \cup \{d\}, x_0^{\mathcal{M}_{\Sigma_o}} := (x_0^H, x_0^H, x_0^G), \Sigma^{\mathcal{M}_{\Sigma_o}} := \Sigma \cup \Sigma'$. Let us define the set of conditions at state (x_1, x_2, x_3) that we call the $(*)$ conditions.

$$\left.\begin{array}{l} \delta^H(x_1, \sigma) \text{ is defined if } \sigma \in \Sigma_c \\ \delta^H(x_2, \sigma) \text{ is not defined} \\ \delta^G(x_3, \sigma) \text{ is defined} \end{array}\right\} \qquad (*)$$

The nondeterministic transition function $\delta^{\mathcal{M}_{\Sigma_o}}$ is defined as follows.
For $\sigma \notin \Sigma_o$ and its Σ' equivalent, σ',

$$\delta^{\mathcal{M}_{\Sigma_o}}((x_1, x_2, x_3), \sigma') = \left\{ \begin{array}{l} (\delta^H(x_1, \sigma), x_2, x_3) \\ (x_1, \delta^H(x_2, \sigma), \delta^G(x_3, \sigma)) \end{array} \right\}$$

For $\sigma \in \Sigma$,

$$\delta^{\mathcal{M}_{\Sigma_o}}((x_1, x_2, x_3), \sigma) = \left\{ \begin{array}{l} (\delta^H(x_1, \sigma), \delta^H(x_2, \sigma), \delta^G(x_3, \sigma)) \\ d \quad \text{if } (*) \end{array} \right\}$$

For $\sigma \in \Sigma$, $\delta^{\mathcal{M}_{\Sigma_o}}(d, \sigma)$ is undefined, and $\delta^{\mathcal{M}_{\Sigma_o}}$ can be extended in the usual manner to be defined over strings of events.

The state d is reachable in $\mathcal{M}_{\Sigma_o}$ if and only if $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c. The $\mathcal{M}_{\Sigma_o}$-machine here is modified from the original in Tsitsiklis (1989) in that Σ' transitions replace some Σ transitions.

Effectively $\mathcal{M}_{\Sigma_o}$ is a nondeterministic simulation of an observer's estimate of a system's behavior with respect to a specification. In a state (x_1, x_2, x_3) of $\mathcal{M}_{\Sigma_o}$, the first state is an observer's estimate of the specification state and the second and third states are the true states of the specification and system, respectively. In $\mathcal{M}_{\Sigma_o}$ a Σ transition occurs if an event occurrence in H and G is correctly predicted by the observer. A Σ' transition occurs if the prediction is not correct. Therefore, if an event is observed, it is predicted correctly and σ' transitions in $\mathcal{M}_{\Sigma_o}$ would be removed if σ were made observable. This implies that $\mathcal{M}_{\Sigma_o \cup \{\sigma\}}$ can be constructed from $\mathcal{M}_{\Sigma_o}$ by cutting all σ' transitions, and conversely, cutting all occurrences of σ' transitions in $\mathcal{M}_{\Sigma_o}$ corresponds to adding σ to Σ_o.

This realization shows that the sensor selection problem is really a type of colored cut problem. Suppose we consider $\mathcal{M}_\emptyset$ to be a colored directed graph as introduced above where the transition labels are really colors. A colored $x_0^{\mathcal{M}_\emptyset} d$-cut for $\mathcal{M}_\emptyset$ where we only select Σ' transitions for cutting corresponds to a sufficient sensor selection for observability to hold. This prompts the following proposition.

Proposition 3. $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c if and only if $\Sigma'_o \subseteq \Sigma'$ is a colored $x_0^{\mathcal{M}_\emptyset}$ *d*-cut $\mathcal{M}_\emptyset$.

However, the $\mathcal{M}_\emptyset$ cut problem is not in the same form as in Problem 2 because we are not allowed to cut Σ transitions from $\mathcal{M}_\emptyset$ in Proposition 3. To counter this difference, we use the following construction which performs a form of state condensation and hides the Σ transitions in $\mathcal{M}_\emptyset$.

Construct $\mathcal{M}_{\Sigma_o}$ from H, G, Σ_c and Σ_o. Define $X_x^{\mathcal{M}_{\Sigma_o}} = \delta^{\mathcal{M}_{\Sigma_o}}(x^{\mathcal{M}_{\Sigma_o}}, t)$ which represents all states that could be reached from $x^{\mathcal{M}_{\Sigma_o}}$ in $\mathcal{M}_{\Sigma_o}$ if only Σ transitions were allowed. The states in $X_x^{\mathcal{M}_{\Sigma_o}}$ would be reachable from $x^{\mathcal{M}_{\Sigma_o}}$ in $\mathcal{M}_{\Sigma_o}$ no matter what events were added to the observability set because only Σ'_o transitions can be cut. With this in mind, let us build the nondeterministic machine, $\tilde{\mathcal{M}}_{\Sigma_o}$ from $\mathcal{M}_{\Sigma_o}$. We assume that $d \notin X_{x_0}^{\mathcal{M}_{\Sigma_o}}$. Let $\tilde{\mathcal{M}}_{\Sigma_o} = (X^{\tilde{\mathcal{M}}_{\Sigma_o}}, x_0^{\tilde{\mathcal{M}}_{\Sigma_o}}, \Sigma^{\tilde{\mathcal{M}}_{\Sigma_o}}, \delta^{\tilde{\mathcal{M}}_{\Sigma_o}})$, where $X^{\tilde{\mathcal{M}}_{\Sigma_o}} := X^H \times X^H \times X^G \cup \{d\}, x_0^{\tilde{\mathcal{M}}_{\Sigma_o}} := (x_0^H, x_0^H, x_0^G)$ and $\Sigma^{\tilde{\mathcal{M}}_{\Sigma_o}} := \Sigma$.

The transition relation $\delta^{\tilde{\mathcal{M}}_{\Sigma_o}}$ is defined as follows. Suppose there exists $x^{\mathcal{M}_{\Sigma_o}}, y^{\mathcal{M}_{\Sigma_o}}, z^{\mathcal{M}_{\Sigma_o}} \in X^{\mathcal{M}_{\Sigma_o}}$ and $\sigma \in \Sigma$ such that $z^{\mathcal{M}_{\Sigma_o}} \in X_x^{\mathcal{M}_{\Sigma_o}}$, $\delta^{\mathcal{M}_{\Sigma_o}}(z^{\mathcal{M}_{\Sigma_o}}, \sigma') = y^{\mathcal{M}_{\Sigma_o}}$.

$$\delta^{\tilde{\mathcal{M}}_{\Sigma_o}}(x^{\mathcal{M}_{\Sigma_o}}, \sigma) = \begin{cases} y^{\mathcal{M}_{\Sigma_o}} & \text{if } d \notin X_y^{\mathcal{M}_{\Sigma_o}} \\ d & \text{if } d \in X_y^{\mathcal{M}_{\Sigma_o}} \end{cases}$$

Note that $\tilde{\mathcal{M}}_{\Sigma_o}$ is effectively a colored directed graph, with states as vertices, transitions as directed edges and transition labels for colors. This prompts one of the main results of this paper.

Theorem 4. Given a $\tilde{\mathcal{M}}_\emptyset$ machine ultimately constructed from H, G, Σ_c and $\emptyset$ as the set of observable events, $\mathcal{L}(H)$ is observable with respect to $\mathcal{L}(G)$, Σ_o and Σ_c if and only if Σ_o is a colored $x_0^{\tilde{\mathcal{M}}_{\Sigma_o}}$ d-cut in the colored directed graph $\tilde{\mathcal{M}}_\emptyset$.

With these conversions, any methods developed to approximate solutions to the colored cut problem could also be used to produce approximate solutions to the sensor selection problem and vice-versa.

4. INAPPROXIMABILITY RESULTS

To our knowledge the graph cutting problem has not been explored in the standard literature (from graph theory or computer science). Unfortunately, although many other types of graph cutting problems are computationally simple, it is shown here that solutions to Problem 2 are most likely difficult to approximate. Because of the above results, solutions to the sensor selection problem are similarly difficult to approximate.

Consider a bipartite graph $K = (V_1, V_2, E)$ where the sets V_1 and V_2 are partitioned into a disjoint union of q sets, $V_1 = \cup_{i=1}^q Y_i$ and $V_2 = \cup_{j=1}^q Z_i$. The sets $\{Y_1, Z_1, \ldots, Y_q, Z_q\}$ all have size N. E is the set of edges in K. The bipartite graph and the partitions of V_1 and V_2 induce a super-graph $\mathcal{H}$ where the vertices of $\mathcal{H}$ are the sets $\{Y_1, Z_1, \ldots, Y_q, Z_q\}$. Y_i and Z_j are connected by a (super)edge in $\mathcal{H}$ if and only if there exists $(a, b) \in Y_i \times Z_j$ that are adjacent in K. For our purposes, it is convenient to assume that $\mathcal{H}$ is d-regular.

Given $X \subseteq V_1 \cup V_2$, we say that the super-edge (Y_i, Z_j) in $\mathcal{H}$ is *covered* by X if there exists two nodes $a \in X \cap Y_i$ and $b \in X \cap Z_j$ such that $(a, b) \in E$. We can now formally introduce the MIN-REP problem.

Problem 5. MIN-REP: Given a bipartite graph K as introduced above, find a subset X of minimum size covering all superedges of $\mathcal{H}$.

The subset $X \cap Y_i$ is referred to as the *representatives* of Y_i in X (similarly for Z_j). We say that X is a set of *unique representatives* if $|X \cap Y_i| = 1$ and $|X \cap Z_j| = 1$ for all i, j.

In addition we will need is the following result from Kortsarz (1998) that shows the approximation difficulty of Problem 5. Consider the following theorem based on the well-known satisfiability problem (SAT) from computer science.

Theorem 6. Raz (1998). Let T be an instance of SAT. For any $0 < \epsilon < 1$, there exists a (quasi-polynomial) reduction of T to K, an instance of MIN-REP with n vertices so that if T is satisfiable, then there exists a set of unique representatives which cover all super-edges. If T is not satisfiable, then the size of any MIN-REP solution has at least $2q2^{\log^{(1-\epsilon)} n}$ vertices.

In the reduction in Theorem 6, if T is not satisfiable, then the average number of representatives needed per supernode is $\Omega(2^{\log^{(1-\epsilon)} n})$. Also, n is quasi-polynomial in the size of the SAT formula. This is used in Kortsarz (1998) to show the approximation difficulty of the MIN-REP problem.

Theorem 7. Kortsarz (1998). The MIN-REP problem admits no $2^{\log^{(1-\epsilon)} n}$ approximation for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.

We can now show a conversion of the MIN-REP problem, $K = (V_1, V_2, E)$, to the colored cut problem, $D = (V, A, C)$ that establishes a similar connection between SAT and the colored cut problem and ultimately the sensor selection problem. For simplicity, we assume that D may be a multi-graph, but by using a standard trick of subdividing edges, we can remove this assumption. We construct D in three parts.

First we construct the vertices V of D. Initially assign $V := V_1 \cup V_2$. For the edges $\{e_{ij}^1, \ldots, e_{ij}^b\}$ where $e_{ij}^k = (y^k, z^k)$ and $y^k \in Y_i$ and $z^k \in Z_j$, add a vertices $\{x_{ij}^1, \ldots, x_{ij}^b\}$ to V. For each superedge $e_{ij} = (Y_i, Z_j)$ in $\mathcal{H}$, add a vertex x_{ij}^{end} to V. Finally, add nodes s and t to V.

We now construct the edges A. Initially assign $A := E$, with all edges directed from the V_1 vertices to the V_2 vertices. Also add directed edges from s to all V_1 vertices. For some superedge $e_{ij} = (Y_i, Z_j)$, arbitrarily order the edges between Y_i and Z_j, $\{e_{ij}^1, \ldots, e_{ij}^b\}$. For each vertex in $z \in Z_j$, add an edge from z to x_{ij}^1. Add two parallel edges from x_{ij}^1 to x_{ij}^2, from x_{ij}^2 to x_{ij}^3 and so on until x_{ij}^b is reached. Then, add an edge from x_{ij}^b to x_{ij}^{end}. We call this construction of edges from the vertices in Z_j to x_{ij}^{end} the *chain* of superedge e_{ij}. Let X_{end} be the set of all x_{ij}^{end} vertices. Add additional edges from all of the X_{end} vertices to t. Observe that there is one chain associated with Z_j for every Y_i adjacent to it.

We finally color the edges of D. For each $y \in Y_i$, $z \in Z_j$ let there be colors c_y and c_z. For the paired edges $e_{ij}^k = (y^k, z^k)$ used to construct the chains, color one of the edges (x_{ij}^k, x_{ij}^{k+1}) in the chain c_{y^k} and the other c_{z^k}. The edges touching s and t are colored new distinct colors along with the edges originally in E.

Make the edges from s into V_1, from X_{end} into t and the edges of E *hard to cut* as follows. Suppose there are m vertices in the graph before edges are made hard to cut. An edge (x, x') can be made hard to cut by adding vertices $\{r_{xx'}^1, \ldots, r_{xx'}^{m^3}\}$. An edge is added from x to each $r_{xx'}^i$, and an edge is added from each $r_{xx'}^i$ to x'. Each of the new $2m^3$ edges is given a new distinct color. Now, cutting edge (x, x') corresponds to adding a very large number of new colors that also need to be cut, effectively making the cutting of (x, x') impossible in the solution.

Now that we have shown this construction, we discuss why this shows the colored cut problem is difficult to approximate. Consider a super-edge $e_{ij} = (Y_i, Z_j)$ and its chain. Let I be any feasible colored st-cut of D such that I does not contain colors for hard to cut edges.

Claim 8. There exists at least one parallel pair of edges in the chain corresponding to e_{ij} so that I contains both colors corresponding to this parallel pair of edges.

PROOF. Suppose that for each pair of parallel edges, at most one color is in I. This implies that even with I colored edges, there is a path from all Z_j vertices to t. All edges on paths to the Z_j from the vertices from s are hard to cut as are all paths from the X_{end} vertices to t. Therefore, if there is not a pair of parallel edges both cut in every chain, then I is not an st-cut. □

Let I be a colored st-cut for the colored cut problem constructed above.

Corollary 9. Let X be the vertices corresponding to the colored cut I. X is a feasible solution to the MIN-REP problem and $|X| = |I|$.

PROOF. Consider a superedge (Y_i, Z_j). There is a parallel pair of edges such that its colors are in I by Claim 8. The color of each edge corresponds to some vertex $y \in Y_i$ and $z \in Z_j$ such that $(y, z) \in E$. Hence, I defines in a natural way a collection $X \subseteq V$ that is a solution for MIN-REP. Therefore, the X collection covers all the superedges. Furthermore, the number of colors equals the number of vertices from X. □

Claim 10. Given any MIN-REP solution X for K, we can find a subset I of $|X|$ colors that solves the colored cut problem constructed from K.

PROOF. Given that $x \in X$, this defines a color c_x corresponding to x in the construction. Let I be the collection of all these colors. We show that I is a feasible solution for D.

Let $e_{ij} = (Y_i, Z_j)$. Consider the chain of e_{ij}. As e_{ij} is covered by X, there are $y \in X \cap Y_i$ and $z \in Z_j$ so that $(y, z) \in E$. By definition, there is a parallel pair of edges on this chain from x_{ij}^k corresponding to (y, z) and its colors are in I. The cutting of edges colored c_y and c_z cuts both of these parallel edges from D. It follows that there cannot be a path through the chain of e_{ij} to t. Since all superedges are covered by X, it follows that s cannot reach t after the edges colored by I are removed. Furthermore, $|I| = |X|$. □

The following is thus immediate.

Theorem 11. Let T be an instance of SAT. For any $0 < \epsilon < 1$, there exists a (quasi-polynomial) reduction of T to D, an instance of colored cut with n vertices so that if T is satisfiable, then there exists a set of $2q$ colors forming a feasible solution to the colored cut problem. If T is not satisfiable, then any colored cut solution has at least $2q2^{\log^{(1-\epsilon)} n}$ colors.

Corollary 12. The colored cut problem admits no $2^{\log^{(1-\epsilon)} n}$ approximation for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.

Corollary 13. The sensor selection problem admits no $2^{\log^{(1-\epsilon)} n}$ approximation for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.

It follows that if we can approximate solutions to either the colored cut or sensor selection problems with better than a $2^{\log^{(1-\epsilon)} n}$-approximation, then we have found a method for solving NP-complete problems in quasipolynomial time. This lower bound is generally considered to be a very poor lower bound in the computer science community. Indeed, as ϵ approaches 0, then $2^{\log^{(1-\epsilon)} n}$ approaches n.

5. ALGORITHM FOR A SPECIAL CASE

Now that we have shown that colored cuts and sensor selections are in a sense difficult to approximate, we show an approximation method for a special case with a reasonable approximation bound. Suppose we have an instance of the colored cut problem $D = (V, A, C)$ with vertices $s, t \in V$

such that there are no parallel edges in the same direction. K is a prespecified value which we will optimize in the following proof.

Algorithm 1.
Input: $D = (V, A, C), s, t \in V$ and K, an integer.
Output: A colored st-cut I for D.
Initialize: $I := \emptyset$, a set of colors.
As long as there is a st path of length at most K, add the colors associated with this path to I and remove all edges in D with colors in I.
After all paths of length at most K have been removed, let k be the distance from s to t.
Perform a depth-first-search to convert D to a layered directed graph D' of depth k.
Let V_i denote the set of vertices at distance i from s in D' and let A_i denote the the set of edges from V_i to V_{i+1}.
Select the set of edges A_j that uses the least number of colors and add those colors to I.
Return: I.

Theorem 14. Algorithm 1 gives a $|V|^{2/3}$ approximation for colored cut problems where D contains no parallel edges in the same direction.

PROOF. Let I^{min} represent the minimal colored cut. Suppose in the first phase of the algorithm, paths $\{P_1, \ldots, P_l\}$ are removed. At most lK colors are chosen during this step because there are at most K colors per path. The colors used by the paths are disjoint, but the optimal solutions needs to remove at most one color for each path. Hence, $l \leq |I^{min}|$. If $\{I_1, \ldots, I_l\}$ are the colors removed by each path, then $\Sigma_{i=1}^{l}|I_i| \leq K|I^{min}|$.

Let I_{A_j} represent the colors cut in the last part of the algorithm. Also, let $|A_p| = \min_i |A_i|$. It is known that $\Sigma_{i=1}^{k}|V_i| < N$ where k is the distance from s to t, $k \geq K$. Also, $A_i \leq |V_i||V_{i+1}|$. Therefore, $|I_{A_j}| \leq |I_{A_p}| \leq |A_p| \leq \left(\frac{N}{K}\right)^2$.

The optimal solution has at least one color from this cut that is different from the colors that have been added to I. $I = \Sigma_{i=1}^{l}|I_i| + |I_{A_j}| \leq K|I^{min}| + \left(\frac{N}{K}\right)^2$. By choosing $K = N^{2/3}$, we get a $N^{2/3}$ approximation. □

When a sensor selection problem is converted to a colored cut problem using the conversion outlined above, methods developed to solved the colored cut problem can be used to solve the sensor selection problem. In Algorithm 1, the restriction that D contains no parallel edges in the same direction is generally fairly restrictive for solving many interesting sensor selection problems. However, this method is still interesting and relevant from a graph theoretic and computer science point of view. We hope that by presenting this algorithm we might spur further research into approximation algorithms for other special cases of the colored cut and sensor selection problems. Our conversion of the sensor selection problem to the colored cut problem is helpful for analysis of the sensor selection problem as graph cutting problems are much more intuitive and this would helpfully aid further research on these problems.

6. DISCUSSION

We have shown how the theory of approximation algorithms can be applied in a formal manner for the analysis of the sensor selection problem in supervisory control. Although not discussed here, our contribution of using graph cutting methods to perform sensor selections allows us to force some events to be observable or unobservable depending on the construction of $\mathcal{M}_{\Sigma_o}$. Beyond the sensor selection problem, our methods could also be directly applied for the analysis of the corresponding actuator selection problem. There are also extensions to decentralized control and hybrid systems.

REFERENCES

S. Arora and C. Lund. Hardness of approximations. In D.S. Hochbaum, editor, *Algorithms for NP-Hard Problems*, pages 46–93. PWS Publishing, 1997.

G. Ausiello, P. Crescenzi, G. Gambosi, V. Kahn, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation Combinatorial Optimization Problems and Their Approximability Properties*. Springer-Verlag, Berlin Heidelberg, 1999.

C.G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, Boston, MA, 1999.

G. Kortsarz. On the hardness of approximating spanners. In *The First International Workshop on Approximation*, pages 135–146, 1998.

F. Lin and W. M. Wonham. On observability of discrete-event systems. *Information Sciences*, 44:173–198, 1988.

R. Raz. A parallel repetition theorem. *SIAM Journal of Computing*, 27(3):763–803, 1998.

K. Rohloff. *Computations on Distributed Discrete-Event Systems*. PhD thesis, The University of Michigan, Ann Arbor, 2004. Available at `http://www.eecs.umich.edu/umdes`.

J. Tsitsiklis. On the control of discrete-event dynamical systems. *Mathematics of Control, Signals and Systems*, 2:95–107, 1989.

T.-S. Yoo and S. Lafortune. NP-completeness of sensor selection problems arising in partially-observed discrete-event systems. *IEEE Trans. Auto. Contr.*, 47(9):1495–1499, 2002.