

Searching the Nodes of a Graph: Theory and Algorithms

Ath. Kehagias, G. Hollinger and A. Gelastopoulos*

19 May, 2009

Abstract

One or more *searchers* must *capture* an invisible *evader* hiding in the *nodes* of a graph. We study this version of the *graph search* problem under additional restrictions, such as monotonicity and connectedness. We emphasize that we study *node search*, i.e., the capture of a *node-located* evader; this problem has so far received much less attention than *edge search*, i.e., the capture of an *edge-located* evader.

We show that in general graphs the problem of node search is easier than that of edge search, namely, every edge clearing search is also node clearing, but the converse does not hold in general (however node search is NP-complete, just like edge search). Then we concentrate on the *internal monotone connected* (IMC) node search of *trees* and show that it is essentially equivalent to IMC edge search; hence Barriere's tree search algorithm [2], originally designed for edge search, can also be used for node search.

We return to IMC node search on general graphs and present (several variants of) a new algorithm: *GSST* (Guaranteed Search by Spanning Tree). *GSST* clears a graph \mathbf{G} by performing all its clearing moves along a spanning tree \mathbf{T} of \mathbf{G} . Because spanning trees can be generated and cleared very quickly, *GSST* can test a large number of spanning trees and find one which clears \mathbf{G} with a small (though not necessarily minimal) number of searchers. We prove the existence of probabilistically complete variants of *GSST* (i.e., these variants are guaranteed to find a minimal IMC node clearing schedule if run for sufficiently long time). Our experiments also indicate that *GSST* can efficiently node-clear large graphs given only a small running time. *An implementation of GSST* (running on Windows and Linux computers) *is also provided and made publicly available*.

1 Introduction

In this paper we study a variant of *graph search*. Our motivation comes from applied robotics problems. As an introduction, consider the following two problems, which can both be modeled as graph searches.

1. An *evader* is hiding in a *building* and a team of one or more *searchers* are trying to *capture* him. The building is represented by a graph \mathbf{G} : nodes are *rooms* and edges are *doors* between rooms. Both searchers and evader occupy the nodes of the graph and move from node to node by sliding along the edges. The evader is captured when he is located in the same node as a searcher. It is assumed that the evader has the following characteristics: (a) he wants to avoid capture, (b) he is invisible to the searchers (unless he is located in the same node as one of them), (c) he is always aware of the searchers' locations and (d) he is arbitrarily fast.
2. An evader and a team of searchers are located inside a *cave*. The cave is again represented by a graph \mathbf{G} : edges are *tunnels* and nodes are the *intersections* of tunnels. The searchers move from node to node, sliding along the edges; capture takes place if a searcher slides along an edge occupied by the evader, or if the evader moves through a node occupied by a searcher. The previously mentioned properties of the evader are assumed to hold in this case as well.

*emails: kehagiat@auth.gr, geoff.hollinger@gmail.com, alexgelas@hotmail.com.

The two problems are very similar but differ in one important respect. In the first problem the evader is assumed to reside in the *nodes* of the graph (the rooms); in the second problem he is assumed to reside in the *edges* (the tunnels). Accordingly, we will call the first problem *node search* and the second *edge search*. Edge- or node-location of the evader is an important aspect of the graph search problem and can be used to categorize the related literature. Historically, research has concentrated on the problem of “*edge-located* evader”. To the best of our knowledge, very little work has been done on the problem of “*node-located* evader”. Interestingly, the term “node search” has been used in the past to denote a version of the graph search for edge-located evader [25]. Hence, to avoid confusion, let us repeat that *we use “node search” to denote the search for a node-located evader*.

Another important aspect of the graph search problem is the “visibility” of the evader. Namely, we say that the evader is *visible* if the searchers always know his location; we say that the evader is *invisible* if the evaders may become aware of his location only when he is in the same node (or edge) as they.

Additional aspects of the graph search include whether it is *internal* (the searchers can only move along the edges of the graph) or not (a searcher can, in one step, move to any node of the graph – we call this “*teleporting*”), *monotone* (the evader can never return to a searched part of the graph) and *connected* (the searched parts form a connected subgraph). These terms will be defined rigorously in Section 2.

We now present a brief overview of the graph search literature. Early papers include [6] where the problem of cave searching was posed and [31] where the first mathematical study of the problem was presented. The problem was independently rediscovered a little later in [32]. The edge search problem was placed in the context of graph theory in [29]. The already mentioned “node search” variant (actually dealing with edge-located evader) appeared a little earlier in [25]. Another variant, called “mixed search” (and again dealing with edge-located evader) appeared in [5] and was further studied in [35, 38]. The study of *connected* edge search started relatively recently; see [2, 3, 11, 13, 14, 39]. There is very little published on “true” node search, i.e., the search for a node-located evader. There is a considerable literature on pursuit of *visible* node-located evaders, which we will not discuss, since visible evader search is outside the scope of the current paper; let us simply mention that an important early paper is [30]. Graph search is also related to several *graph parameters*, most notably *pathwidth* and *vertex separation*. These connections are discussed in many papers, e.g., in [4, 8, 9, 24, 25, 36]. Finally there are several reviews of the graph search literature: an old and deep one is [4], a more recent one is [1] and a very recent and very extensive one is [12].

The above mentioned papers adopt a theoretical point of view. While there is much discussion about graph search algorithms, we have found few actually implemented algorithms which can tackle relatively large graphs (with the exception of *tree* search algorithms). A much more applied line of research appears in the *robotics* literature, for example in [15, 16, 17, 27] and the book [28]. These papers present actual implementations of algorithms, as well as numerical and even physical experiments, but they provide little (if any) theoretical justification of their methods. In particular, the distinction between edge and node search often appears to be misunderstood. We attempt to clarify this distinction and also combine the robotics and graph theoretic points in two technical reports we have previously published [20, 23].

In the current paper we study graph search for an invisible, node-located evader. We also examine connections to the problem of invisible, edge-located evader. We do *not* deal with the case of visible evader. We are mainly interested in *internal, monotone, connected* (IMC) node search. This version of the problem is the one most relevant to robotics applications which, as already stated, provide our main motivation. In general, we try to provide a balanced combination of theory and implementation.

The main contribution of the paper is the introduction of (several variants of) the node-search algorithm GSST (*Guaranteed Search by Spanning Tree*). This algorithm is presented in detail, theoret-

ically motivated and analyzed and also practically evaluated by a number of numerical experiments¹. The *main idea* behind the algorithm is that every IMC node search generates a spanning tree of the graph; conversely, every spanning tree can be used to perform a node search. Trees can be searched much more efficiently than general graphs; GSST exploits this fact to quickly compute a large number of searches and then output the one which requires the least number of searchers². Our experiments also indicate that GSST can efficiently node-clear large graphs in a small running time. From the theoretical point of view, we establish several results about node search and its relation to edge search. These results are used to motivate the GSST algorithm and also to prove that some variants of GSST are probabilistically complete (i.e., they are guaranteed to find a minimal IMC node clearing schedule if given sufficient time).

The paper is organized as follows. Preliminary concepts and notation are introduced in Section 2; edge search and node search are compared in Section 3; IMC node and edge search on *trees* is studied in Section 4; IMC node search on general graphs is studied in Section 5 where the GSST search algorithm is also introduced; the algorithm is evaluated by numerical experiments in Section 6; conclusions and future research directions appear in Section 7. Some edge search results by Barriere et al. [2, 3] are presented in Appendix A; the connection of node search to *mixed* edge search is discussed in Appendix B; connections to *pathwidth* are discussed in Appendix C; an implementation of GSST as an actual executable program (which runs on Windows and Linux computers) is discussed in Appendix D.

2 Preliminaries

2.1 Basic concepts and notation

We denote graphs by boldface letters, e.g. $\mathbf{G} = (V, E)$ where V is the *node set* and E is the *edge set*. We will always label the nodes of \mathbf{G} as $V = \{1, 2, \dots, N\}$ (hence the graph contains N nodes, i.e. $|V| = N$). Edges are denoted as $\{u, v\}$ where $u, v \in V$; often we will write uv instead, for simplicity, but note that vu and uv are the same edge (i.e. we study *undirected* graphs). We only consider connected graphs without loops or multiple edges. *Also, to avoid some trivial cases, we will always consider graphs with at least one edge (and at least two nodes).*

Nodes u, v are said to be *neighbors* iff $uv \in E$.

Given a node sequence $u_1u_2\dots u_L$, where $u_iu_{i+1} \in E$ for $i = 1, 2, \dots, L - 1$, we say that $u_1u_2\dots u_L$ is

1. a *path* iff $u_i \neq u_j$ for $i, j \in \{1, 2, \dots, L\}$ and $i \neq j$;
2. a *cycle* iff $u_i \neq u_j$ for $i, j \in \{1, 2, \dots, L\}$ (and $i \neq j$) except that $u_1 = u_L$.

A *tree* is a connected graph without any cycles. Equivalent definitions are: a tree is a graph for which there is a unique path between every pair of nodes; a tree is a connected graph with N nodes and $N - 1$ edges. The *leaves* of the tree are the nodes which have exactly one neighbor.

A *rooted tree* is a tree with a distinguished node u_0 called the *root* of the tree. Given a tree $\mathbf{T} = (V, E)$, we will denote the same tree rooted at u_0 by $\mathbf{T}_{u_0} = (V, E, u_0)$. Take any node $u_L \neq u_0$ and let $u_0u_1\dots u_{L-1}u_L$ be the unique path from u_0 to u_L ; then u_{L-1} is the *parent* of u_L ; u_L is the *child* of u_{L-1} . Given a node x , its children, their children and so on are the *descendants* of x . Given a rooted tree $\mathbf{T}_x = (V, E, x)$ and a node y , consider the node set

$$V[y] = \{z : z \text{ is } y \text{ or a descendant of } y\}$$

¹An implementation of the algorithm (for Windows and Linux computers) is available in the public domain; for details see Appendix D.

²In fact GSST is an “anytime” algorithm [40]. Namely, the longer it runs the better solutions are provided, but a reasonable solution becomes available even in the initial stages of the algorithm.

and the edge set

$$E[y] = \{zu : zu \in E \text{ and } z, u \in V[y]\};$$

note that both $V[y]$ and $E[y]$ are determined by *both* y and x . The graph $(V[y], E[y])$ is a tree (a *sub-tree* of \mathbf{T}_x). The *rooted* tree $(V[y], E[y], y)$ will be denoted by $\mathbf{T}_x[y]$.

A *search schedule* on the graph $\mathbf{G} = (V, E)$ is a sequence of *ordered* pairs of nodes:

$$\mathbf{S} = ((u(1), v(1)), (u(2), v(2)), \dots, (u(t_{fin}), v(t_{fin}))),$$

subject to $\{u(t), v(t)\} \in E$ for $t = 1, 2, \dots, t_{fin}$. We call $\mathbf{S}(t) = (u(t), v(t))$ the t -th move (or step) of the search schedule. The t -th move can be any one of the following

1. *placing* a searcher at node v (in which case $u(t) = 0$ and $v(t) \in V$) or
2. *sliding* a searcher from node $u(t) \in V$ to node $v(t) \in V$ or
3. *removing* a searcher from node u (in which case $u(t) \in V$ and $v(t) = 0$).

We will also use the more evocative notation $\mathbf{S}(t) = (u \rightarrow v)$, i.e. a searcher is moved from node u to node v . Moves of the form: $0 \rightarrow v$ (a new searcher is introduced in the graph) and $u \rightarrow 0$ (a searcher is removed from the graph) involve the fictitious “source node” 0 (it is not an element of V), in which searchers are kept whenever they are not actively involved in the graph search; the evader does not have access to the source node.

Given a search \mathbf{S} , the number of searchers inside the graph at time t will be denoted by $sn(\mathbf{S}, t)$. The maximum number of searchers used by \mathbf{S} will be denoted by $\overline{sn}(\mathbf{S})$, i.e.

$$\overline{sn}(\mathbf{S}) = \max_t sn(\mathbf{S}, t).$$

These numbers must not be confused with the search number of a *graph*, which will be defined in Section 2.2.

Finally, we say that a node u is *guarded* at time t iff a searcher is located at u (at time t); otherwise we say u is *unguarded*. A path $u_1u_2\dots u_L$ is called *n-unguarded* (node unguarded) iff nodes u_i ($i = 1, 2, \dots, L$) are unguarded (at time t); otherwise it is called *n-guarded*. The path is called *e-unguarded* (edge unguarded) iff nodes u_i ($i = 2, 3, \dots, L - 1$) are unguarded (at time t) and e-guarded otherwise. It is easy to see that

$$\begin{aligned} u_1u_2\dots u_L \text{ is e-unguarded} &\Rightarrow u_2u_3\dots u_{L-1} \text{ is n-unguarded,} \\ u_1u_2\dots u_L \text{ is n-unguarded} &\Rightarrow u_1u_2\dots u_L \text{ is e-unguarded.} \end{aligned}$$

The reason for the two different definitions is that the first pertains to *node* recontamination and the second to *edge* recontamination, as will be seen in Section 2.2.

2.2 Node and edge search

We repeat the assumptions introduced in Section 1 regarding the evader. Namely, the evader wants to avoid capture, he is invisible to the searchers (unless located in the same node), he is always aware of the searchers’ locations and arbitrarily fast. The net result of all these assumptions is that (in both node and edge search) the evader can (and will) always avoid capture if an escape route is available. Hence, from the searchers’ point of view, we can think of graph search as a process of eliminating escape routes. This is expressed as follows:

1. a node / edge is *dirty* if it can *possibly* contain the evader and *clear* otherwise (e.g., a node occupied by a searcher is clear);

2. a previously clear node / edge can become dirty (e.g., when a previously guarded path between a clear and a dirty node becomes unguarded) – this is called *recontamination*;
3. graph search is the process of gradually decreasing the *dirty set* (of nodes or edges) until it becomes the empty set (i.e., the evader has no escape route left).

This is a *worst case approach* which essentially eliminates the evader from the graph search, introducing in his place the dirty set. We can think of node (edge) search as a *one-player* “node game” (“edge game”). In both these games the player controls all searchers, much like pieces in a game of chess. We now present the rules of the two games. For reasons which will be explained presently, we substitute the terms “clear” and “dirty” with the terms “n-clear” and “n-dirty” (in the node game) and “e-clear” and “e-dirty” (in the edge game).

Rules of the Node Game

- N0** At time $t = 0$ all nodes are *n-dirty* and no searcher is located in the graph.
- N1** At every time $t = 1, 2, \dots$ the player performs *one* of the following moves:
- N1a** place a searcher on a node,
 - N1b** remove a searcher from a node,
 - N1c** slide a searcher along one edge.
- N2** An n-dirty node becomes *n-clear* when occupied by a searcher.
- N3** An n-clear node u becomes n-dirty when it is connected to to an n-dirty node v by an *n-unguarded path*.
- N4** An edge is n-dirty if it is adjacent to an n-dirty node; otherwise it is clear.
- N5** The game is concluded when all nodes (and consequently also all edges) are n-clear.

Rules of the Edge Game

- E0** At time $t = 0$ all edges are *e-dirty* and no searcher is located in the graph.
- E1** At every time $t = 1, 2, \dots$ the player performs *one* of the following moves:
- E1a** place a searcher on a node,
 - E1b** remove a searcher from a node,
 - E1c** slide a searcher along one edge.
- E2** An e-dirty edge becomes e-clear when a searcher slides along the edge.
- E3** An e-clear edge uv becomes e-dirty when it is connected to to a e-dirty edge xy by an *e-unguarded path*.
- E4** A node is e-dirty if it is unguarded and adjacent to an e-dirty edge; otherwise it is e-clear.
- E5** The game is concluded when all edges (and consequently also all nodes) are e-clear.

Remark 2.1. The reason for using “n-clear” and “e-clear” (instead of simply “clear”) is that an edge can be clear in the node game and dirty in the edge game. An example will illustrate this point. Consider the graph of Fig.1 and the search schedule $0 \rightarrow 1, 0 \rightarrow 1, 1 \rightarrow 2, 2 \rightarrow 4, 4 \rightarrow 3$. If this schedule is executed in a node game, after the final move all nodes are n-clear and so all edges are also n-clear. But in an edge game, after the final move, edge $\{1, 3\}$ has still not been traversed and hence it is still e-dirty. More generally, while “n-clear” and “e-clear” refer to similar physical situations, there is no a priori reason that their mathematical definitions are equivalent (in Section 4 we will show that they *are* equivalent *provided certain conditions are satisfied*).




Figure 1: A graph in which node clearing and edge clearing are not equivalent. For example, placing a searcher into node 0 ($0 \rightarrow 1$), then placing a second searcher at 1 and sending him through 2, 4, 3 ($0 \rightarrow 1, 1 \rightarrow 2, 2 \rightarrow 4, 4 \rightarrow 3$) is node clearing but not edge clearing.

Remark 2.2. Note that in the edge game clear edges are “recorded” at the completion of each move. This is important to note in the case that an edge is e-cleared and e-dirtied *during the same move* (i.e. for a single value of t). To clarify this, suppose that in the graph of Fig.1 the following moves are executed: $0 \rightarrow 1$ at $t = 1$; $1 \rightarrow 2$ at $t = 2$. In this case edge $\{1, 2\}$ is first e-cleared (because traversed) and then e-dirtied, *both during* $t = 2$. After completion of the second move, *edge* $\{1, 2\}$ *is e-dirty*.

We now introduce more detailed notation regarding clear/dirty nodes/edges. For the node game we will use the following notation.

$$\begin{aligned}
 V_N^C(t) &= \{u \in V : \text{at time } t, u \text{ is n-clear}\}; \\
 V_N^D(t) &= \{u \in V : \text{at time } t, u \text{ is n-dirty}\}; \\
 E_N^C(t) &= \{uv \in E : \text{at time } t, uv \text{ is n-clear}\}; \\
 E_N^D(t) &= \{uv \in E : \text{at time } t, uv \text{ is n-dirty}\}; \\
 \mathbf{G}_N^C(t) &= (V_N^C(t), E_N^C(t)) \text{ is the } n\text{-clear graph at time } t.
 \end{aligned}$$

We use exactly analogous notation for the edge game, but with an E (rather than N) subscript: $V_E^C(t), V_E^D(t), E_E^C(t), E_E^D(t), \mathbf{G}_E^C(t) = (V_E^C(t), E_E^C(t))$. Obviously we have

$$\begin{aligned}
 V_N^C(t) \cup V_N^D(t) &= V, & V_N^C(t) \cap V_N^D(t) &= \emptyset, & E_N^C(t) \cup E_N^D(t) &= E, & E_N^C(t) \cap E_N^D(t) &= \emptyset, \\
 V_E^C(t) \cup V_E^D(t) &= V, & V_E^C(t) \cap V_E^D(t) &= \emptyset, & E_E^C(t) \cup E_E^D(t) &= E, & E_E^C(t) \cap E_E^D(t) &= \emptyset.
 \end{aligned}$$

Remark 2.3. The n-clear / e-clear graph definitions are consistent, i.e. $(V_N^C(t), E_N^C(t)) / (V_E^C(t), E_E^C(t))$ are graphs. More specifically:

1. in the node game $\mathbf{G}_N^C(t)$ consists of the n-clear nodes and all edges between such nodes (these are exactly the set $E_N^C(t)$);
2. in the edge game $\mathbf{G}_E^C(t)$ consists of the e-clear edges and none of these can be left “dangling”, i.e. if $uv \in E_E^C(t)$ then $u, v \in V_E^C(t)$ (we will prove this as Lemma 3.1 in Section 3); in addition $V_E^C(t)$ may contain guarded nodes adjacent to e-dirty edges; such nodes may appear as disconnected components of $\mathbf{G}_E^C(t)$.

In the current paper we will concentrate on some *restricted* versions of the edge and node games. More specifically, the restrictions concern the allowable search schedules and they are the following.

1. We call a search schedule *rooted* iff searchers can be placed only into a single, prespecified node u_0 , called the *root* of the search.
2. We call a search schedule *internal* iff searchers once placed into the graph (a) can only move along the edges and (b) are never removed from the graph. In other words, “teleporting” (the arbitrary movement of a searcher from one node to another, irrespective of the graph connectivity) is not allowed³.
3. In the node game we call a search schedule \mathbf{S} *monotone* iff the clear node set is never decreasing: $V_N^C(t) \subseteq V_N^C(t+1)$ for all t . In other words, once a node is n-cleared it never becomes n-dirty again. The definition is similar for the edge game; i.e., \mathbf{S} is monotone in the edge game iff $E_E^C(t) \subseteq E_E^C(t+1)$ for all t .⁴
4. In the node game (resp. edge game) we call a search schedule *connected* iff the clear graph $\mathbf{G}_N^C(t)$ (resp. $\mathbf{G}_E^C(t)$) is connected for $t = 1, 2, \dots$

In Section 3 we will study general, unrestricted graph search; in the remaining sections we will concentrate on internal monotone (*IM*), internal connected (*IC*) and, especially, internal monotone connected (*IMC*) node / edge search. Our focus originates in our interest in practical *pursuit / evasion* problems which arise in robotics. The nature of the problem dictates the restrictions on the search.

1. The internality constraint arises from the fact that a robot cannot “teleport”; it can only move through rooms and doors connecting these rooms; in the graph model of the physical situation this is reflected by the use of edge-sliding moves only.
2. Similarly, rootedness reflects the fact that usually an environment has a single entrance through which robots must enter.
3. Connectedness and monotonicity are desirable (but not mandatory) characteristics of robotic search. For example, in a hostile environment it is preferable that the cleared area consists of a single connected region, which is easier to guard and control than multiple regions. Similarly, if a high cost is associated with clearing an area, it is desirable that this area is cleared only once, i.e., that the search is monotone.

³Since no searchers are ever removed during an internal search \mathbf{S} , we have $\overline{sn}(\mathbf{S}) = sn(\mathbf{S}, t_{fin})$ where t_{fin} is the length of the search.

⁴An important detail is worth mentioning at this point. There is a property, call it *strong monotonicity*, which characterizes schedules in which a *traversed* edge never becomes e-dirty again. The schedule of Remark 2.2. is monotone (since $E_E^C(1) = E_E^C(2) = \emptyset$) but *not* strongly monotone. We stress this detail mostly for the sake of completeness; in what follows we will *not* use strong monotonicity.

A given graph \mathbf{G} with N nodes can always be node-cleared using N searchers: just place one searcher in every node. Usually \mathbf{G} can be cleared with much fewer searchers; the *node search number* of \mathbf{G} , denoted by $s_N(\mathbf{G})$, is the minimum number of searchers required to node-clear \mathbf{G} . Introducing additional constraints on the search schedule we get additional search numbers:

$s_N^i(\mathbf{G})$:	min. nr. of searchers required to clear \mathbf{G} with	internal node search schedule;
$s_N^{im}(\mathbf{G})$:	-//-	IM node search schedule;
$s_N^{ic}(\mathbf{G})$:	-//-	IC node search schedule;
$s_N^{imc}(\mathbf{G})$:	-//-	IMC node search schedule.

For the edge game we define the corresponding edge search numbers: $s_E(\mathbf{G})$, $s_E^i(\mathbf{G})$, $s_E^{im}(\mathbf{G})$, $s_E^{ic}(\mathbf{G})$, $s_E^{imc}(\mathbf{G})$.

In case of a search rooted at node x , we will modify the notation as follows: $s_N(\mathbf{G};x)$ is the minimum number of searchers required to node-clear \mathbf{G} by a search rooted at x ; similarly for $s_E(\mathbf{G};x)$, $s_E^i(\mathbf{G};x)$ and so on. If the graph \mathbf{G} is a rooted tree \mathbf{T}_x , then we will assume it is always searched by rooted searches; hence $s_E(\mathbf{T}_x) = s_E(\mathbf{T};x)$, $s_E^i(\mathbf{T}_x) = s_E^i(\mathbf{T};x)$ and so on.

A *minimal* node search schedule is one which clears \mathbf{G} using $s_N(\mathbf{G})$ searchers; a *minimal* IM node search schedule is an IM schedule which clears \mathbf{G} using $s_N^{im}(\mathbf{G})$ searchers and similarly for all other types of either node or edge search.

3 Comparison of node and edge search

In this section we compare “unrestricted” edge and node search. In other words, we compare the outcome of a search schedule \mathbf{S} when it is used first in a node game and then in edge game, both played on the same graph \mathbf{G} and without requiring rootedness, internality, monotonicity or connectedness.

Our main result is Theorem 3.3 which, informally, says the following: every search schedule \mathbf{S} gives at least as good results in the node game as in the edge game. More precisely, at every step t of the search, the set of e-clear nodes (and edges) is a subset of the set of n-clear nodes (and edges). An immediate corollary is that an edge clearing search schedule is also node clearing; The converse does not hold, as we have already seen by the example of Remark 2.1.

To prove Theorem 3.3 we will need the following lemmas.

Lemma 3.1. *Given a graph \mathbf{G} and a search schedule \mathbf{S} , for $t = 0, 1, 2, \dots$ we have*

$$u \in V_E^D(t) \Rightarrow (\forall ux \in E : ux \in E_E^D(t)) \quad (1)$$

$$uv \in E_E^C(t) \Rightarrow (u, v \in V_E^C(t)) \quad (2)$$

In other words: if a node u is e-dirty, then all edges ux , adjacent to u , are also e-dirty; and if an edge uv is e-clear, then nodes u, v are e-clear.

Proof. For (1): since u is e-dirty, it must be unguarded. We distinguish two cases.

(1.i) Suppose u is adjacent to a single edge, call it uv . If uv is e-clear then u is not adjacent to any e-dirty edges and hence must be e-clear, which is contrary to the original assumption. Hence uv is e-dirty.

(1.ii) On the other hand, if u is adjacent to more than one edges, they cannot all be e-clear (then u would also be e-clear); so at least one edge, call it uw , is e-dirty. Suppose there is also an e-clear edge, call it uv . But then the e-unguarded path wuv connects e-dirty wu to e-clear uv , which leads to contradiction. Hence all edges adjacent to u are e-dirty.

For (2): suppose $u \in V_E^D(t)$. Then u is unguarded and there exists some $ux \in E_E^D(t)$ and $x \neq v$. But the path xuv is e-unguarded and so $uv \in E_E^D(t)$ which is a contradiction. \square

Lemma 3.2. *If a node u is e-cleared exactly at time t , then u is entered by a searcher at t .*

Proof. Suppose $u \in V_E^D(t-1) \cap V_E^C(t)$. This can happen in two ways. (i) u was unguarded and e-dirty at $t-1$ and then entered by a searcher at t . This is exactly the case described by the theorem. (ii) The other way for u to be cleared exactly at t is if u was unguarded and e-dirty at $t-1$ and in addition a single edge, call it ux_0 , was e-dirty at $t-1$ and e-cleared at t (while all other edges ux were e-clear at $t-1$ and remained e-clear at t). For this to happen, the move at time t is either $u \rightarrow x_0$ or $x_0 \rightarrow u$. But $u \rightarrow x_0$ is not possible, because then u would be guarded and e-clear at $t-1$. Hence the move at t is $x_0 \rightarrow u$ and u was entered at t (which shows that (ii) is a sub-case of (i)). \square

Now we are ready to prove the following.

Theorem 3.3. *Given a graph \mathbf{G} and a search schedule \mathbf{S} , we have*

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) \subseteq V_N^C(t) \quad \text{and} \quad E_E^C(t) \subseteq E_N^C(t). \quad (3)$$

Proof. We will actually show the relationships (equivalent to (3))

$$V_N^D(t) \subseteq V_E^D(t), E_N^D(t) \subseteq E_E^D(t).$$

The proof is by induction. At $s = 0$ all nodes are dirty in both the edge game and the node game, hence $V_N^D(0) = V_E^D(0) = V$. Suppose $V_N^D(s) \subseteq V_E^D(s)$ for $s = 1, 2, \dots, t$. Now we take $s = t + 1$ and consider first the cleared and then the recontaminated nodes.

I. Cleared Nodes. By Lemma 3.2, a node becomes e-cleared *exactly* at $t + 1$ only if entered at $t + 1$. Hence at most one such node (call it x_0) exists and, since entered at $t + 1$, it is also n-clear at $t + 1$. In other words: $x_0 \in V_E^D(t) \cap V_E^C(t + 1) \Rightarrow x_0 \in V_N^C(t + 1)$.

II. Recontaminated Nodes. Denote by A the set of nodes recontaminated (in the node game) *exactly* at $t + 1$; i.e., $A = V_N^C(t) \cap V_N^D(t + 1)$. By the definition of n-dirty node, every $u \in A$ must have a path (n-unguarded at $t + 1$) to a node $v \in V_N^D(t + 1)$. Furthermore, at least one $u \in A$ must have a path (n-unguarded at $t + 1$) to a node $v \in V_N^D(t)$ (if this did not hold for any $u \in A$, then none of the nodes of A would get recontaminated to begin with) and in fact this shows *every* $u \in A$ must have a path (n-unguarded at $t + 1$) to a node $v \in V_N^D(t)$. So take any $u \in A$, then there exists a path $v_1 v_2 \dots v_K u$ satisfying:

$$v_1 \in V_N^D(t) \cap V_N^D(t + 1), \quad (4)$$

$$v_1 v_2 \dots v_K u \text{ is n-unguarded at } t + 1. \quad (5)$$

Now we will show that v_1 must belong to $V_E^D(t) \cap V_E^D(t + 1)$ as well. First, $v_1 \in V_N^D(t) \subseteq V_E^D(t)$. Second, if $v_1 \notin V_E^D(t + 1)$ then v_1 was cleared exactly at $t + 1$, which (by Lemma 3.2) means a searcher entered v_1 at $t + 1$; but then $v_1 \notin V_N^D(t + 1)$ which is contrary to (4). Hence $v_1 \in V_E^D(t) \cap V_E^D(t + 1)$. Since v_1 is e-dirty at time $t + 1$, $v_1 v_2$ is e-dirty at $t + 1$ by Lemma 3.1. Also $v_1 v_2 \dots v_K u$ is e-unguarded at $t + 1$. Hence $u \in V_E^D(t + 1)$. Since this is true for every $u \in A$, we conclude

$$A \subseteq V_E^D(t + 1). \quad (6)$$

If the move at $t + 1$ was placing or sliding a searcher into node x_0 , let $B = \{x_0\}$; else let $B = \emptyset$. Since $V_N^D(t) \subseteq V_E^D(t)$ we have

$$V_N^D(t) - B \subseteq V_E^D(t) - B. \quad (7)$$

Since the e-dirty nodes at $t + 1$ include all the e-dirty nodes at t except the node x_0 (which was e-cleared at $t + 1$), we have

$$V_E^D(t) - B \subseteq V_E^D(t + 1). \quad (8)$$

And so we have

$$V_N^D(t) - B \subseteq V_E^D(t+1). \quad (9)$$

The n-dirty nodes at $t+1$ are exactly the n-dirty nodes at t , excluding nodes n-cleared at $t+1$, plus the nodes recontaminated at $t+1$. In other words

$$V_N^D(t+1) = (V_N^D(t) - B) \cup A \quad (10)$$

Combining eqs. (6), (9) and (10) we get

$$V_N^D(t+1) = (V_N^D(t) - B) \cup A \subseteq V_E^D(t+1). \quad (11)$$

Hence by induction we get $V_N^D(t) \subseteq V_E^D(t)$ (and $V_E^C(t) \subseteq V_N^C(t)$) for $t = 0, 1, 2, \dots$.

Let us now consider the edge sets $E_N^D(t)$ and $E_E^D(t)$. Take any t and any edge $uv \in E_N^D(t)$. Without loss of generality we can assume $u \in V_N^D(t) \subseteq V_E^D(t)$ and hence, by Lemma 3.1, $uv \in E_E^D(t)$. Hence, for $t = 0, 1, 2, \dots$, we have

$$E_N^D(t) \subseteq E_E^D(t) \quad \text{and} \quad E_E^C(t) \subseteq E_N^C(t),$$

which completes the proof. \square

Corollary 3.4. *Given a graph \mathbf{G} and an internal search schedule \mathbf{S} , if \mathbf{S} is edge clearing it is also node clearing.*

Proof. Suppose the final move of \mathbf{S} took place at time t_{fin} , then $E_E^D(t_{fin}) = \emptyset$. Since $E_N^D(t_{fin}) \subseteq E_E^D(t_{fin}) = \emptyset$ we have $E_N^D(t_{fin}) = \emptyset$ which also means $V_N^D(t_{fin}) = \emptyset$ (if a node u was n-dirty, an adjacent edge uv would also be n-dirty). \square

Remark 3.5. From Theorem 3.3 we see that edge search is “weaker” than node search, i.e., for every search schedule, and at every time step, the clear set in the edge game is smaller or equal than the one in the node game. There is a variant of edge search, the so-called *mixed edge search* which, as we will see in Appendix B, is *equivalent* to node search, i.e. every search schedule produces the same clear and dirty sets at every step of both the edge and node game (this is stated and proved as Theorem B.1). From this fact follow two consequences. First, *node search is NP-complete* (since mixed edge search is NP-complete [5]). Second, Theorem 3.3 can be obtained as a corollary of Theorem B.1. A more detailed discussion of these issues appears in Appendix B.

4 Search on Trees

In this section we will focus on *rooted IMC (internal, monotone, connected) search on trees* and we will compare node and edge search in this context.

Several authors have studied rooted IMC *edge search* [2, 3, 11, 13, 14, 39]. Our main results in this section are Theorems 4.1 and 4.2 which establish the equivalence of rooted IMC node and edge search on trees. Hence several results already established for edge search also hold for node search (on trees) as will be seen in the sequel.

Theorem 4.1. *Given a tree \mathbf{T} and an internal rooted search schedule \mathbf{S} , if \mathbf{S} is monotone connected in the node game then*

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) = V_N^C(t) \quad \text{and} \quad E_E^C(t) = E_N^C(t)$$

and \mathbf{S} is monotone connected in the edge game.

Proof. The proof is inductive. We have $V_E^C(0) = V_N^C(0) = \emptyset$ and $E_E^C(0) = E_N^C(0) = \emptyset$. Also, since the search is rooted and internal, the move at $t = 1$ will be the placement of a searcher into the root u_0 ; hence $V_E^C(1) = V_N^C(1) = \{u_0\}$ and $E_E^C(1) = E_N^C(1) = \emptyset$. Now suppose that at time $t \geq 1$ we have

$$V_E^C(t) = V_N^C(t), \quad E_E^C(t) = E_N^C(t) \quad (12)$$

and let us consider the move $u \rightarrow v$, performed at $t + 1$, $t > 1$.

We first dispose of the case where the move involves the node 0. Because of internality, $v \neq 0$ (no searcher is removed from the graph). If the move is $0 \rightarrow v$, by rootedness we must have $v = u_0$. Also by rootedness, u_0 is the first node cleared (at some $t = 1$). Now, in the node game: (a) by monotonicity, no node is recontaminated at $t + 1$ and (b) no node is n-cleared either (since u_0 remains n-clear, the move $0 \rightarrow u_0$ implies $V_N^C(t + 1) - V_N^C(t) = \emptyset$); hence $V_N^C(t) = V_N^C(t + 1)$ which also means $E_N^C(t) = E_N^C(t + 1)$. In the edge game, $0 \rightarrow u_0$ (a) does not clear any new edges (so it does not clear any new nodes either) and (b) does not remove any searchers from any node, so neither edge nor node recontamination is possible; hence $V_E^C(t) = V_E^C(t + 1)$ and $E_E^C(t) = E_E^C(t + 1)$. In short, if the $(t + 1)$ -th move involves node 0, from (12) follows

$$V_E^C(t + 1) = V_N^C(t + 1), \quad E_E^C(t + 1) = E_N^C(t + 1).$$

Next we examine the case where the $(t + 1)$ -th move is $u \rightarrow v$, with $u, v \neq 0$. We will examine the effects of this move separately for the node and edge game.

I. Node Game. u was guarded at time t , so $u \in V_N^C(t)$. Suppose, first, that at t only one searcher was located on u , so u is unguarded at $t + 1$. However, since \mathbf{S} was supposed monotone in the node game, there is no node $x \in V_N^C(t) \cap V_N^D(t + 1)$ and, especially, u remains n-clear. Since no node has become n-dirty at $t + 1$, no edge has become n-dirty either (an edge becomes n-dirty iff incident on an n-dirty node). As for clearings, v is n-clear at $t + 1$, since entered by a searcher; and uv is also n-clear since $\{u, v\} \subseteq V_N^C(t + 1)$. No node other than v was entered at $t + 1$, so no node other than v was n-cleared. Is there an edge (other than uv) which was n-cleared exactly at $t + 1$? This is only possible if (a) the edge has the form vx_1 and (b) $v \in V_N^D(t) \cap V_N^C(t + 1)$ (v became n-clear *exactly* at $t + 1$) and (c) $x_1 \in V_N^C(t) \cap V_N^C(t + 1)$. Suppose $x_1 \neq u$. Since $\mathbf{G}_N^C(t)$ is connected, there is a path $ux_{L-1} \dots x_1$ which is *inside* $\mathbf{G}_N^C(t)$ and hence *does not contain* v . Then the tree \mathbf{T} contains a cycle $ux_{L-1} \dots x_1vu$ which is a contradiction. Hence the only edge *possibly* n-cleared at $t + 1$ (if not already n-clear at t) is uv . In short

$$V_N^C(t + 1) = V_N^C(t) \cup \{v\} \quad \text{and} \quad E_N^C(t + 1) = E_N^C(t) \cup \{uv\}. \quad (13)$$

The case that u is still guarded at $t + 1$ (i.e. u contained more than one searcher at time t) is omitted, since it is similar but easier to treat than the case of unguarded u .

II. Edge Game. Again, we first examine the case that only one searcher was located in u at time t . So u was guarded (and e-clear) at t and is unguarded at $t + 1$. We will now show that u remains e-clear at $t + 1$; in other words that

$$u \in V_E^C(t) \cap V_E^C(t + 1). \quad (14)$$

Suppose, on the contrary, that we have

$$u \in V_E^C(t) \cap V_E^D(t + 1); \quad (15)$$

There are only two ways for (15) to hold, each of which we examine separately.

II.1 $\exists x_1 \neq v$ such that $ux_1 \in E_E^D(t) \cap E_E^D(t + 1)$. Then $ux_1 \in E_E^D(t) = E_N^D(t)$. Since the move at $t + 1$ was $u \rightarrow v$, u was guarded at t and $x_1 \in V_N^D(t)$. Now, from $u \rightarrow v$ and $x_1 \neq v$ we get $x_1 \in V_N^D(t + 1)$; from this and u unguarded at $t + 1$ we conclude that $u \in V_N^D(t + 1)$ which contradicts monotonicity of the node search.

II.2 Alternatively, there exists a path $ux_1\dots x_{K-1}x_K$ (with $K > 1$) such that $x_Kx_{K-1} \in E_E^D(t) \cap E_E^D(t+1)$ and $ux_1\dots x_{K-1}x_K$ is e-unguarded at $t+1$. First note that, since $ux_1\dots x_K$ is e-unguarded at $t+1$ and the $(t+1)$ -th move is $u \rightarrow v$, it follows that $v \notin \{x_1, \dots, x_{K-1}\}$; and since \mathbf{T} is a tree, $v \neq x_K$ too; in short $v \notin \{x_1, \dots, x_K\}$. Also $u \notin \{x_1, \dots, x_K\}$, since $ux_1\dots x_K$ is a path. Now

$$\left. \begin{array}{l} \text{the } (t+1)\text{-th move is } u \rightarrow v \\ u \notin \{x_1, \dots, x_K\} \\ ux_1\dots x_K \text{ is e-unguarded at } t+1 \end{array} \right\} \Rightarrow ux_1\dots x_K \text{ is e-unguarded at } t$$

and this together with $x_Kx_{K-1} \in E_E^D(t)$ implies $x_{K-1} \in V_E^D(t) = V_N^D(t)$. From $v \neq x_{K-1}$ follows that x_{K-1} was not entered at $t+1$ and so $x_{K-1} \in V_N^D(t+1)$. From this and $ux_1\dots x_{K-1}$ being n-unguarded at $t+1$ we conclude that $u \in V_N^D(t+1)$ which contradicts monotonicity of the node search.

In short, u cannot become e-dirty at $t+1$, i.e., we have proved (14). Using this we now will show that no previously e-clear edge xy can become e-dirty at $t+1$. Because this would require the existence of some edge $pq \in E_E^D(t) \cap E_E^D(t+1)$ and a path $xy\dots pq$ which is e-unguarded at $t+1$ but e-guarded at t (otherwise xy would already be e-dirty at t). But any such path would have to contain u (no searcher was removed from any other node at $t+1$) which would in turn imply $u \in V_E^D(t+1)$, which contradicts (14). Since no edge is recontaminated no node is recontaminated either (the only node guarded at t and unguarded at $t+1$ is u , for which we have (14)).

On the other hand, the only edge *possibly* e-cleared at time $t+1$ (if not already e-clear at t) is uv , since no other edge was traversed at $t+1$. Hence the only node *possibly* e-cleared at time $t+1$ (if not already e-clear at t) is v .

In short we have shown

$$V_E^C(t+1) = V_E^C(t) \cup \{v\} \quad \text{and} \quad E_E^C(t+1) = E_E^C(t) \cup \{uv\}. \quad (16)$$

The treatment of the case that u is still guarded at $t+1$ (i.e., at t it contained more than one searcher) is omitted, since it is similar but easier to the case of u unguarded. From (12), (13) and (16) we obtain

$$V_E^C(t+1) = V_N^C(t+1) \quad \text{and} \quad E_E^C(t+1) = E_N^C(t+1) \quad (17)$$

and, proceeding inductively, we obtain the required result:

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) = V_N^C(t) \quad \text{and} \quad E_E^C(t) = E_N^C(t) \quad (18)$$

Monotonicity and connectedness in the edge game follow from the fact that these hold in the node game and from eq.(18). \square

Theorem 4.2. *Given a tree \mathbf{T} and an internal rooted search schedule \mathbf{S} with root u_0 , if \mathbf{S} is monotone connected in the edge game and satisfies either of the following constraints*

$$\mathbf{C1} \quad E_E^C(1) = \emptyset, \quad E_E^C(2) = \{u_0v\},$$

$$\mathbf{C2} \quad E_E^C(1) = \emptyset, \quad E_E^C(2) = \emptyset, \quad E_E^C(3) = \{u_0v\},$$

then

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) = V_N^C(t) \quad \text{and} \quad E_E^C(t) = E_N^C(t) \quad (19)$$

and \mathbf{S} is monotone connected in the node game.

Remark 4.3. Before proving the theorem, let us discuss the significance of constraints **C1** and **C2**. Three remarks must be made

1. Regarding the kind of searches which will satisfy either of the constraints: **C1** will be satisfied by searches with $\mathbf{S}(1) = 0 \rightarrow u_0$, $\mathbf{S}(2) = u_0 \rightarrow v$ (provided u_0 has a single neighbor, namely v); **C2** will be satisfied by searches with $\mathbf{S}(1) = 0 \rightarrow u_0$, $\mathbf{S}(2) = 0 \rightarrow u_0$, $\mathbf{S}(3) = u_0 \rightarrow v$; in both cases u_0 is, obviously, the root of the search.
2. The constraints are imposed to exclude situations similar to the one discussed in Remark 2.2, where an edge is e-cleared and e-dirtied during the *same* time step t .
3. Finally, **C1** and **C2** are not exceedingly restrictive; as will be seen later, every “interesting” IMC search on a tree either satisfies **C1** / **C2** or can easily be converted to an equivalent search which does.

Proof of Theorem 4.2. The theorem will be proved by induction. It is easy to check that

$$\text{for } t \in [0, t_1]: \quad V_E^C(t) = V_N^C(t) \text{ and } E_E^C(t) = E_N^C(t) \quad (20)$$

with $t_1 = 2$ when **C1** holds and $t_1 = 3$ when **C2** holds. Now suppose that at time $t \geq t_1$ we have

$$V_E^C(t) = V_N^C(t), \quad E_E^C(t) = E_N^C(t) \quad (21)$$

and consider the move performed at $t + 1$.

We first dispose of the case where the move involves the node 0. Because of internality, $v \neq 0$ (no searcher is removed from the graph). If the move is $0 \rightarrow v$, by rootedness we must have $v = u_0$. By **C1** and **C2**, u_0 is the first node cleared (at $t = 1$). In the edge game, at $t + 1$, no edge is recontaminated (by monotonicity) and no edge is e-cleared (since no edge is traversed); hence $E_E^C(t) = E_E^C(t + 1)$; since no searcher is removed from a node, $V_E^C(t) = V_E^C(t + 1)$ as well. In the node game, u_0 is already n-clear at t_1 and no other node is entered at $t + 1$, so no node is n-cleared; no searcher is removed from the graph, so no node is recontaminated; hence $V_N^C(t) = V_N^C(t + 1)$ which implies $E_N^C(t) = E_N^C(t + 1)$ as well. In short, if the $(t + 1)$ -th move involves node 0, from (21) and the above arguments follows

$$V_E^C(t + 1) = V_N^C(t + 1), \quad E_E^C(t + 1) = E_N^C(t + 1).$$

Next we examine the case where the $(t + 1)$ -th move is $u \rightarrow v$, with $u, v \neq 0$. We will examine the effects of this move separately for the node and edge game.

I. Edge Game. By monotonicity of the search, no previously e-clear edge becomes e-dirty at $t + 1$; and hence the only node that can *possibly* become e-dirty is u . We will show however that u also remains e-clear. Since $E_E^C(t_1) = \{u_0v\}$ and $t \geq t_1$, from edge monotonicity follows that $E_E^C(t) \neq \emptyset$ for all $t \geq t_1$. And, since $u \in V_N^C(t) = V_E^C(t)$ and $\mathbf{G}_E^C(t)$ is connected, there exists $uz \in E_E^C(t)$; if u becomes e-dirty at $t + 1$, then $uz \in E_E^D(t + 1)$, which contradicts monotonicity of the edge search. Hence no previously e-clear node becomes e-dirty at $t + 1$.

Regarding new clearings, edge uv and node v *may* become e-clear at $t + 1$ (if they were not already e-clear at t). No edge other than uv was traversed, hence no edge other than uv (and no node other than v) can be e-cleared at $t + 1$.

In short, in the edge game we have

$$V_E^C(t + 1) = V_E^C(t) \cup \{v\} \text{ and } E_E^C(t + 1) = E_E^C(t) \cup \{uv\}. \quad (22)$$

II. Node Game. Again, we first examine the case that only one searcher was located on u at time t . So u was guarded (hence n-clear) at t and is unguarded at $t + 1$. We will now show that u remains n-clear at $t + 1$; in other words that

$$u \in V_N^C(t) \cap V_N^C(t + 1). \quad (23)$$

Suppose on the contrary that we have

$$u \in V_N^C(t) \cap V_N^D(t+1); \quad (24)$$

(24) can happen in only two ways, each of which we examine separately.

II.1 $\exists x_1 \neq v$ such that $ux_1 \in E$ and $x_1 \in V_N^D(t) \cap V_N^D(t+1)$. Then $x_1 \in V_N^D(t) = V_E^D(t) \Rightarrow ux_1 \in E_E^D(t)$ and hence (since the move at $t+1$ was $u \rightarrow v$ and $x_1 \neq v$) $ux_1 \in E_E^D(t+1)$.

Since $E_E^C(t) \neq \emptyset$ (for all $t \geq t_1$) and $u \in V_N^C(t) = V_E^C(t)$ and $\mathbf{G}_E^C(t)$ is connected, there exists $uz \in E_E^C(t)$. However, since $ux_1 \in E_E^D(t+1)$ and u is unguarded at $t+1$, we conclude $uz \in E_E^D(t+1)$, which contradicts monotonicity of the edge search.

II.2 Alternatively, there exists a path $ux_1 \dots x_{K-1}x_K$ (with $K > 1$) such that $x_K \in V_N^D(t) \cap V_N^D(t+1)$ and $ux_1 \dots x_{K-1}x_K$ is n-unguarded at $t+1$ (so $v \notin \{u, x_1, \dots, x_K\}$). First note that, since $ux_1 \dots x_K$ is a path, $u \notin \{x_1, \dots, x_K\}$. Now

$$\left. \begin{array}{l} \text{the } (t+1)\text{-th move is } u \rightarrow v \\ u \notin \{x_1, \dots, x_K\} \\ ux_1 \dots x_K \text{ is n-unguarded at } t+1 \end{array} \right\} \Rightarrow ux_1 \dots x_K \text{ is e-unguarded at } t$$

Since $x_K \in V_N^D(t) = V_E^D(t)$ it follows that $x_{K-1}x_K \in E_E^D(t)$; since also $x_{K-1}x_K$ is not traversed at $t+1$, it follows that $x_{K-1}x_K \in E_E^D(t+1)$. By the same reasoning as in the previous case, there exists $uz \in E_E^C(t)$ and, since $ux_1 \in E_E^D(t+1)$ and u is unguarded at $t+1$, we conclude $uz \in E_E^D(t+1)$, which contradicts the edge monotonicity of \mathbf{S} .

In short, u cannot become n-dirty at $t+1$, i.e., we have proved (23). Using this fact and an argument similar to that of Theorem 4.1, we conclude that no previously n-clear node can become n-dirty at $t+1$. Since no node is recontaminated, no edge is recontaminated either.

On the other hand, the only node *possibly* n-cleared at time $t+1$ is v (no other node was entered) and hence the only edge *possibly* n-cleared at time $t+1$ is uv (if another edge was cleared, by an analysis similar to that of Theorem 4.1, we conclude that a cycle must exist in \mathbf{T} , which is impossible). Hence the only edge possibly n-cleared at $t+1$ is uv .

In short, in the node game we have

$$V_N^C(t+1) = V_N^C(t) \cup \{v\} \text{ and } E_N^C(t+1) = E_N^C(t) \cup \{uv\}. \quad (25)$$

The treatment of the case that u is still guarded at $t+1$ (i.e., it contained more than one searcher at t) is omitted, since it is similar to but easier than the case of u unguarded. From (21), (22) and (25) we obtain

$$V_E^C(t+1) = V_N^C(t+1) \text{ and } E_E^C(t+1) = E_N^C(t+1) \quad (26)$$

and, proceeding inductively, we obtain the required result

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) = V_N^C(t) \text{ and } E_E^C(t) = E_N^C(t). \quad (27)$$

Monotonicity and connectedness in the node game follow from the fact that these hold in the edge game and eq.(27). \square

We will now review some results (obtained by Barriere et al. [2, 3]) regarding IMC *edge* search on trees and will show that similar results hold for IMC *node* search on trees. In the following presentation the terminology and notation of [2, 3] is somewhat changed, to conform with the one used in the current paper.

The first basic result of [2] is the following.

Theorem 4.4. [2] For every tree \mathbf{T} there is an IMC edge clearing search schedule \mathbf{S} which uses $s_E^{ic}(\mathbf{T})$ searchers. Moreover, in \mathbf{S} all searchers are initially placed at the same node u_0 (i.e., the search schedule is rooted) and the first step (after placing the searchers) consists in clearing an edge incident to u_0 .

Corollary 4.5. For every tree \mathbf{T} we have $s_E^{ic}(\mathbf{T}) = s_E^{imc}(\mathbf{T})$. There exists a rooted IMC search schedule which achieves this bound.

Theorem 4.4 can be extended to the following sequence of inequalities[3].

Theorem 4.6. [3] For every tree \mathbf{T} we have.

$$s_E(\mathbf{T}) = s_E^i(\mathbf{T}) = s_E^m(\mathbf{T}) \leq s_E^{im}(\mathbf{T}) = s_E^c(\mathbf{T}) = s_E^{ic}(\mathbf{T}) = s_E^{mc}(\mathbf{T}) = s_E^{imc}(\mathbf{T}) \leq 2s_E(\mathbf{T}) - 2. \quad (28)$$

Furthermore, there are trees \mathbf{T} for which the inequality $s_E^m(\mathbf{T}) \leq s_E^{im}(\mathbf{T})$ is strict.

Barriere et al. also present the **Search** algorithm [2], which computes a minimal edge-clearing rooted IMC schedule \mathbf{S} for every tree \mathbf{T} . The next Theorem shows that \mathbf{S} is also a minimal node-clearing rooted IMC schedule.

Theorem 4.7. For every tree \mathbf{T} , the search \mathbf{S} produced by the **Search** algorithm is

1. a minimal edge-clearing IMC search of \mathbf{T} , which is also rooted;
2. a minimal node-clearing IMC search of \mathbf{T} , which is also rooted.

Proof. Part 1 of the Theorem is proved as Lemma 9 in [2]. Let us now prove part 2. Since \mathbf{S} is an edge-clearing strategy it is also node-clearing. We next show it is minimal.

Let us first consider the case of trees \mathbf{T} such that $s_E^{imc}(\mathbf{T}) = 1$. It is easy to see that such trees are paths of the form $u_1u_2\dots u_N$, and **Search** produces the “obvious” edge clearing schedule: $0 \rightarrow u_1$, $u_1 \rightarrow u_2$, ... (or the reverse, starting at u_N) which is also a rooted minimal node clearing IMC schedule.

Suppose now that $s_E^{imc}(\mathbf{T}) \geq 2$. By part 1, **Search** will produce a rooted minimal edge clearing IMC schedule \mathbf{S} . Also, by relaxing the requirement that all searchers are placed into the graph in the first move (in other words, by interspersing searcher placements with edge slidings) we can obtain from \mathbf{S} a new search schedule \mathbf{S}' which

1. uses the same number of searchers as \mathbf{S} (i.e., $\overline{sn}(\mathbf{S}) = \overline{sn}(\mathbf{S}')$);
2. is rooted (i.e. all searcher placements are into u_0);
3. satisfies either **C1** or **C2** of Theorem 4.2;
4. produces the same sequence of clear graphs $\mathbf{G}_N^C(t)$ as \mathbf{S} (with a slight time adjustment, to account for the changes in searcher placement times).

Hence \mathbf{S}' satisfies the conditions of Theorem 4.2 and is IMC in the edge game which means that, using \mathbf{S}' , we have

$$\text{for } t = 0, 1, 2, \dots : \quad V_E^C(t) = V_N^C(t) \quad \text{and} \quad E_E^C(t) = E_N^C(t). \quad (29)$$

From this also follows that \mathbf{S}' is a node clearing rooted IMC search schedule of \mathbf{T} .

Now take another node clearing IMC search of \mathbf{T} , call it \mathbf{S}'' , which is minimal for the node game (i.e., $\overline{sn}(\mathbf{S}'') = s_N^{imc}(\mathbf{T})$). Then, by Theorem 4.1, \mathbf{S}'' is also edge clearing, hence

$$\overline{sn}(\mathbf{S}'') \geq s_E^{imc}(\mathbf{T}) = \overline{sn}(\mathbf{S}).$$

On the other hand, since \mathbf{S}' is node clearing,

$$\overline{s_n}(\mathbf{S}'') = s_N^{imc}(\mathbf{T}) \leq \overline{s_n}(\mathbf{S}') = \overline{s_n}(\mathbf{S}).$$

Hence

$$s_N^{imc}(\mathbf{T}) = \overline{s_n}(\mathbf{S}'') = \overline{s_n}(\mathbf{S}) = s_E^{imc}(\mathbf{T})$$

which completes the proof. \square

Corollary 4.8. *For every tree \mathbf{T} we have $s_E^{imc}(\mathbf{T}) = s_N^{imc}(\mathbf{T})$.*

Remark 4.9. Theorem 4.7 shows that the **Barriere Search** algorithm can be used to compute a rooted IMC node clearing schedule for every tree \mathbf{T} . How good is such a schedule? In other words, can we node clear a tree \mathbf{T} with fewer than $s_N^{imc}(\mathbf{T})$ searchers? We can certainly do this with an internal (non-monotone, non-connected) search. For example, let $\mathbf{T}^{(2M)}$ be the complete binary tree of height $2M$; then $s_N^i(\mathbf{T}^{(2M)}) = M$, while $s_N^{imc}(\mathbf{T}^{(2M)}) = 2M$; but we do not know the values of $s_N^{im}(\mathbf{T}^{(2M)})$ and $s_N^{ic}(\mathbf{T}^{(2M)})$. More generally, we do not have the analog for node search of the inequalities (28). It is easy to see that

$$s_N(\mathbf{G}) = s_N^i(\mathbf{G}) \leq s_N^m(\mathbf{G}) \leq \frac{s_N^{im}(\mathbf{G})}{s_N^c(\mathbf{G})} = s_N^{ic}(\mathbf{G}) \leq s_E^{mc}(\mathbf{G}) = s_E^{imc}(\mathbf{G}) \quad (30)$$

for every graph \mathbf{G} and, in particular, for every tree \mathbf{T} (and the inequalities in (30) can be strict). Refining (30) to something like (28) is a subject of our future research. The question is of special interest to us because, as already mentioned, the one *absolute* requirement for robotic pursuit / evasion is internality; monotonicity and connectedness are desirable but not indispensable.

Remark 4.10. A basic component of Barriere's **Search** algorithm is the “ λ labeling” of edges, performed by the auxiliary **Label** algorithm[2]. As will be seen in the next section, the Barriere λ labels are also used by our GSST algorithm to node-clear *general graphs*.

5 Search on Graphs

We now turn to the study of rooted IMC node searches of an arbitrary graph $\mathbf{G} = (V, E)$ (which will always be assumed to have N nodes, i.e., $|V| = N$). In every such search, the t -th move has the form $\mathbf{S}(t) = u \rightarrow v$, where $u, v \in V$ or, perhaps, $u = 0$, the “source” node.

We present several variants of a basic node clearing algorithm. All variants are guaranteed to find a node clearing search schedule; we prove that two of the variants will find a minimal schedule with probability $1 - \alpha^M$ where $\alpha \in (0, 1)$ and M is the number of iterations of the algorithm; even non-minimal search schedules require a reasonably small number of searchers, as will be seen by the experiments of Section 6.

5.1 Motivation

Our basic algorithmic idea is motivated by the following rather simple observation: *every rooted IMC node-clearing search of \mathbf{G} generates a spanning tree*. This observation can be refined in the form of the following theorem.

Theorem 5.1. *Given a graph $\mathbf{G} = (V, E)$ and a rooted IMC node clearing search \mathbf{S} of \mathbf{G} . The clearing moves of \mathbf{S} generate a sequence of trees, $(\mathbf{T}_0, \mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_N)$, where (for $n = 1, 2, \dots, N$) $\mathbf{T}_n = (V_n, E_n)$ and the following hold:*

D1 \mathbf{T}_0 is the empty graph ($V_0 = \emptyset, E_0 = \emptyset$);

D2 \mathbf{T}_N is a spanning tree of \mathbf{G} ($V_N = V$, $E_N \subseteq E$);

D3 for $n = 1, 2, \dots, N$: $V_{n-1} \subseteq V_n$, $E_{n-1} \subseteq E_n$ (in other words \mathbf{T}_{n-1} is a subtree of \mathbf{T}_n);

D4 for $n = 1, 2, \dots, N$: $V_n = V_{n-1} \cup \{u_n\}$, and for $n = 2, 3, \dots$: $E_n = E_{n-1} \cup \{u_i u_n\}$, with $i \in [1, n-1]$.

Proof. Inductively. Since \mathbf{S} is monotone, it involves N clearing moves. \mathbf{T}_0 is the empty graph. \mathbf{T}_1 is formed by the first move of \mathbf{S} , which consists in placing a searcher at the root node. So \mathbf{T}_1 is the tree with a single node and trivially has \mathbf{T}_0 as a subgraph. Suppose **D3** and **D4** hold up to $m = n$ and consider the $(n+1)$ -th clearing move of \mathbf{S} . Since \mathbf{S} is connected, we add to \mathbf{T}_n one node u_{n+1} and (as explained in the Proof of Theorem 4.1) *exactly* one edge $u_i u_{n+1}$ (with $i \in [1, n]$) to obtain a new tree \mathbf{T}_{n+1} (of $n+1$ nodes and n edges) which also satisfies **D3** and **D4**. Hence **D3** and **D4** hold for $m = 1, 2, \dots, N$. At $m = N$, V_N contains N nodes, hence $V_N = V$; since \mathbf{T}_N is a tree, it is a spanning tree of \mathbf{G} . \square

Remark 5.2. An abbreviated statement of Theorem 5.1 could be: “Every rooted IMC node clearing search of \mathbf{G} specifies a spanning tree of \mathbf{G} and an order of clearing the nodes”. But the order must be consistent with the edge structure of the graph and the spanning tree (e.g., we cannot use an order of node clearing which requires non-existent edges). This consistence is exactly what conditions **D3** and **D4** describe.

Theorem 5.1 gives the motivation for our algorithm GSST (*Guaranteed Search with Spanning Trees*), which is informally described below.

Algorithm 1 GSST, Informal Description

Input: Graph \mathbf{G}

Select a spanning tree \mathbf{T} of \mathbf{G} and a root u_0 of \mathbf{T}

Find a rooted IMC node-clearing search \mathbf{S}' of \mathbf{T}

Apply \mathbf{S} to \mathbf{G}

if at some step of \mathbf{S}' a move $u \rightarrow v$ would result in recontamination in \mathbf{G} **then**

send a “guard” to u

execute $u \rightarrow v$

end if

Output: The search \mathbf{S} obtained by combining \mathbf{S}' with the guard moves

Note that, by construction, all clearing moves take place along edges of the spanning tree \mathbf{T} .

In the above description we have used the term “guard”. Stated informally, the searchers play two roles: “*tree searchers*” perform the clearing moves, always along the edges of the spanning tree; “*guards*” are stationary and block potential recontamination paths. However, note that a particular searcher can change roles during the course of the search.

The main advantage of the GSST algorithm is that it is *fast*. A random spanning tree \mathbf{T} can be quickly generated and searched. Since \mathbf{T} and \mathbf{G} have the same node set, node-clearing \mathbf{T} in a *node-monotone manner* results in node-clearing \mathbf{G} as well. The main issue is: *how many guards will be required to block recontamination through non-tree edges of G ?* The number usually turns out to be quite reasonable, because (a) guards can be reused and (b) tree searchers can also be used as guards when they do not perform clearing moves.

Both the labeling and the traversal phase of GSST can be executed in either a centralized or distributed manner; the latter is useful for robotic applications, where each robot can share some of the computational load. In the distributed implementation, all searchers share the underlying spanning tree and labeling. When a searcher reaches a node, he checks to see if he can move without recontamination.

If he can, he determines his next move based on the traversal strategy, and he shares this move with the team.

Hence GSST has short execution time and can be run repeatedly (in reasonable time), using many different (randomly selected) spanning trees. Our algorithm depends on the quick discovery of a spanning tree corresponding to a minimal or near-minimal search (Theorem 5.1). Our experiments in Section 6 show that GSST discovers near-minimal searches in only a short time for several families of complex graphs.

An important characteristic of the GSST algorithm is its “anytime-ness”. Anytime algorithms return a partial answer before completion and they keep providing improved answers, the improvement increasing with computation time [40]. GSST has these characteristics, as will be understood in Section 5.2. Namely, GSST is characterized by monotonicity (the solution only improves over time), recognizable quality (the quality of the solution, i.e. number of searchers, can be determined at run time), consistency (the algorithm will not spend too much time finding a single solution), and interruptibility.

5.2 The GSST Algorithms

A detailed description of the GSST algorithm is given by the following pseudocode on p.19. The notation $\mathbf{S} = \mathbf{S} | (u \rightarrow v)$ means that the move $(u \rightarrow v)$ is appended to the previously determined search schedule \mathbf{S} (i.e., becomes the next move of \mathbf{S}). Several subroutines appearing in the following listing will be discussed presently.

The following remarks explain the operation of the algorithm.

1. Generate a random spanning tree \mathbf{T} by **GenerateTree**. We have used two different methods of random spanning tree generation.
 - (a) The *uniform* method is an implementation of Wilson’s [37] *loop erased random walk algorithm*.
 - (b) The *DFS* method selects a root node and randomly moves down the tree in a depth-first manner. At each node, a random incident edge is chosen and set as an edge in the spanning tree. A visited list is maintained, and when a node is visited more than once, the edge used to reach it the second time is set as a non-tree edge. This eliminates cycles in the graph and thus generates a tree. When a leaf is reached, the algorithm recurses to ensure that all nodes are included in the tree (i.e, it is a spanning tree of the original graph). The motivation for this method is to bias towards spanning trees that require fewer guards. The intuition is that a DFS traversal will generate only a few nodes with non-tree edges, thus leading to few required guards.
2. Label the edges of \mathbf{T} by the **R-Label** algorithm (this, a modification of Barriere’s **Label** algorithm, is listed and discussed in Appendix A).
3. While n-dirty nodes still exist, select an n-dirty edge uv of \mathbf{T} (by **SelectEdge**, to be discussed presently). Let a searcher traverse uv if this does not cause node recontamination⁵; if no such searcher exists, then use a new searcher (originally placed at the root) to traverse uv .
4. Repeat the process until all nodes are n-cleared. (Hence the algorithm does not allow node recontamination and will generate as many searchers as necessary to prevent it).
5. When all nodes have been n-cleared, a rooted IMC node clearing search has been generated, which performs all its clearing moves along the edges of \mathbf{T} .

⁵The line “Move that searcher to u staying inside the clear graph” is actually a simplification, i.e., in the interest of brevity, we do not indicate how such a path is obtained; however this is always possible, usually in more than one ways.

Algorithm 2 GSST

Input: \mathbf{G} : a graph; M : no. of spanning trees to use.

$\mathbf{S}_{min} = \emptyset$

$s_{min} = \infty$

for $m = 1 : M$ **do**

$\mathbf{S} = \emptyset$

$\mathbf{T} = \mathbf{GenerateTree}(\mathbf{G})$

 Randomly choose root u_0

$\mathbf{S} = \mathbf{S} \mid (0 \rightarrow u_0)$

$V_N^C = \{u_0\}, E_N^C = \emptyset$

$V_N^D = V - \{u_0\}, E_N^D = E$

$\mathbf{G}_N^C = (V_N^C, E_N^C)$

$\lambda = \mathbf{R-Label}(\mathbf{T})$

while $V_N^D \neq \emptyset$ **do**

$uv = \mathbf{SelectEdge}(\mathbf{G}_N^C, \mathbf{T}, E_N^D, \lambda)$

if a searcher can traverse uv without recontamination **then**

 Move that searcher to u *staying inside the clear graph*

$\mathbf{S} = \mathbf{S} \mid (u \rightarrow v)$

$V_C = V_N^C \cup \{v\}, E_N^C = E_C \cup \{uv\}$

$V_N^D = V_N^D - \{v\}, E_N^D = E_N^D - \{uv\}$

$\mathbf{G}_N^C = (V_N^C, E_N^C)$

else

$\mathbf{S} = \mathbf{S} \mid (0 \rightarrow u_0)$

end if

end while

if $\overline{sn}(\mathbf{S}) < s_{min}$ **then**

$\mathbf{S}_{min} = \mathbf{S}$

$s_{min} = \overline{sn}(\mathbf{S})$

end if

end for

Output: Node clearing schedule \mathbf{S}_{min} .

6. Go back to step 1 and repeat the process with a new spanning tree.

7. After the maximum number of trees and (corresponding searches) has been generated, return a search \mathbf{S}_{min} which attains s_{min} , the minimum value of $\overline{sn}(\mathbf{S})$.

The subroutine **SelectEdge** chooses an n-dirty edge uv *belonging to* \mathbf{T} and adjacent to the current clear graph \mathbf{G}_N^C . There are several ways to perform this selection.

1. **Labeled Selection (L)**. Select the next edge uv of \mathbf{T} to be traversed according to the Barriere λ labels (as in algorithm **R-Search**, presented in Appendix A); however, if traversing uv would cause recontamination then select the next edge in the Barriere sequence; if, at some stage of the search, traversing *any* edge of \mathbf{T} would cause recontamination (i.e., if all searchers are stuck) then introduce a new searcher at the root.
2. **Labeled Selection with Randomized Tie-breaking (LR)**. Same as the previous except that ties of edge labels are broken randomly.

3. **Randomized Selection (R)**. Choose uv randomly (without making *any* use of the *Barriere* labeling) by a uniform probability on all n-dirty edges *belonging to* \mathbf{T} and adjacent to \mathbf{G}_N^C :

$$\Pr(uv) = \begin{cases} c & \text{if } u \text{ is n-clear and } uv \text{ is n-dirty and an edge of } \mathbf{T}; \\ 0 & \text{else.} \end{cases}$$

4. **Labeled Weighted Selection (LW)**. This rule is intermediate between R and L: edge selection is still random but, instead of a uniform probability distribution, the probability of an edge uv being selected is inversely proportional to its *Barriere* λ label.
5. **Label Dominated Selection (LD)**. This can be done by labeling edges that lead to parts of the graph that are trees (subtrees of the graph). A list of searchers who can move without recontamination can be maintained during search. If an edge adjacent to $V_N^C(t)$ leads to a subtree of the graph, and enough free searchers are available, clearing this subtree can only improve the search strategies.

By using each of the above rules in the “basic” GSST algorithm, we obtain *ten* GSST variants: uniform GSST-L, uniform GSST-LR, ... , uniform GSST-LD, DFS GSST-L, ..., DFS GSST-LD. These variants (except for the two GSST-R’s) utilize the *Barriere* labeling, *but do not necessarily produce a Barriere traversal of the spanning tree*⁶. In a sense Randomized Selection is the simplest or most naive rule that can be used to select the next move of the search schedule: every n-dirty edge of \mathbf{T} (adjacent to the clear graph) is equally likely to be selected. The remaining three rules can be understood as ways to bias the probability by which edges are selected in some meaningful way. The effectiveness of these rules will be judged by the experiments of Section 6. From the theoretical point of view, we will show in Section 5.3 that the uniform GSST-R and uniform GSST-LD will find a minimal node clearing schedule with probability $1 - \alpha^M$ where M is the number of iterations and $\alpha \in (0, 1)$; we conjecture that this property does not hold for *Barriere* selection. As a practical matter, the issue is how large M has to be for $1 - \alpha^M$ to be sufficiently close to 1. However, the experiments of Section 6 indicate that the above rules find good search schedules in very reasonable time.

Variants of GSST can also be produced by replacing the **GenerateTree** subroutine with an *exhaustive generation* of all spanning trees of \mathbf{G} (it can be used in conjunction with any of the above variants of **SelectEdge**). To do this we have used Char’s spanning tree enumeration algorithm [21]. Exhaustive enumeration is feasible only for relatively small graphs.

Finally note that the search schedules produced by (every variant of) GSST are IMC. This holds for the search of both \mathbf{T} and \mathbf{G} . Indeed, for the search of \mathbf{G} to work, the search of \mathbf{T} *must* be IMC. In other words, no obvious modification of GSST will produce, for example, an internal, connected, *non-monotone* node-clearing search of \mathbf{G} . The question arises: how good is a minimal IMC node clearing of \mathbf{G} (as compared to, for example, an internal but not monotone / connected node clearing)? This is the question already hinted at in Remark 4.9.

Let us close this section by repeating that *the basic idea of GSST is to perform all clearing moves along the edges of a spanning tree*. This idea exploits the facts that (a) spanning trees can be both generated and searched quickly and (b) blocking recontamination does not require an excessively large number of guards (because a searcher can change roles as a guard and a tree searcher).

5.3 Completeness

Definition 5.3. *Given a graph \mathbf{G} and a search \mathbf{S} of \mathbf{G} , the frontier at t (under \mathbf{S}) is*

$$V_N^F(t) = \{u : u \in V_N^C(t) \text{ and } \exists v : v \in V_N^D(t), uv \in E\},$$

⁶This is true even of GSST-L, because an edge which would be next in *Barriere*’s traversal order may be temporarily skipped if its traversal (at the current stage of the search) would cause recontamination.

i.e., the n -clear nodes which are connected to n -dirty nodes.

Lemma 5.4. *In a node search, for every t , the nodes $u \in V_N^F(t)$ are guarded.*

Proof. A frontier node u is by definition n -clear and adjacent to an n -dirty node v . This is only possible if u is guarded. \square

Definition 5.5. *Consider a rooted IMC node clearing search \mathbf{S} of \mathbf{G} . Let $t_0 = 0$ and suppose the clearing moves of \mathbf{S} take place at times t_1, \dots, t_N ; let also $t_0 = 0$. The n -th phase of \mathbf{S} (for $n = 1, 2, \dots, N$) is the time interval $[t_{n-1} + 1, t_n]$, i.e. the interval between the $(n - 1)$ -th and n -th clearing move.*

The following remarks are rather obvious. For $m = 2, 3, \dots$, a target edge $u_i u_m$ (with $i \in [1, m - 1]$) corresponds to the m -th phase (here we take the root node to be u_1). While $t \in [t_{m-1} + 1, t_m]$ the algorithm moves a searcher towards $u_i u_m$. For $t \in [t_{m-1} + 1, t_m - 1]$, $u_i u_m$ is n -dirty, u_i is n -clear, u_m is n -dirty. At $t = t_m$ we have $\mathbf{S}(t_m) = u_i \rightarrow u_m$ and $u_i u_m, u_m$ are n -cleared.

Lemma 5.6. *Given a graph $\mathbf{G} = (V, E)$ and a rooted IMC node clearing search \mathbf{S} of \mathbf{G} , produced by either GSST-R or GSST-LW, let t_1, t_2, \dots, t_N be the times at which clearing moves take place; let also $t_0 = 0$. Then for $n = 2, \dots, N$:*

1. for $t \in [t_{n-1}, t_n - 1]$: $V_N^F(t) = V_N^F(t_{n-1})$;
2. for $t \in [t_{n-1}, t_n - 1]$: every $u \in V_N^F(t)$ contains exactly one searcher, except one node $u(t)$ which possibly contains two searchers;
3. for $t = t_n$: every $u \in V_N^F(t_n)$ contains exactly one searcher.

Proof. The proof is by induction on n . Items 1, 2, 3 of the theorem hold for $n = 1$, $t \in [t_0, t_1] = \{0, 1\}$. Suppose they also hold up to the m -th phase. In the interval $[t_m + 1, t_{m+1}]$ a target edge $u_i u_{m+1}$ (with $i \in [1, m]$) is selected and an available searcher is sent towards $u_i u_{m+1}$. Because GSST-R / GSST-LW always selects for clearing an edge adjacent to the clear graph, $u_i \in V_N^F(t_m)$, $u_{m+1} \in V_N^D(t_m)$. There are three cases.

I. u_i is neighbor of a single n -dirty node, namely u_{m+1} . In this case $t_{m+1} = t_m + 1$ (i.e., u_{m+1} is n -cleared in one step) and $[t_m, t_{m+1}] = \{t_m, t_m + 1\}$. Node $u_i \notin V_N^F(t_{m+1})$; node u_{m+1} may or may not belong to $V_N^F(t_{m+1})$ but, at any rate, it contains exactly one searcher; no searchers enter or exit any other nodes, hence (by the inductive hypothesis) all $u \in V_N^F(t_{m+1})$ contain exactly one searcher.

II. u_i is neighbor of more than one n -dirty nodes, one of which is u_{m+1} , and there are free searchers. Since u_i is a frontier node, by hypothesis it contains a single searcher who, consequently, is stuck. However, if free searchers are available inside n -clear, non-frontier nodes, one of these searchers will be sent to $u_i u_{m+1}$ by a sequence of moves. At every t during this sequence, the searcher may enter a frontier node w ; for that particular t , w will be the only frontier node which contains two searchers. At $t = t_{m+1} - 1$ the searcher will be located at u_i (which will now contain two searchers) and at t_{m+1} he will enter u_{m+1} ; this leaves at t_{m+1} every $u \in V_N^F(t_m)$ with a single searcher and also places a searcher at u_{m+1} . The only node which may have been added to the frontier is u_{m+1} which contains a single searcher; every other node $u \in V_N^F(t + 1)$ was previously in the frontier, contained a single searcher and, if a searcher entered u at some $t \in [t_m + 1, t_{m+1} - 1]$, it exited u at $t + 1$; hence every such node at t_{m+1} contains a single searcher.

III. The final case is when all frontier-located searchers are stuck and there are no searchers inside n -clear, non-frontier nodes. In this case a new searcher is placed at the root node and the rest of the analysis is identical to that of Case II. Hence items 1, 2, 3 also hold for $[t_m + 1, t_{m+1}]$ and we can complete the induction for $n = 1, 2, \dots, N$. \square

Lemma 5.7. *Given a graph $\mathbf{G} = (V, E)$ and a tree sequence $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$ which satisfies conditions **D1-D4** of Theorem 5.1. Then uniform GSST-R / GSST-LW with $M = 1$ (i.e., using a single spanning tree) has a nonzero probability of producing a search \mathbf{S} which generates $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$.*

Proof. The probability that Algorithm 1 generates the tree sequence $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$ is

$$\Pr(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N) = \left[\prod_{n=1}^N \Pr(\mathbf{T}_n | \mathbf{T}_N, \mathbf{T}_0, \dots, \mathbf{T}_{n-1}) \right] \Pr(\mathbf{T}_0 | \mathbf{T}_N) \Pr(\mathbf{T}_N).$$

Note that the conditioning in the above expression *always* includes \mathbf{T}_N , since this is the first choice made in running GSST-R / GSST-LW. Now obviously, $\Pr(\mathbf{T}_0 | \mathbf{T}_N) = 1$. By Wilson's Theorem 1 [37], $\Pr(\mathbf{T}_N) > 0$ for every spanning tree \mathbf{T}_N . Also, $\Pr(\mathbf{T}_n | \mathbf{T}_N, \mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_{n-1})$ is the probability of expanding (at the n -th step) \mathbf{T}_{n-1} by the edge $u_i u_n \in \mathbf{E}_n - \mathbf{E}_{n-1}$ which, by the construction of both GSST-R and GSST-LW, is always positive. Finally, $\Pr(\mathbf{T}_N | \mathbf{T}_N, \mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_{N-1}) = 1$. Hence $\Pr(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N) > 0$ for every sequence $\mathbf{T}_1, \dots, \mathbf{T}_N$. \square

Lemma 5.8. *Given a graph $\mathbf{G} = (V, E)$ and a rooted IMC node clearing search \mathbf{S} of \mathbf{G} ; let $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$ be the tree sequence generated by \mathbf{S} . Let \mathbf{S}' be a search produced by either GSST-R or GSST-LW and also generating $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$. Then $\overline{sn}(\mathbf{S}) \geq \overline{sn}(\mathbf{S}')$.*

Proof. The proof is exactly the same for GSST-R and GSST-LW, so we only prove the first one, by induction. Let t_1, \dots, t_N be the clearing times of \mathbf{S} and t'_1, \dots, t'_N be the clearing times of \mathbf{S}' . Also let $t_0 = t'_0 = 0$.

At $t_0 = t'_0 = 0$ we have $sn(\mathbf{S}', 0) = sn(\mathbf{S}, 0) = 0$.

The only times at which $sn(\mathbf{S}', t)$ may change are $1, t'_1 + 1, \dots, t'_{N-1} + 1$. Suppose that

$$sn(\mathbf{S}, t_n) \geq sn(\mathbf{S}', t'_n).$$

Further, suppose that at $t'_n + 1$ a new searcher is introduced in \mathbf{S}' . This can only happen (in the \mathbf{S}' search) if all of the following hold:

1. at t'_n exactly $|V_N^F(t)|$ searchers exist in \mathbf{G} ;
2. there are no searchers inside nodes $u \in V_N^C(t'_n) - V_N^F(t'_n)$ (i.e., all searchers are located inside frontier nodes);
3. all searchers are stuck (i.e., moving a searcher out of a frontier node u exposes u to recontamination).

The sequence $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$ along with the clearing times determines the frontier $V_N^F(t)$ for every t . Hence \mathbf{S}' at t'_n has the same frontier as \mathbf{S} at t_n . If conditions 1-3 above hold in \mathbf{S}' , then every searcher is located in a frontier node and is stuck. It is possible that non-stuck searchers exist in \mathbf{S} (located either in frontier or non-frontier nodes) but this also means that $sn(\mathbf{S}, t_n) \geq sn(\mathbf{S}', t'_n) + 1$; hence adding a searcher in \mathbf{S}' at $t'_n + 1$ preserves

$$sn(\mathbf{S}, t_n) \geq sn(\mathbf{S}', t'_n + 1).$$

Since no searchers are added in \mathbf{S}' for $t \in [t'_n + 2, t'_{n+1}]$ and no searchers are ever removed in \mathbf{S} (i.e., $sn(\mathbf{S}, t_{n+1}) \geq sn(\mathbf{S}, t_n)$) we also get

$$sn(\mathbf{S}, t_{n+1}) \geq sn(\mathbf{S}', t'_{n+1}).$$

From the above inequality inductively we get $sn(\mathbf{S}, t_N) \geq sn(\mathbf{S}', t'_N)$ which proves the Lemma. \square

Theorem 5.9. *Given a graph $\mathbf{G} = (V, E)$.*

1. *Uniform GSST-R will generate a minimal rooted IMC clearing node search of \mathbf{G} with probability greater than or equal to $1 - \alpha_1^M$ where M is the number of iterations and $\alpha_1 \in (0, 1)$.*
2. *Uniform GSST-LW will generate a minimal rooted IMC clearing node search of \mathbf{G} with probability greater than or equal to $1 - \alpha_2^M$ where M is the number of iterations and $\alpha_2 \in (0, 1)$.*

Proof. The proof is exactly the same for GSST-R and GSST-LW, so we only prove the first one. \mathbf{G} has at least one minimal rooted IMC node clearing search \mathbf{S} of \mathbf{G} . Let $(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_N)$ be the tree sequence generated by \mathbf{S} . By Lemma 5.7, GSST-R has a nonzero probability, call it β_1 , of generating in a single iteration a search \mathbf{S}' with the same tree sequence as \mathbf{S} . Then, by Lemma 5.8,

$$\overline{sn}(\mathbf{S}) \geq \overline{sn}(\mathbf{S}').$$

Since \mathbf{S} is minimal, $\overline{sn}(\mathbf{S}) = \overline{sn}(\mathbf{S}')$ and so \mathbf{S}' is minimal too. Now, the probability of *not* generating \mathbf{S}' in a single iteration is $\alpha_1 = 1 - \beta_1$; and the probability of *not* generating \mathbf{S}' in M iterations is $\alpha_1^M = (1 - \beta_1)^M$, while the probability of generating \mathbf{S}' in M iterations is $1 - \alpha_1^M$. \square

Finally, let us mention that the GSST algorithm can be modified to produce an edge- rather than node-clearing search using the following theorem.

Theorem 5.10. *Given a graph \mathbf{G} and an IMC node clearing search \mathbf{S} using K searchers, there is an edge clearing search \mathbf{S}' using either K or $K + 1$ searchers.*

Proof. Suppose that \mathbf{G} contains L_0 edges, that the length (i.e., number of moves) of \mathbf{S} is t_{fin} and that $\overline{sn}(\mathbf{S}) = K$. The new search \mathbf{S}' will consist of the \mathbf{S} moves (executed at integer times $t = 1, 2, \dots, t_{fin}$) combined with the moves of an extra searcher, the “edge cleaner”, who will only (if at all) move at *fractional* time steps of the form $t + \frac{l}{3L_0+1}$, $l = 1, 2, \dots, 3L_0$ (fractional times are introduced to preserve the “alignment” of \mathbf{S} and \mathbf{S}' , i.e. to ensure $\mathbf{S}(t) = \mathbf{S}'(t)$ at integer times; of course, once \mathbf{S}' has been obtained, the time scale can be renormalized, so that *all* moves occur at integer times). We will use the notation $t^- = t - \frac{1}{3L_0+1}$.

We will describe the moves of of the edge cleaner on a step-by-step basis, for $t = 1, 2, \dots, t_{fin}$ in such a manner that at the same time we will complete an inductive proof of the fact that

$$\text{for } t = 1, 2, \dots, t_{fin}+1 : \quad V_N^C(t-1) = V_E^C(t-1), \quad V_N^F(t-1) = V_E^F(t-1), \quad E_N^C(t^-) = E_E^C(t^-);$$

recall that V_N^F is the frontier, i.e., the set of n-clear nodes connected to n-dirty nodes (and similarly for V_E^F) and that the frontier nodes are always guarded in both the node and edge game.

For $t = 1$ we have $V_N^C(0) = V_E^C(0) = \emptyset$, $V_N^F(0) = V_E^F(0) = \emptyset$, and the edge cleaner is not used, so $E_N^C(1^-) = E_N^C(0) = E_E^C(0) = E_E^C(1^-)$. Now suppose that

$$\text{for } s = 1, 2, \dots, t \quad V_N^C(t-1) = V_E^C(t-1), \quad V_N^F(t-1) = V_E^F(t-1), \quad E_N^C(t^-) = E_E^C(t^-) \quad (31)$$

and consider $s = t + 1$. Let the $(t + 1)$ -th move of \mathbf{S} be $u \rightarrow v$. We consider three cases.

1. **Case I.** If u is an interior node (i.e., a non-frontier node: $u \in V_N^C(t-1) - V_N^F(t-1) = V_E^C(t-1) - V_E^F(t-1)$), then uv is an edge of $\mathbf{G}_N^C(t) = \mathbf{G}_E^C(t)$ and no new nodes/edges are

cleared, either in the node or edge game. No path becomes unguarded, and so no recontamination is possible either. Hence

$$\begin{aligned} V_N^C(t) &= V_N^C(t-1) = V_E^C(t-1) = V_E^C(t), \\ V_N^F(t) &= V_N^F(t-1) = V_E^F(t-1) = V_E^F(t). \end{aligned} \quad (32)$$

The edge cleaner is not used, hence also

$$E_N^E((t+1)^-) = E_N^E(t^-) = E_E^E(t^-) = E_N^E((t+1)^-). \quad (33)$$

2. **Case II.** If u is a frontier node ($u \in V_N^F(t-1) = V_E^F(t-1)$) which contains a *single* searcher at $t-1$, then u becomes unguarded at t . Clearly

$$V_N^C(t) = V_N^C(t-1) \cup \{v\}. \quad (34)$$

It is also easy to see that there is no edge uy with $y \in V_N^D(t-1) = V_E^D(t-1)$ and $y \neq v$: if such an edge existed, then we would have $y \in V_N^D(t)$ (y was not entered at t) and so $u \in V_N^D(t)$; but $u \in V_N^D(t-1)$ (it was guarded at $t-1$) and so node monotonicity of \mathbf{S} would be violated.

Take any edge $pq \in E_E^C(t^-) = E_N^C(t^-)$. Edge pq cannot be recontaminated in the node game (no node was recontaminated). Edge pq cannot be recontaminated in the edge game either; for this to happen there must exist an e-unguarded path from pq to some $xz \in E_E^D(t^-)$; but all such paths must go through u (no other node became unguarded at t) and hence must include uv (the only edge e-dirty at t^- and incident on u) but v is guarded at t . Hence no edge is e-dirtied at t and so no node is e-dirtied either. In other words

$$V_E^C(t) = V_E^C(t-1) \cup \{v\}. \quad (35)$$

which, together with (34) shows that $V_N^C(t) = V_E^C(t)$. Also, $V_N^F(t) = V_E^F(t)$; in both the node and edge game, u was removed from the frontier and v was *possibly* added to it. In short, we have established (32) for Case II as well.

From the previous remarks we know that, in the edge game and at time t , no edge was recontaminated and edge uv was e-cleared. I.e.,

$$E_E^C(t) = E_E^C(t^-) \cup \{vu\}.$$

In the node game, on the other hand,

$$E_N^C(t) = E_N^C(t^-) \cup \{vx_0, vx_1, \dots, vx_L\} = E_N^C((t+1)^-) \quad (36)$$

where $L \geq 0$, $x_0 = u$ and there *may* exist other nodes x_1, x_2, \dots, x_L which must (a) be neighbors of v and (b) belong to $V_N^F(t-1) \subseteq V_N^C(t-1)$. Hence x_1, \dots, x_L were guarded at $t-1$ and remain so at t . In other words, at time t the edges vx_1, \dots, vx_L have both endpoints guarded and are n-clear but e-dirty. Now we invoke the edge cleaner who (at times $t + \frac{1}{3L_0+1}, t + \frac{2}{3L_0+1}, \dots$) moves to v and then performs the moves $v \rightarrow x_1, x_1 \rightarrow v, v \rightarrow x_2, x_2 \rightarrow v, \dots, x_L \rightarrow v$. This entire sequence can be performed in no more than $3L_0$ moves, so at $t = (t+1)^-$ the edges vx_0, vx_1, \dots, vx_L have been e-cleared and so

$$E_E^C((t+1)^-) = E_N^C(t^-) \cup \{vx_0, vx_1, \dots, vx_L\}. \quad (37)$$

Combining (36) and (37) we get $E_N^C((t+1)^-) = E_E^C((t+1)^-)$. In short, we have established (33) for Case II as well.

3. **Case III.** The final case to examine is when u is a frontier node which contains more than one searcher at $t - 1$. We omit a detailed treatment because the proof combines elements from the previous two cases; namely, recontamination does not happen (for the same reasons as in Case I) but the edge cleaner may possibly be required (as in Case II).

Hence, in all three cases considered, starting from (31) we have established (32) and (33). Hence we can complete the induction up to time $t = t_{fin} + 1$ which means

$$V_N^C(t_{fin}) = V_E^C(t_{fin}), \quad E_N^C((t_{fin} + 1)^-) = E_E^C((t_{fin} + 1)^-). \quad (38)$$

Eqs.(38) imply that all nodes are e-cleared at t_{fin} but a few extra steps may be required to e-clear all edges (by time $(t_{fin} + 1)^-$ at most). \square

6 Graph Search Experiments

In this section we evaluate the performance of the GSST algorithm by numerical experiments. Some of these experiments involve specific graphs (Section 6.1) and others involve families of graphs (Sections 6.2 and 6.3 – in which case we present average results). We use ten *variants* of GSST, obtained by using two different methods of spanning tree generation (uniform and DFS) and five methods of edge traversal (GSST-L, GSST-R, GSST-LR, GSST-LW, and GSST-LD).

6.1 Experiments using Individual Graphs

6.1.1 Simple Graph

The first graph we have used appears in Fig.2. This is a relatively simple graph (with IMC node search number $s_N^{imc} = 3$) which we use to illustrate the basic principles of GSST operation.




Figure 2: A simple graph.

We node-search the graph using the ten GSST variants, each with $M = 5000$ generated spanning trees; to each spanning tree corresponds a search \mathbf{S}_m , $m = 1, 2, \dots, 5000$. In Table 2 we list for each GSST variant: (a) $\min_{m=1,2,\dots,M} \bar{s}_n(\mathbf{S}_m)$ (i.e., the minimum number of searchers achieved by the

specific combination), (b) the proportion of minimal searches out of the the M total searches and (c) the time⁷ (in sec) required to run the M searches.

	Uniform ST generation			DFS ST generation		
Edge Traversal	Min	Prop. of min	Time	Min	Prop. of min	Time
GSST-L	3	0.1562	0.203125	3	0.1804	0.203125
GSST-R	3	0.2031	0.171875	3	0.2848	0.203125
GSST-LR	3	0.2656	0.265625	3	0.2096	0.265625
GSST-LW	3	0.2292	0.218750	3	0.2854	0.250000
GSST-LD	3	0.2031	0.218750	3	0.3020	0.234375

Table 2. Node-clearing the “simple graph” by the various GSST variants: minimum number of searchers attained and proportion of minimal solutions; number of spanning trees generated is $M = 5 \cdot 10^3$.

The true node search number (i.e., 3) has been found by every variant of GSST. Generally, the DFS variants perform better than the uniform ones, as can be seen by the higher proportion of minimal solutions achieved. A better understanding of the distribution of the number of searchers required by each search can be obtained by looking at the *histogram* of the distribution; one such histogram (for the variant with BH edge traversal and uniform spanning tree generation) is plotted in Fig.3. We can see that this simple graph has a high proportion of spanning trees which yield minimal schedules.




Figure 3: The histogram of the distribution of minimum number of searchers required to node clear the “simple graph” by the uniform GSST-L variant.

An additional tool to evaluate the anytime performance of GSST is the plot of $R(m) = \min_{i \leq m} \overline{sn}(\mathbf{S}_i)$ (the minimum node clearing number of searchers achieved by the first m searches). $R(m)$ is decreasing with m . The overall minimum achieved by GSST is $R(M)$ (having tried M searches). If this minimum is achieved for a small value of m , then the minimal solution has been obtained quickly. A graph of $R(m)$ (for the variant with labeled edge traversal and uniform spanning tree generation) appears in Fig.4; as can be seen a minimal solution (clearing the graph with three searchers) is achieved by the seventh computed search, approximately at time $t = 0.203125 \cdot 7/5000 = 2.8438 \times 10^{-4}$ sec.

The “simple graph” is simple enough to have a relatively small number of spanning trees (namely 272, as computed by Kirchoff’s theorem [18]) and so we can also run GSST on *all* spanning trees

⁷All computations were performed by the gsearch.exe program (see Appendix D), on a PC with Intel Dual Core E7500 CPU, running at 2.93 GHz with 3 GB RAM; computation time is reported in seconds, unless otherwise indicated.

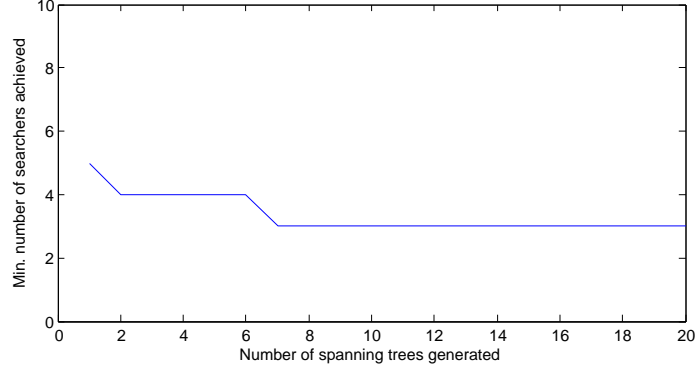


Figure 4: Plot of $R(m)$ vs. m for the “simple graph” and the uniform GSST-L variant. The m axis is truncated at $m = 20$, since the rest of the plot shows no change.

(exhaustive enumeration). The computation takes 0.03125 (for the GSST-L variant) sec and shows that 83 out of the 272 spanning trees (*rooted at node 1*) yield minimal searches⁸.

6.1.2 Tree/Grid Graph

The second graph we have used appears in Fig.5. It consists of a “root” node and two branches under it; the left branch is a tree and the right one a grid; hence the name “tree/grid”. The graph has $s_N^{imc} = 4$ and we have deliberately designed it to “trick” the GSST algorithm. For the sake of definiteness consider GSST-L. If the root of the search is node 6, then GSST-L will find a four searcher node clearing IMC schedule. However, if the root is node 1, then GSST-L will only find a five searcher schedule, even after enumerating all spanning trees. A four searcher IMC node clearing schedule is possible from either starting node; but it requires the use of a non-Barrier edge traversal (for example, one produced by the GSST-R variant); but GSST-L will always first send the searchers down the right branch (towards the grid); actually, going first to the left branch, towards the tree, is better (i.e. yields a four searcher schedule).

Table 3 summarizes the results of our experiment, using the ten variants of GSST and $M = 5 \cdot 10^4$ spanning trees per variant. Note that in this table (as in all others) the starting node is selected randomly. The uniform variants are able to generate minimal search schedules, but not the DFS ones.

Edge Traversal	Uniform ST generation			DFS ST generation		
	Min	Prop. of min	Time	Min	Prop. of min	Time
GSST-L	4	0.00054	5.921875	5	0.06380	6.687500
GSST-R	4	0.00328	3.734375	5	0.06956	4.390625
GSST-LR	4	0.00076	8.328125	5	0.06976	8.765625
GSST-LW	4	0.00206	4.812500	5	0.07252	5.609375
GSST-LD	4	0.00232	4.968750	5	0.09376	5.656250

Table 3. Node-clearing the “tree/grid” by the various GSST variants: minimum number of searchers attained and proportion of minimal solutions; number of spanning trees generated is $M = 5 \cdot 10^4$.

⁸Note that all GSST variants, except GSST-L, use a randomized element in the order of edge traversal – hence multiple runs with the same spanning tree can yield different searches.




Figure 5: The “tree/grid” graph.

In this case the uniform variants find four-searcher clearing schedules and hence outperform the DFS generated ones, which can only clear the graph with five or more searchers⁹. Note also the lower proportion of minimal solutions, even for the uniform variants. For example, uniform GSST-L finds $0.00054 \cdot 50000 = 27$ minimal solutions. This is an indication that this graph is indeed harder than the “simple” one, at least for the GSST algorithm. Also, the tree/grid graph has 31529 spanning trees (and correspondingly many cycles) which is an additional indication that it is (much) harder to search than the simple graph, which has 272 spanning trees. The small number of minimal solutions can also be appreciated by looking at the histogram (for uniform GSST-L variant it is plotted in Fig.6).




Figure 6: The histogram of the distribution of minimum number of searchers required to node clear the “hard graph” by the uniform GSST-L variant. There is a small, barely noticeable, bar at $k = 4$, corresponding to the four-searcher clearing schedules.

However, despite the small proportion of minimal solutions, the first one is always found after a relatively small number of iterations. For example, the uniform GSST-L variant finds a four searcher schedule with the 420-th spanning tree generated (out of a total of $5 \cdot 10^4$ spanning trees), approximately at time $t = 5.921875 \cdot 420/50000 = 4.9744 \times 10^{-2}$ sec.

⁹We have experimented with larger values of M (e.g., $M = 10^6$) but the DFS variants still cannot achieve a four searcher clearing schedule.

6.1.3 NSH Graph

The next graph we use has been obtained by discretization of an actual floorplan, namely the first floor of the Newel-Simon building in the Carnegie Mellon University campus. In Fig.7 we present the actual floorplan and its discretization; in Fig.8 we present the resulting graph (the node numbers in Fig.8 correspond to the cell numbers in Fig.7).




Figure 7: The NSH floorplan.




Figure 8: The NSH graph.

This graph, with 60 nodes and 64 edges, appears to be more complex than the previous ones; however it has a relatively small number of spanning trees (namely 3604, as computed by Kirchoff’s theorem). By visual inspection it appears very unlikely that the graph can be node-cleared with two searchers; on the other hand, as will be seen presently, GSST can compute three-searcher clearing schedules. Hence we conjecture that the search number s_N^{imc} is three.

Applying the “standard” variants of GSST and using $M = 5 \cdot 10^4$ spanning trees per variant we obtain the results of Table 4.

Edge Traversal	Uniform ST generation			DFS ST generation		
	Min	Prop. of min	Time	Min	Prop. of min	Time
GSST-L	3	0.00224	14.265625	3	0.01038	14.578125
GSST-R	3	0.00002	9.406250	3	0.00006	10.140625
GSST-LR	3	0.00478	21.937500	3	0.01026	21.390625
GSST-LW	3	0.00328	11.546875	3	0.00608	12.593750
GSST-LD	3	0.00544	12.625000	3	0.00880	13.578125

Table 4. Node-clearing the NSH graph by the various GSST variants: minimum number of searchers attained and proportion of minimal solutions; number of spanning trees generated is $M = 5 \cdot 10^4$.

We see that all the GSST variants using uniform spanning tree generation achieve the true search number, namely three. The random edge traversal variants perform poorest of all: the uniform GSST-R finds only one minimal search and the DFS GSST-R variant three. Generally, while the time required to complete M searches is higher for the NSH graph than for the “tree/grid” graph, the proportion of correct solutions is in some cases higher (at least for the uniform variants). A histogram of the searchers required appears in Fig.9, for the uniform GSST-L variant; for this variant the first minimal search schedule is computed at step 845 (out of 50000), approximately at time $14.265625 \cdot 845/50000 = 0.24109$ sec.




Figure 9: The histogram of the distribution of minimum number of searchers required to node clear the NSH graph by the uniform GSST-L variant.

6.1.4 National Art Gallery Graph

The next graph we use has also been obtained by discretization of an actual floorplan, namely the first floor of the National Gallery of Art, in Washington, DC. In Fig.10 we present the actual floorplan and the discretization we have used; in Fig.11 we plot the corresponding graph.

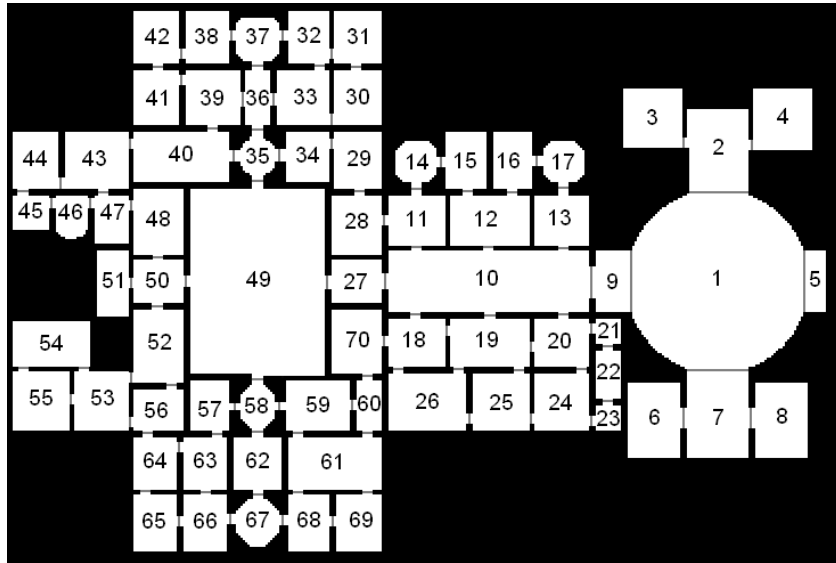


Figure 10: The National Gallery of Art floorplan.




Figure 11: The National Gallery of Art graph.

This graph, has 70 nodes and 93 edges and is more complex than any of the previously used graphs. In particular, it is more complex than the NSH graph. While the National Art Gallery graph has about 50% more edges than the NSH graph, it has a *much* larger number of spanning, approximately $5.3 \cdot 10^{14}$, as compared to 3604; this precludes use of the GSST with exhaustive enumeration of the spanning trees. Also, we have no obvious way to compute the true node search number. Nevertheless, we proceed to apply the ten GSST variants using $M = 1.5 \cdot 10^5$ spanning trees per combination. The results obtained appear in Table 5.

Edge Traversal	Uniform ST generation			DFS ST generation		
	Min	Prop. of min	Time	Min	Prop. of min	Time
GSST-L	5	0.0000066	74.378237	5	0.0004466	77.214947
GSST-R	6	0.0013133	45.500000	5	0.0004800	50.984375
GSST-LR	5	0.0000066	78.172631	5	0.0002000	87.250000
GSST-LW	5	0.0000200	62.859375	5	0.0007600	77.500000
GSST-LD	5	0.0000066	64.656250	5	0.0009800	77.546875

Table 5. Node-clearing the National Art Gallery graph by the various GSST variants: minimum number of searchers attained and proportion of minimal solutions; number of spanning trees generated is $M = 1.5 \cdot 10^5$.

We see that all the GSST variants¹⁰ find node clearing searches with five searchers. While we cannot be sure that the true node search number of the National Art Gallery is five, we have been unable to find a lower search number using any method (including extensive inspection by the authors). In addition, the middle part of the graph resembles a five-by-six grid, which is known to have a node search number of five (see Section 6.3). A histogram of the searchers required appears in 12, for the uniform GSST-L variant; for this variant the first minimal search schedule is found at step 101823 (out of a total $1.5 \cdot 10^5$). Similar results hold for the other variants.




Figure 12: The histogram of the distribution of minimum number of searchers required to node clear the National Gallery of Art graph by the uniform GSST-L variant.

The National Art Gallery graph results (as well as the NSH ones) show the applicability of GSST on graphs that are derived from representations of real indoor environments (which has been the main motivation for our research). The methods that incorporate Barriere labeling improve performance on the NSH map because it is similar to a tree (it becomes a tree if we remove only a *few* edges). In this case, Barriere labeling, which is based on trees, helps to improve the schedules generated by GSST.

¹⁰Except for uniform GSST-R; however 5-searcher clearing schedules can also be found by this variant if it is run for a sufficiently large M .

The National Art Gallery, on the other hand, is more similar to a grid, which lessens the advantage of using Barriere labeling to guide traversal (though a significant improvement is still obtained).

6.2 Interval Graphs

The next experiment presented involves, unlike the ones of Section 6.1, a large number of graphs belonging to the same *family*. Our goal is to evaluate the *average* performance of the GSST algorithm. While, strictly speaking, the results are specific to the family of interval graphs, they also suggest general properties of the GSST algorithm.

To evaluate the performance of GSST we must have some estimate of the actual search number of each graph we use; then we can compare this search number with the minimum $\overline{sn}(\mathbf{S})$ achieved by GSST. To satisfy this condition, we will work with *interval graphs*. Briefly, an interval graph is obtained from a *system of intervals*, i.e. a collection of intervals of real numbers; each interval corresponds to a graph node and two nodes are connected by an edge iff the corresponding intervals intersect. It is known [10, 22] that the (*non-monotone, non-connected*) *edge search number* of an interval graph \mathbf{G} (i.e., $s_E(\mathbf{G})$) is equal to its *interval width*, which is defined to be the size of the largest clique of the graph. It is also known that the interval width of an interval graph \mathbf{G} of N nodes can be computed in time $O(N)$ [10, 22]. Of course, while we know the *edge search number* $s_E(\mathbf{G})$ exactly, this only provides an estimate of the IMC *node-search number* $s_N^{imc}(\mathbf{G})$, in which we are really interested. There is no strict inequality connecting $s_E(\mathbf{G})$ and $s_N^{imc}(\mathbf{G})$; we just know that $s_N(\mathbf{G}) \leq s_E(\mathbf{G})$ and $s_N(\mathbf{G}) \leq s_N^{imc}(\mathbf{G})$. In general, we expect that the discrepancy between $s_E(\mathbf{G})$ and $s_N^{imc}(\mathbf{G})$ is not too large and hence $s_E(\mathbf{G})$ can be used to evaluate GSST performance.

We randomly generate interval systems (and the corresponding interval graphs) by the following mechanism. First we select two parameters of the family: N , the number of intervals, and Δ , the average interval length. Then we generate N intervals, the n -th interval having its left endpoint at n and the right endpoint at $n + \delta$, where δ follows an exponential probability law $f(\delta) \sim e^{-\delta/\Delta}$. We form the interval graph \mathbf{G} corresponding to this interval system and compute its interval width. We repeat the process 100 times to obtain 100 interval graphs; these form a family characterized by the parameters N, Δ (specified by us) and also by the average number of edges and the average interval width. We repeat the process for five different choices of (N, Δ) ; this information is summarized in Table 6.

N	Δ	Average num. of Edges	Average Interval Width
30	3	74.22	5.49
30	5	113.06	7.66
25	7	182.01	9.51
35	10	231.84	11.68
40	15	366.79	15.53

Table 6. Characteristics of the various families of interval graphs used for the experiments.

Finally, we apply the ten GSST variants to each of the five families (using $M = 2 \cdot 10^4$ spanning trees per variant). The results are summarized in Tables 7.a (for the uniform variants) and 7.b (for the DFS variants). In these tables every row corresponds to a family, the families being indexed by their average interval widths (appearing in the first column). Each of the remaining columns corresponds to one GSST variant and lists the average (over the 100 graphs) minimum number of searchers achieved by the respective variant. Hence, an estimate of the efficiency of the GSST variants can be obtained by comparing the first column of each table to the remaining ones. For example, in the first row of Table 7.a we see that all uniform GSST variants node clear the graphs of the first family with fewer searchers (on the average) than those “predicted” by the interval width of the graphs (e.g., $5.15 < 5.49$). As

we proceed down the rows of Table 7.a to graphs of higher interval width (and, presumably, of greater complexity) the average number of searchers required to clear a graph increases above the average interval width. For example, in the last row the average interval width is 15.53 and the average minimum number of searchers required by the uniform GSST-L variant is 19.05; in other words the uniform GSST-L requires $\frac{19.05-15.53}{15.53} = 22.66\%$ more searchers than expected by the interval width estimate. Things get better with the DFS variants; for example, for the ($N = 40, \Delta = 15$) family, the searcher overhead incurred by the DFS GSST-L variant is $\frac{17.68-15.53}{15.53} = 13.84\%$, which is actually quite good for graphs of such high complexity.

Av. Int. Width	GSST-L	GSST-R	GSST-LR	GSST-LW	GSST-LD
5.49	5.15	5.32	5.36	5.14	5.16
7.66	7.82	8.08	8.03	7.80	7.87
9.51	11.01	11.42	11.45	11.02	11.30
11.68	13.61	13.91	13.91	13.60	13.63
15.53	19.05	19.37	19.50	19.06	19.25

Table 7.a: Average minimum number of searchers required to node-clear each family of interval graphs by the uniform GSST variants; number of spanning trees generated is $M = 2 \cdot 10^4$.

Av. Int. Width	GSST-L	GSST-R	GSST-LR	GSST-LW	GSST-LD
5.49	5.04	5.13	5.04	5.08	5.06
7.66	7.34	7.47	7.39	7.44	7.36
9.51	10.23	10.34	10.25	10.25	10.24
11.68	12.70	12.80	12.64	12.73	12.62
15.53	17.68	17.87	17.75	17.72	17.66

Table 7.b: Average minimum number of searchers required to node-clear each family of interval graphs by the DFS GSST variants; number of spanning trees generated is $M = 2 \cdot 10^4$.

The results on interval graphs show that GSST can yield near-minimal schedules on a large class of graphs. They also demonstrate that GSST scales well with increasing complexity both in performance and computation time. This is a direct result of the linear scalability of the algorithm in the number of nodes in the environment.

6.3 Grid Graphs

6.3.1 Full Grids

We now present some experiments involving grid graphs, i.e. graphs with nodes located at points with integer coordinates (an example appears in Fig. 13). It is easily seen that a grid graph of dimension $J_1 \times J_2$ (i.e. containing nodes with coordinates $(j_1, j_2) \in \{1, 2, \dots, J_1\} \times \{1, 2, \dots, J_2\}$) can be node cleared using $J_0 = \min(J_1, J_2)$ searchers; the corresponding search schedule is obvious¹¹. However, for a *general* graph search algorithm (such as GSST) which must work without assuming any special structure of the graph, grid graphs are potentially extremely hard, because they have many cycles and many spanning trees, only a few of which correspond to minimal search schedules.

In this experiment we use 6 grid graphs, with dimensions starting at 5×5 and going up to 10×10 . To each such graph we apply the ten GSST variants; in Tables 8.a (uniform ST generation variants) and 8.b (DFS ST generation variants) we list the minimum search number attained by each variant on

¹¹Namely (assuming that the height J_1 of the grid is less than or equal to the width J_2) place J_1 searchers on one of the vertical sides of the grid and slide them horizontally to the other vertical side.

each graph (M , the number of spanning trees generated per graph is also listed, in the last column). Let us also note that the total execution time for this experiment is approximately 25 hours, reflecting (a) the large number of graphs used, (b) the high complexity of many graphs, (c) the large number of spanning trees used (the *much smaller* execution time for each individual graph and GSST variant is not listed, for economy of space).

Uniform random spanning tree generation						
Graph Dim.	GSST-L	GSST-LR	GSST-R	GSST-LW	GSST-LD	No.Trees
5×5	5	5	5	5	5	10 ⁵
6×6	7	7	7	7	7	2·10 ⁵
7×7	8	8	8	8	8	3·10 ⁵
8×8	10	9	9	9	9	3·10 ⁵
9×9	11	11	11	10	11	4·10 ⁵
10×10	13	13	13	13	13	5·10 ⁵

Table 8.a. Minimum search number attained by the uniform GSST variants on the full-grid graphs.

DFS random spanning tree generation						
Graph Dim.	GSST-L	GSST-LR	GSST-R	GSST-LW	GSST-LD	No.Trees
5×5	6	6	6	6	6	10 ⁵
6×6	7	7	8	7	7	2·10 ⁵
7×7	9	9	9	9	9	3·10 ⁵
8×8	11	11	11	11	11	3·10 ⁵
9×9	12	13	12	13	12	4·10 ⁵
10×10	14	14	14	14	15	5·10 ⁵

Table 8.b. Minimum search number attained by the DFS GSST variants on the full-grid graphs.

We see that, on the grids, the uniform variants perform better than the DFS ones; this is exactly the opposite situation from what happens in interval graphs. With few exceptions, the uniform variants (Table 8.a) either attain the actual minimum search number (for the 5 × 5 graph) or incur a small overhead of 1 or 2 extra searchers. The situation changes with the 10 × 10 graph, where the minimum attained search number is 13, with a 30% overhead over the true $s_N^{imc} = 10$. However, the 10 × 10 grid graph is very complex, with 100 nodes, 180 edges and $5.6943 \cdot 10^{42}$ spanning trees. Judging from the distribution of the attained search numbers (which, for economy of space, is not displayed here) lower search numbers *can* be attained but they require a much larger number of spanning trees. To get an idea of the complexity of the full grid graphs, we list in Table 8.c. the number of spanning trees for each of the six graphs used. This provides a measure of the hardness of node-clearing the graph. This actually quite reasonable: a graph with many cycles (and many spanning trees) is harder to clear because it contains many escape routes for the evader.

Graph Dim.	No. of spanning trees
5×5	5.57560·10 ⁸
6×6	3.2566·10 ¹³
7×7	1.9872·10 ¹⁹
8×8	1.2623·10 ²⁶
9×9	8.3266·10 ³³
10×10	5.6943·10 ⁴²

Table 8.c. Number of spanning trees for each of the full-grid graphs

On the positive side, we see that node-clearing schedules can be computed for the full grid graphs using slightly more than the minimum number of searchers and in reasonable computation time. For example, the uniform GSST-L variant finds the first 8-searcher clearing schedule for the 7×7 graph in under 2 mins and an 11-searcher clearing schedule for the 9×9 graph in under 10 mins. Considering that (due to the large number of cycles) grid graphs are potentially some of the hardest graphs for GSST¹², we find these results to be quite satisfactory.

6.3.2 Depleted Grids

Our final experiment involves what we call “depleted grid graphs”. These are obtained as follows. First we choose two parameters, J_1 and J_2 , the length and width of the grid. Just like with the full grid graphs, we place nodes at the positions $(j_1, j_2) \in \{1, 2, \dots, J_1\} \times \{1, 2, \dots, J_2\}$. We also connect all nearest neighbor nodes along the following lines

1. one horizontal line: $\{(j_1, 1)\}_{j_1 \in \{1, 2, \dots, J_2\}}$,
2. J_1 vertical lines: $\{(j_1, j_2)\}_{j_2 \in \{1, 2, \dots, J_2\}}$, for $j_1 = 1, 2, \dots, J_1$,

obtaining a tree of the form indicated by the solid edges in Fig.13. Finally, we consider all pairs of nearest neighbors $(i_1, i_2), (j_1, j_2)$ which are not already connected and add to the graph an edge connecting each such pair with probability p .




Figure 13: Generating a depleted grid graph. The solid lines correspond to edges which are always present in the graph. Each dotted line becomes an edge with probability p .

The resulting graph \mathbf{G} is a subgraph of the full $J_1 \times J_2$ grid and has $s_N^{imc}(\mathbf{G}) \leq \min(J_1, J_2)$; in other words, $\min(J_1, J_2)$ yields an upper bound (useful for estimating the performance of GSST) of $s_N^{imc}(\mathbf{G})$. It seems reasonable that, on the average, $s_N^{imc}(\mathbf{G})$ is an increasing function of p and for $p = 1$ the bound is tight, i.e. $s_N^{imc}(\mathbf{G}) = \min(J_1, J_2)$; it will be useful to keep this in mind when evaluating the results of the experiment. Note also that, as p approaches 0, the graph becomes more similar to a tree (has fewer cycles).

We generate twelve families of depleted grids; namely we use dimensions 5×5 , 7×7 , 8×8 , 10×10 and p values 0.4375, 0.7500 and 0.9375, yielding twelve combinations. We generate *fifty* graphs from each family and apply to each of these the ten GSST variants, *using* $2 \cdot 10^5$ *spanning trees for each graph*. The results appear in Tables 9.a (uniform variants) and 9.b (DFS variants). Namely, for each

¹²It is interesting to note that a somewhat related fact is well known in the probabilistic inference literature: inference is easier on graphs with low tree-width [7].

combination we present the minimum search number attained, averaged over the fifty graphs of the corresponding family. It can be seen that the resulting numbers are quite low, especially for the uniform variants, often going under the $\min(J_1, J_2)$ bound.

Graph Dim.	p	GSST-L	GSST-R	GSST-LR	GSST-LW	GSST-LD
5×5	0.4375	2.86	2.86	2.86	2.86	2.86
5×5	0.7500	3.68	3.68	3.50	3.48	3.48
5×5	0.9375	4.53	4.53	4.12	4.18	4.20
7×7	0.4375	3.63	3.56	3.64	3.57	3.62
7×7	0.7500	5.24	5.14	5.32	5.08	5.06
7×7	0.9375	6.62	6.40	6.56	6.30	6.28
8×8	0.4375	4.22	4.14	4.08	4.08	4.00
8×8	0.7500	6.10	6.14	5.90	5.86	5.84
8×8	0.9375	7.70	7.76	7.52	7.40	7.46
10×10	0.4375	5.44	5.50	5.32	5.30	5.28
10×10	0.7500	7.94	8.02	7.84	7.68	7.80
10×10	0.9375	10.20	10.40	10.24	10.06	10.94

Table 9.a. Average minimum number of searchers required to node-clear each family of depleted-grid graphs by the uniform GSST variants; number of spanning trees generated is $M = 2 \cdot 10^5$.

Graph Dim.	p	GSST-L	GSST-R	GSST-LR	GSST-LW	GSST-LD
5×5	0.4375	3.36	3.36	3.36	3.36	3.36
5×5	0.7500	4.38	4.36	4.36	4.36	4.36
5×5	0.9375	5.04	5.04	5.02	5.04	4.98
7×7	0.4375	4.24	4.24	4.24	4.24	4.24
7×7	0.7500	5.82	5.78	5.84	5.80	5.77
7×7	0.9375	7.24	7.30	7.28	7.28	7.32
8×8	0.4375	4.68	4.66	4.66	4.63	4.64
8×8	0.7500	6.56	6.62	6.62	6.64	6.62
8×8	0.9375	8.58	8.52	8.62	8.54	8.56
10×10	0.4375	5.80	5.80	5.78	5.78	5.80
10×10	0.7500	8.62	8.58	8.68	8.72	8.62
10×10	0.9375	11.34	11.38	11.54	11.38	11.26

Table 9.b. Average minimum number of searchers required to node-clear each family of depleted-grid graphs by the DFS GSST variants; number of spanning trees generated is $M = 2 \cdot 10^5$.

The depleted grid results show the high performance of GSST on graphs with large numbers of spanning trees and many cycles. Better results are obtained for smaller values of p , i.e. when the graph becomes more similar to a tree; this is to be expected, since a basic component of the GSST algorithm is search along a spanning tree. It is interesting to note that even for a quite high value of p , namely 0.9375, results (at least for uniform GSST variants) are quite good; taking into account also the results of the full grids, it appears that the increase of attained search number over $\min(J_1, J_2)$ happens only when p gets *very* close to 1. Once again we observe that the uniform variants perform better than the DFS ones.

6.4 Discussion

Our experiments demonstrate the performance of GSST on several different classes of graphs. On many complex graphs, GSST was able to find a minimal search schedule in reasonable time. On *all* graphs, at least one near-minimal schedule is computed early in the execution of the algorithm. In fact, the “anytime” operation of GSST (see also the discussion in Section 5.1) yields reasonably good solutions in very short time and then keeps improving on these solutions as long as additional computation time is available. While the minimal schedules form a small percentage of the computed solutions, many more near-minimal solutions are available. These properties of GSST are highly desirable, especially given the fact that (as far as we are aware) no other algorithm has been *implemented* which can tackle graphs of the size and complexity presented here.

We can highlight some specific conclusions supported by the experiments on the various graphs considered in Section 6.1 – 6.3.

1. The NSH and National Art Gallery graphs show the application of GSST on graphs that are derived from representations of indoor environments. The methods that incorporate Barriere labeling improve performance on the NSH map because it is similar to a tree. The National Art Gallery, on the other hand, is more similar to a grid, which lessens the advantage of using Barriere labeling to guide traversal.
2. The results on interval graphs show that GSST can yield near-minimal schedules on a large class of graphs. They also demonstrate that, with increasing complexity, GSST scales well both in performance and computation time.
3. Finally, the grid results show the high performance of GSST on graphs with large numbers of spanning trees and many cycles. These are potentially some of the hardest graphs for GSST. Also, it appears that the increase of attained search number over $\min(J_1, J_2)$ happens only when p gets *very* close to 1.

Comparing spanning tree generation methods, we see that DFS generation performs better on all graphs considered except for grid graphs (and the tree/grid graph) where the uniform ST generation is better. At this point we do not have a good explanation but we intend to further research this fact because we believe it will give us a better understanding of which are the “good” spanning trees, i.e. the ones associated with minimal searches. See Section 7 for some additional remarks on this issue.

Regarding edge traversal methods, there is no single GSST variant which consistently outperforms all others. However, uniform GSST-L never does too poorly and is fast; GSST-LW has the same advantages and, in addition, is provably complete. Hence, in case node-clearing schedules must be quickly produced (i.e. computation time is at premium), one can use a “reduced suite” of GSST variants, consisting of GSST-L and GSST-LW (with both uniform and DFS spanning tree generation) and expect to obtain nearly as good results (at a fraction of the computing time) as when using the “full” suite of the ten GSST variants.

Let us also stress that all variants of GSST can attain lower search numbers than the ones presented here if sufficient running time is available.

7 Conclusion

Motivated by the problem of robotic pursuit / evasion, in the current paper we have investigated *node search*, i.e., the capture by a team of searchers of an invisible evader located in the *nodes* of a graph. This problem has so far received little attention in the graph theoretic literature, with most related publications concentrating on the problem of edge-located evader. The basic contributions of the current paper are of two kinds.

1. From the theoretical point of view, we have shown that, in general graphs, the problem of node search is easier than that of edge search, in the sense that every edge clearing search is also a node clearing one; but the converse does not hold in general. We have then concentrated on the internal monotone connected (IMC) node search of trees and shown that it is essentially equivalent to edge search under the same restrictions; hence Barriere’s tree search algorithm, originally designed for edge search, can also be used for node search.
2. From the algorithmic point of view, we have presented GSST, a new algorithm which performs IMC search on general graphs. This algorithm is based on the fact that every node clearing search generates a spanning tree; hence node-clearing a preselected spanning tree (by the use of *tree searchers*) and simultaneously blocking recontamination (by the use of *guards*) monotonically node-clears the graph. Because spanning tree generation and search can both be performed very efficiently, a large number of spanning trees can be tried by GSST until the one yielding the minimum number of searchers is discovered. Experiments have shown that GSST can quickly node-clear large and complex graphs using a small number of searchers.

Many issues touched upon in the current paper require additional research. We conclude by listing some future research directions, categorized as follows.

1. Theory.

- (a) *Optimality.* Given a graph \mathbf{G} with IMC node search number $s_N^{imc}(\mathbf{G})$. Suppose we have K searchers with $K \geq s_N^{imc}(\mathbf{G})$. What is the best we can do? Of course “best” depends on some optimality criterion. For example we may want to clear the graph in the minimum number of steps. Or, we may want a search strategy which minimizes the cumulative “area” of the dirty set (i.e., $\sum_{t=1}^T V_N^D(t)$). We would like to obtain algorithms which solve such optimization problems.
- (b) Our experiments show that the uniform GSST variants perform better than DFS on grids and grid-like graphs. We want to discover a mathematical explanation of this fact, e.g., to find necessary and / or sufficient conditions under which uniform GSST outperforms DFS GSST. A result of this type will probably have useful computational implications – see also item 2.(a) below.
- (c) We want to establish a sequence of inequalities between node search numbers. In other words, to refine (30), either for a general graph \mathbf{G} , or for the special case when \mathbf{G} is a tree. For the latter case, we conjecture that we can re-establish the Barriere et al. inequalities (28) so as to hold for *node-* rather than edge-search numbers.
- (d) Finally, we are interested in questions of *consistency*. Suppose that we are given a graph along with a *configuration* of each nodes (i.e., a specification of the clear and dirty nodes). Can this configuration be obtained by a sequence of moves in the node game? In the edge game? Also, given a configuration of clear and dirty nodes and another one of clear and dirty edges, are these configurations compatible in the node game? In the edge game?

2. **Computation.** The GSST algorithm searches the space of spanning trees in a random manner. Is there a way to bias the search in some useful way? For this we need a characterization of “good” spanning trees, i.e., trees which yield a low number of searchers (this includes tree searchers *and* guards). We have no such characterization at present. Perhaps we can obtain it by understanding why uniform variants perform better than DFS on grid graphs (and conversely on interval graphs).
3. **Extensions.** We want to study theoretically and develop algorithms for the following variants of the search problem.

- (a) For less restricted types of node search, e.g., internal monotone (but not connected), internal connected (but not monotone); or for completely unrestricted node search.
- (b) For different types of pursuer / evader behavior; e.g., for pursuers with extended (non-local) visibility, evaders with finite speed etc.
- (c) How does the graph search problem look from the evader's point of view? Given a graph \mathbf{G} and K searchers, what is the best the evader can do if $K < s(\mathbf{G})$? What if $K \geq s(\mathbf{G})$? In the latter case capture is guaranteed, but can the evader, for example, maximize the number of steps until capture? Many other interesting problems can be posed from the evader's point of view.

Acknowledgement. We gratefully acknowledge Sanjiv Singh for many useful discussions and valuable comments on the current paper.

A Appendix: Edge Search on Trees

In this Appendix we present the **R-Label** and **R-Search** algorithms, which perform rooted IMC search of *trees* and are also used by the GSST algorithm. **R-Label** and **R-Search** are simplified versions of corresponding algorithms presented by Barriere et al. [2]; some theoretical results from [2] are also reviewed. In our presentation the terminology and notation of [2] are somewhat changed, to conform with the current paper.

The next Lemma and Definition play a crucial role in computing the IMC search number of a tree.

Lemma A.1. [2] *Given a rooted tree \mathbf{T}_x and a node y of \mathbf{T}_x , if y has more than one children, enumerate them as y_1, y_2, \dots, y_K , so that they satisfy $s_E^{imc}(\mathbf{T}_x[y_k]) \geq s_E^{imc}(\mathbf{T}_x[y_{k+1}])$ for $k = 1, 2, \dots, K - 1$. Then*

$$s_E^{imc}(\mathbf{T}_x[y]) = \max(s_E^{imc}(\mathbf{T}_x[y_1]), s_E^{imc}(\mathbf{T}_x[y_2]) + 1). \quad (39)$$

Definition A.2. [2] *Given a tree $\mathbf{T} = (V, E)$ and a node $x \in V$, consider all the edges xy incident on x and label them as follows.*

1. *If y is a leaf, then $\lambda_x(xy) = 1$.*

2. *If y is not a leaf, let z_1, z_2, \dots, z_K be the neighbors of y other than x , enumerated so that $\lambda_y(yz_k) \geq \lambda_y(yz_{k+1})$ (for $k = 1, 2, \dots, K - 1$) and define*

$$\lambda_x(xy) = \max(\lambda_y(yz_1), \lambda_y(yz_2) + 1). \quad (40)$$

Note that the above definition assigns *two* labels to each edge xy , namely $\lambda_x(xy)$ and $\lambda_y(xy)$. Intuitively, $\lambda_x(xy)$ is the number of searchers that must cross the *directed* xy edge in an IMC edge-clearing of the tree. In [2] Barriere et al. present the **Label** algorithm which computes the labels $\{\lambda_x(xy)\}_{xy \in E}$. Furthermore they prove the following.

Lemma A.3. [2] *For every tree \mathbf{T} and every edge xy of \mathbf{T} , we have $s_E^{imc}(\mathbf{T}_x[y]) = \lambda_x(xy)$.*

Lemma A.4. [2] *Given a tree $\mathbf{T} = (V, E)$, for every $x \in V$ define $\mu(x)$ as follows*

$$\mu(x) = \begin{cases} \lambda_x(xy_1) & \text{if } x \text{ has a single neighbor } y_1 \\ \max(\lambda_x(xy_1), \lambda_x(xy_2) + 1) & \text{otherwise;} \end{cases}$$

in the second line of the definition, y_1, y_2, \dots are the neighbors of x , enumerated so that they satisfy $s_E^{imc}(\mathbf{T}_x[y_k]) \geq s_E^{imc}(\mathbf{T}_x[y_{k+1}])$ for $k = 1, 2, \dots$. Then

$$s_E^{imc}(\mathbf{T}_x) = \mu(x) \quad \text{and} \quad s_E^{imc}(\mathbf{T}) = \min_{x \in V} \mu(x).$$

The algorithms and results of [2] are geared towards computing $s_E^{imc}(\mathbf{T})$. The GSST algorithm, on the other hand, requires the computation of $s_E^{imc}(\mathbf{T}_x)$ and rooted searches. Hence we introduce the **R-Label** and **R-Search** algorithms, which are simplified, “rooted” versions of the corresponding Barriere algorithms. The pseudocode of these algorithms is listed in the next page.

The **R-Label** algorithm takes as input a tree $\mathbf{T} = (V, E)$ and a root node $u_0 \in V$. The algorithm is a straightforward implementation of Definition A.2. **R-Label** does *not* compute all the labels $\{\lambda_x(xy)\}_{xy \in E}$ but only the ones for which y is a child of x in the rooted tree \mathbf{T}_{u_0} ; for these edges, **R-Label** produces the same λ labels as **Label**. **R-Label** uses two subroutines: given that u_0 is the root, the subroutine **Depth**(\mathbf{T}, u_0) returns a partition $\{V^{(l)}\}_{l=1}^L$ of the node set V , where $V^{(l)}$ contains all nodes of depth l ; the subroutine **SortChildren**($y, \mathbf{T}, u_0, \lambda$) returns the children of y sorted in decreasing order of their λ labels.

The **R-Search** algorithm uses the λ labels to produce a rooted IMC (edge and node) clearing search of \mathbf{T} . **R-Search** is almost identical to Barriere's **Search**, the only difference being that the starting node u_0 is given, rather than chosen by the algorithm (this implies that **R-Search** produces a minimal *rooted* IMC search of \mathbf{T}_{u_0}). **R-Search** uses Barriere's subroutine **Move**. The notation $\mathbf{S} = \mathbf{S} | (u, v, J)$ means: append J moves of the form $u \rightarrow v$ to the search schedule \mathbf{S} .

Algorithm 3 R-Label(\mathbf{T}, u_0)

Input: Tree $\mathbf{T} = (V, E)$, start node u_0
 $\{V^{(l)}\}_{l=1}^L = \mathbf{Depth}(\mathbf{T}, u_0)$
for $l = L - 1$ **with step** -1 **until** 1 **do**
 for $x \in V^{(l)}$ **do**
 for $y \in \mathbf{SortChildren}(x, \mathbf{T}, u_0, \lambda)$ **do**
 $[z_1, z_2, \dots, z_K] = \mathbf{SortChildren}(y, \mathbf{T}, u_0, \lambda)$
 if $\mathbf{Length}([z_1, z_2, \dots, z_K]) = 0$ **then**
 $\lambda_x(xy) = 1$
 else if $\mathbf{Length}([z_1, z_2, \dots, z_K]) = 1$ **then**
 $\lambda_x(xy) = \lambda_y(yz_1)$
 else
 $\lambda_x(xy) = \max(\lambda_y(yz_1), \lambda_y(yz_2) + 1)$
 end if
 end for
 end for
end for
Output: Edge labeling $\{\lambda_x(xy)\}_{xy \in E}$

Algorithm 4 R-Search(\mathbf{T}, u_0, λ)

Input: Tree $\mathbf{T} = (V, E)$, start node u_0 , edge labeling λ
 $\mathbf{S} = \emptyset$
 $[y_1, y_2, \dots, y_K] = \mathbf{SortChildren}(u_0, \mathbf{T}, u_0, \lambda)$
for $k = K$ **with step** -1 **until** 1 **do**
 $\mathbf{S} = \mathbf{Move}(x, y_k, \lambda_x(xy_k), \mathbf{S})$
end for
Output: Search \mathbf{S}

Algorithm 5 $\text{Move}(u, v, J, \mathbf{S})$

Input: Tree $\mathbf{T} = (V, E)$, start node u_0 , edge labeling λ

$\mathbf{S} = \mathbf{S}|(u, v, J)$

$[w_1, w_2, \dots, w_K] = \text{SortChildren}(v, \mathbf{T}, u_0, \lambda)$

for $k = K$ **with step** -1 **until** 1 **do**

$\mathbf{S} = \text{Move}(v, w_k, \lambda_y(yw_k), \mathbf{S})$

end for

$\mathbf{S} = \mathbf{S}|(v, u, J)$

Output: Search \mathbf{S}

B Appendix: The relationship between node search and *mixed* edge search

Theorem 3.3 tells us that edge search is “weaker” than node search in the sense that, for every graph and every search schedule, we have $V_E^C(t) \subseteq V_N^C(t)$ and $E_E^C(t) \subseteq E_N^C(t)$. We will now consider a variant of edge search, the so-called *mixed edge search* which, as we will see, is *equivalent* to node search.

Mixed edge search is obtained by augmenting the clearing rules. In the “*Mixed Search Edge Game*” an edge can be cleared not only by sliding, but also by guarding both its endpoints. More precisely, to obtain the rules of the mixed edge search game, take the rules of the edge game (presented in Section 2), change the terms “e-clear” and “e-dirty” to “m-clear” and “m-dirty” and add the following rule:

E2' An m-dirty edge uv becomes m-clear when both u and v are occupied by searchers.

In other words, we can clear an edge either by traversing it or by guarding both its endpoints.

We will use $V_M^C(t)$ to denote the set of m-clear nodes at time t , and will analogously use the notations $E_M^C(t)$, $V_M^D(t)$, $E_M^D(t)$, $\mathbf{G}_M^C(t)$ etc.

We emphasize that, in the mixed search edge game, the evader resides in the edges of the graph and recontamination occurs by the same rules as in the “plain” edge game.

Theorem B.1. *Given a graph \mathbf{G} and an internal search schedule \mathbf{S} . Then, we have*

$$\text{for } t = 0, 1, 2, \dots : \quad V_N^C(t) = V_M^C(t) \quad \text{and} \quad E_N^C(t) = E_M^C(t). \quad (41)$$

Proof. The proof is inductive and consists of several steps.

I. At $t = 0$ we have

$$V_N^C(0) = V_M^C(0) = \emptyset \quad \text{and} \quad E_N^C(0) = E_M^C(0) = \emptyset. \quad (42)$$

Suppose that we have

$$\text{for } s = 0, 1, \dots, t : \quad V_N^C(s) = V_M^C(s) \quad \text{and} \quad E_N^C(s) = E_M^C(s). \quad (43)$$

and let us next consider time $t + 1$.

II. We will now prove that

$$E_N^C(t + 1) \subseteq E_M^C(t + 1). \quad (44)$$

To this end we will show two things.

II.a *If a previously n-dirty edge becomes n-clear at $t + 1$, then it also becomes m-clear.* Suppose that $uv \in E_N^D(t) \cap E_N^C(t + 1)$. In other words, uv is n-cleared exactly at $t + 1$; this means that one of u, v was already n-clear by t and the other was n-cleared *exactly* at $t + 1$. Without loss of generality, assume

$$u \in V_N^C(t) \cap V_N^C(t + 1), \quad v \in V_N^D(t) \cap V_N^C(t + 1) \quad (45)$$

Since a node can be n-cleared only by moving into it, the $(t + 1)$ -th move must be $w \rightarrow v$. We distinguish the following cases.

(i.1) Suppose $w \neq u$. Since $v \in V_N^D(t)$ and $u \in V_N^C(t)$, u was guarded at t ; since $w \neq u$, u remains guarded at $t + 1$. Hence, at $t + 1$, uv has both ends guarded and so $uv \in E_M^C(t + 1)$. (Note that this analysis holds even when $w = 0$, i.e., when the move was to place a new searcher).

(i.2) Suppose $w = u$ and u is still guarded at $t + 1$ (i.e., at t there was more than one searcher in u). Then at $t + 1$ both u and v are guarded and $uv \in E_M^C(t + 1)$.

(i.3) The remaining possibility is that $w = u$ and u is unguarded at $t + 1$ (i.e., at t there was exactly one searcher in u). Note that $u \rightarrow v$ at $t + 1$ means that u is guarded at t and hence $u \in V_M^C(t)$. If $u \in V_M^C(t + 1)$ as well, then $uv \in V_M^C(t + 1)$. So we now will show that $u \in V_M^D(t + 1)$ is not possible. Indeed $u \in V_M^D(t + 1)$ can only happen in one of the following two ways.

(i.3.1) There exists a node $x_1 \neq v$ such that $ux_1 \in E_M^D(t) \cap E_M^D(t + 1)$. But $ux_1 \in E_M^D(t) = E_N^D(t)$ implies $x_1 \in V_N^D(t)$ (because u is guarded at t) and $x_1 \in V_N^D(t)$ implies $x_1 \in V_N^D(t + 1)$ (because x_1 was not entered at $t + 1$). Finally, from $x_1 \in V_N^D(t + 1)$ and u unguarded at $t + 1$, we conclude $u \in V_N^D(t + 1)$ which contradicts assumption (45).

(i.3.2) Alternatively, there exists a path $vux_1 \dots x_K$ (with $K \geq 2$), e-unguarded at $t + 1$ and with $x_{K-1}x_K \in E_M^D(t) \cap E_M^D(t + 1)$ (note also that, since $vux_1 \dots x_K$ is a path, $\{u, v\} \cap \{x_1, x_2, \dots, x_K\} = \emptyset$). In this case, we have

$$\left. \begin{array}{l} vux_1 \dots x_K \text{ e-unguarded at } t + 1 \\ u \rightarrow v \text{ at } t + 1 \\ u \notin \{x_1, x_2, \dots, x_K\} \end{array} \right\} \Rightarrow ux_1 \dots x_K \text{ is e-unguarded at } t;$$

and

$$\left. \begin{array}{l} ux_1 \dots x_K \text{ e-unguarded at } t \\ x_{K-1}x_K \in E_M^D(t) \end{array} \right\} \Rightarrow x_{K-1} \in V_M^D(t) = V_N^D(t);$$

and finally

$$\left. \begin{array}{l} x_{K-1} \in V_N^D(t) \\ u \rightarrow v \text{ at } t + 1 \\ v \notin \{x_1, x_2, \dots, x_K\} \\ ux_1 \dots x_{K-1} \text{ is n-unguarded at } t + 1 \end{array} \right\} \Rightarrow x_{K-1} \in V_N^D(t + 1) \left. \right\} \Rightarrow u \in V_N^D(t + 1)$$

But $u \in V_N^D(t + 1)$ contradicts assumption (45). Hence $u \notin V_M^D(t + 1)$.

In short, by examining cases (i.1)-(i.3) we have shown that

$$uv \in E_N^D(t) \cap E_N^C(t + 1) \Rightarrow uv \in E_M^C(t + 1). \quad (46)$$

i.e., if at $t + 1$ some uv is n-cleared it is also m-cleared.

II.b If a previously m-clear edge becomes m-dirty at $t + 1$, then it also becomes n-dirty. Suppose that $uv \in E_M^C(t) \cap E_M^D(t + 1)$. Without loss of generality, we can assume that there exists a $x_{K-1}x_K \in E_M^D(t) \cap E_M^D(t + 1)$ and a path $uvx_1 \dots x_K$ which was e-guarded at t but became e-unguarded at $t + 1$. Since the path became e-unguarded at $t + 1$, exactly one node in it was guarded at t (and became unguarded at $t + 1$). Call this node $x_{\bar{k}}$, with $\bar{k} = 0$ if the node in question is v (i.e., $x_0 = v$) and $\bar{k} \in \{1, 2, \dots, K - 1\}$ otherwise. The move at $t + 1$ is $x_{\bar{k}} \rightarrow z$, where $z \notin \{v, x_1, \dots, x_{K-1}\}$ (since $uvx_1 \dots x_K$ is e-unguarded at $t + 1$). Since $x_{K-1}x_K \in E_M^D(t) = E_N^D(t)$, either $x_{K-1} \in V_N^D(t)$ or $x_K \in V_N^D(t)$. We consider the following cases.

(ii.1) If $x_{K-1} \in V_N^D(t)$ then $x_{K-1} \in V_N^D(t + 1)$ as well (it was not entered) and the path $vx_1 \dots x_{K-1}$ is n-unguarded at $t + 1$. So $v \in V_N^D(t + 1)$ and $uv \in E_N^D(t + 1)$.

(ii.2) If $x_{K-1} \notin V_N^D(t)$, then $x_K \in V_N^D(t)$ and so x_{K-1} is the guarded node in the path $uvx_1 \dots x_K$ (i.e., $\bar{k} = K - 1$). We distinguish two subcases.

(ii.2.1) The move at $t + 1$ is $x_{K-1} \rightarrow z$ and $z \neq x_K$. Then $x_K \in V_N^D(t + 1)$ and so $x_{K-1} \in V_N^D(t + 1)$ and, by the same argument as in (ii.1), we get $uv \in E_N^D(t + 1)$.

(ii.2.2) The move at $t+1$ is $x_{K-1} \rightarrow x_K$. Since we have assumed $x_{K-1}x_K \in E_M^D(t) \cap E_M^D(t+1)$, $x_{K-1}x_K$ must be adjacent to some *other* edge $x_{K-1}y \in E_M^D(t) \cap E_M^D(t+1)$; so we can use y in place of x_K , $uvx_1\dots x_{K-1}y$ in place of $uvx_1\dots x_K$ and conclude by the reasoning of case (ii.2.1) that $uv \in E_N^D(t+1)$.

In short, we have shown that

$$uv \in E_M^C(t) \cap E_M^D(t+1) \Rightarrow uv \in E_N^D(t+1), \quad (47)$$

i.e., if at $t+1$ some uv is m-dirtied then it is also n-dirtied.

Hence every edge added at $t+1$ to $E_N^C(t)$ is also added to $E_M^C(t)$ and every edge removed at $t+1$ from $E_M^C(t)$ is also removed from $E_N^C(t)$. These facts, combined with $E_N^C(t) = E_M^C(t)$, yield

$$E_N^C(t+1) \subseteq E_M^C(t+1). \quad (48)$$

III. To strengthen (48) to set equality, we need the reverse set inclusion. In fact we can prove

$$\text{for } s = 0, 1, \dots, t, t+1, \dots : \quad E_M^C(s) \subseteq E_N^C(s) \quad (49)$$

by following the proof of Theorem 3.3 and replacing “e-clear”, “e-dirty”, ... with “m-clear”, “m-dirty”, ... (analogous of Lemmas 3.1, 3.2 can also be proved easily)¹³. From (48) and (49) we obtain

$$E_M^D(t+1) = E_N^D(t+1). \quad (50)$$

IV. Now suppose $u \in V_M^D(t+1)$. Then u is unguarded and adjacent to some $uv \in E_M^D(t+1) = E_N^D(t+1)$. Hence either u or v is n-dirty and so (since u is unguarded) $u \in V_N^D(t+1)$. On the other hand, if $u \in V_N^D(t+1)$, then u is unguarded and there exists a $uv \in E_N^D(t+1) = E_M^D(t+1)$ and (since u is unguarded) $u \in V_M^D(t+1)$. It follows that

$$V_M^D(t+1) = V_N^D(t+1).$$

V. In short, we have established that

$$\left(\begin{array}{l} E_M^C(t) = E_N^C(t) \\ V_M^C(t) = V_N^C(t) \end{array} \right) \Rightarrow \left(\begin{array}{l} E_M^C(t+1) = E_N^C(t+1) \\ V_M^C(t+1) = V_N^C(t+1) \end{array} \right). \quad (51)$$

Eq.(51) and the fact $V_M^C(0) = V_N^C(0)$ and $E_M^C(0) = E_N^C(0)$ complete the proof of the theorem. \square

Several facts follow from Theorem B.1. First, it is well known that mixed edge search is NP-complete [9] hence node search is also NP-complete.

Second, Theorem B.1 can be used to obtain a Theorem 3.3 as a corollary with a short proof (which, of course, presupposes the lengthy proof of Theorem B.1). We give a sketch of such a proof. Suppose a graph \mathbf{G} and a search schedule \mathbf{S} are given. Then, by Theorem B.1 we have (for every t) $V_N^C(t) = V_M^C(t)$ and $E_N^C(t) = E_M^C(t)$. Now suppose \mathbf{S} is applied to \mathbf{G} under the rules of the edge game and, up to time t we have $V_E^C(t) \subseteq V_M^C(t)$ and $E_E^C(t) \subseteq E_M^C(t)$. Then, at time $t+1$, every e-cleared edge is also m-cleared (since the mixed game has all the clearing rules of the edge game and an additional one) and every m-dirtied edge is also e-dirtied (since the edge and mixed games have the same recontamination rules) except if an edge was already e-dirty. Hence $E_E^C(t+1) \subseteq E_M^C(t+1)$, which also shows that $V_E^C(t+1) \subseteq V_M^C(t+1)$.

Finally, the above Theorem B.1 and the monotonicity result proved in [5] (“if there is a mixed edge clearing search of \mathbf{G} using $\leq K$ guards, there is a monotone mixed edge clearing search of \mathbf{G} using $\leq K$ guards”) makes the following conjecture seem almost trivially true.

¹³Basically the proof remains valid because the *recontamination rules* remain the same.

Conjecture B.2. *If there is a node clearing search \mathbf{S} of \mathbf{G} using at most K guards, there is a monotone node clearing search \mathbf{S}' of \mathbf{G} using at most K guards.*

However some additional work is required to prove the conjecture, because of the following detail: edge monotonicity does not necessarily imply node monotonicity. In other words, the implication

$$E_N^C(t) \subseteq E_N^C(t+1) \Rightarrow V_N^C(t) \subseteq V_N^C(t+1) \quad (52)$$

is not necessarily true! Consider again the search schedule of Remark 2.2. In this case $E_N^C(1) = E_N^C(2) = \emptyset$, but $V_N^C(1) \subseteq V_N^C(2)$ does not hold. It is true however (and easy to prove) that

$$(\textit{strong edge monotonicity}) \Rightarrow (\textit{node monotonicity}). \quad (53)$$

Unfortunately, the result proved by Bienstock in [5] concerns “simple”, not strong monotonicity. Hence one of our future research goals is to extend Bienstock’s result to node search and thus prove that $s_N(\mathbf{G}) = s_N^m(\mathbf{G})$.

C Appendix: Vertex separation and pathwidth

Vertex separation and pathwidth are strongly (but not obviously) related to graph search and search number, as discussed in (among other papers) [29, 12], where the author can find a more detailed discussion; here we just give the basic definitions and the main theorem.

Definition C.1. Given a graph $\mathbf{G} = (V, E)$, a path decomposition of \mathbf{G} is a pair (X, \mathbf{P}) , where $X = \{X_1, \dots, X_M\}$ is a family of subsets of V and \mathbf{P} is a path whose nodes are the subsets X_i and they satisfy the following properties

1. $\cup_{m=1}^M X_m = V$.
2. For every edge vw in E , there is a subset X_m that contains both v and w .
3. If $1 \leq i \leq j \leq k \leq M$ then $X_i \cap X_k \subseteq X_j$.

Definition C.2. The width of a path decomposition (X, \mathbf{P}) (with $X = \{X_1, \dots, X_n\}$) is denoted by $pdw(X, \mathbf{P})$ and is defined by

$$pdw(X, \mathbf{P}) = \max_{1 \leq i \leq n} |X_i|.$$

Definition C.3. The pathwidth of a graph $\mathbf{G} = (V, E)$ is denoted by $pw(\mathbf{G})$ and defined by

$$pw(\mathbf{G}) = \min_{\text{all path decompositions } (X, \mathbf{P}) \text{ of } \mathbf{G}} pdw(X, \mathbf{P}).$$

Definition C.4. Given a graph $\mathbf{G} = (V, E)$ (with $|V| = N$) and a permutation \mathcal{P} of $\{1, 2, \dots, N\}$. The vertex separation of \mathbf{G} with respect to \mathcal{P} is denoted by $vs(\mathbf{G}, \mathcal{P})$ and defined by

$$vs(\mathbf{G}, \mathcal{P}) = \max_{1 \leq n \leq N} |\{u \in V : \mathcal{P}(u) \leq n \text{ and } \exists v \text{ such that: } \mathcal{P}(v) > n \text{ and } uv \in E\}|.$$

The vertex separation of G is denoted by $vs(\mathbf{G})$ and defined by

$$vs(\mathbf{G}) = \min_{\text{all permutations } \mathcal{P}} vs(\mathbf{G}, \mathcal{P}).$$

The following theorem shows the connection between $vs(\mathbf{G})$ and $pw(\mathbf{G})$ and search number.

Theorem C.5. For every graph \mathbf{G} , $vs(\mathbf{G}) = pw(\mathbf{G})$ and $|vs(\mathbf{G}) - s_E(\mathbf{G})| \leq 1$.

D Appendix: Implementations of the GSearch Algorithm

We have implemented the GSST algorithm in two forms which are publicly available.

1. As a command line executable, which runs on Windows and Linux computers. The program and supporting material are available at the URL

<http://www.frc.ri.cmu.edu/~gholling/home/software.html>.

2. As a graphical user interface (GUI), available at

<http://users.auth.gr/~kehagiat/KehagiasSoftware.htm>.

We discuss each of these implementations separately.

D.1 The Command Line program

This is an executable, `gsearch.exe`, which has the following usage

```
USAGE: gsearch -m [graph] -n [no.trees] -s [startnode] -g [gen.tree]
      -y [wr-search] -w [wr-tree] -v [visualize] -t [edge traversal]
      -i [improve tree] -l [low number] -r [redundancy check]
EXAMPLE: gsearch -m graphs/Edge01.txt -g exhaustive -s 1 -n 500000 -t bh
```

The various options available are as follows.

```
-m [graph]: string, name of file with edge list of the graph.
-n [no.trees]: int, how many sp.trees to generate (DEFAULT is 1).
-s [start node]: int, which node to start
  (DEFAULT is 1, random choice is 0)
-g [gen.tree]: string, method of generating spanning trees
  (acceptable values: readtree, exhaustive, uniform, dfsrand;
  DEFAULT is uniform)
-y [wr-search]: string, how many best searches to write
  (acceptable values: one, all, none; DEFAULT is one)
-w [wr-tree]: string, how many best sp.trees to write
  (acceptable values: one, all, none; DEFAULT is one)
-t [traversal]: string, how to break edge label ties
  (acceptable values: bh, random, bhrand, bhweight, bhdom;
  DEFAULT is bh).
-i [improve tree]: boolean, use tree improve technique or not.
  (acceptable values are 0 / 1; DEFAULT is 0: do not use it)
-r [redundancy check]: boolean, check for redundant trees
  (acceptable values are 0 / 1; DEFAULT 1:
  check for redundancy)
-l [low number]: break if a tree is found with this number of
  searchers (DEFAULT is 0: do not break)
-v [visualize]: boolean, use visualizer (only supported
  on linux) (DEFAULT is 0: do not use)
-h [help]
```

The `-m` option indicates the file which contains the graph description. This must be an ascii (plain text) file containing a list of the edges of the graph, one edge per line, indicated as a pair of nodes. The nodes must be continuously numbered from 1 to N and these numbers are used as labels. N is assumed to be the largest number appearing in the edge list (the graph is assumed to be undirected and connected). Examples of edge lists can be found in the `graphs` directory.

The remaining options of `gsearch` correspond to the description of the algorithm in Section 5. The `-n` option corresponds to the M parameter of GSST (number of spanning trees), `-s` is the root of the search, `-g` describes the uniform and DFS methods for generating spanning trees (there are also options for exhaustive enumeration of all spanning trees and for reading a specific spanning tree from file) and the `-t` options (`bh`, `random`, `bhrand`, `bhweight`, `bhdom`) corresponds to the traversal methods (L, R, LR, LW and LD, respectively). The `-i`, `-r` and `-l` options are self explanatory. The `-y` option writes one or more minimal strategies in the file(s) `output/strat*.txt`; each row of the file corresponds to one step of the strategy and shows the nodes in which the searchers are currently located. The `-w` option writes the rooted tree(s) corresponding to optimal strategies, in the file(s) `output/tree*.txt`; each row of the file shows one parent and her child. The `-v` visualization option only works on Linux computers (for a visualization of the search on Windows computers use the GUI).

The command line `gsearch` was implemented by G. Hollinger in ANSI C using the gcc 4.1 compiler.

D.2 The Graphical Interface

This is an executable, `gsearchGUI.exe`, which corresponds closely to the command line program. (In fact the GUI is a front end for the command line program.) Launching the program brings up the window of Figure 14. On the left side we see several input boxes. All of these correspond to abovementioned options of the command line program with one exception: in addition to the edge list file (`graphs/Edge11.txt` in Figure 14) there is also a file containing the x - and y -coordinates of the nodes (`graphs/Node11.txt` in Figure 14 – this is required for visualization). The GUI starts with the default selection of `Edge01.txt` and `Node01.txt` ; if you want a different graph, type the corresponding file names in the input boxes.

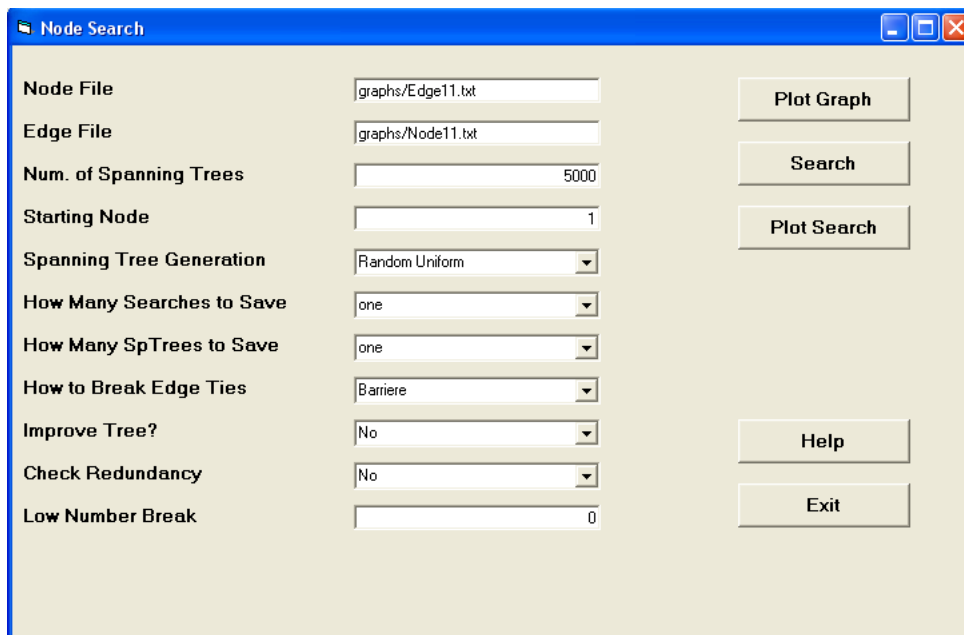


Figure 14: A screenshot of the `gsearch` GUI.

Having completed the input boxes on the left side of the window, you can now use the buttons on the right side.

1. The **Search** button performs the graph search (it actually invokes the `gsearch` executable and stores results in the `output` directory). Time to complete the search depends on the size of the graph and the number of spanning trees used. Do *not* use the exhaustive spanning tree enumeration option on large graphs because the program may take too long to terminate.
2. The **Plot Graph** button plots the graph. An image of the graph is generated and an image viewer is invoked which can be used to view the image file (this is always named `graph.jpg`). A screenshot of the image viewer appears in Figure 15.
3. The **Plot Search** button plots the graph search. Actually, first the search strategy files are used to generate a sequence of image files, one image corresponding to each step of the search (these files are stored in the directory `pix`) and then the image viewer is invoked to view the files. The user can step through the images using the arrow keys, or use the `start / stop` buttons to run a slide show of the search. A screenshot of the image viewer appears in Figure 15.

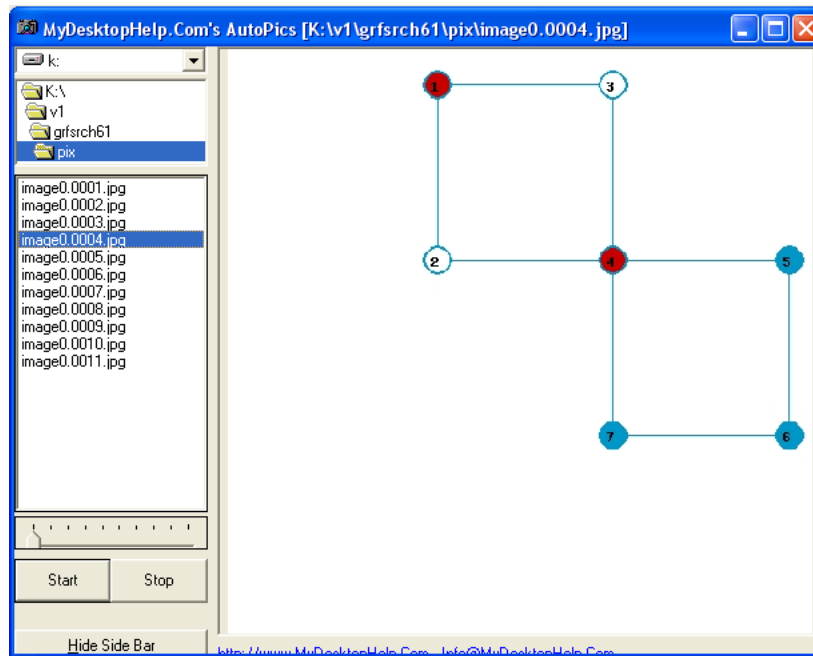


Figure 15: A screenshot of the image viewer.

Nota bene: do not run **Plot Search** before running **Search**! The **Plot Search** function uses the latest results of **Search**; if these correspond to an earlier graph (from a previous run of the GUI) the plotted search will produce nonsense results.

The GUI was implemented by Ath. Kehagias in MS Visual Basic 5.0 (and a little bit of C, using the dev-cpp 4.9.9.2 environment). The image viewer is a freeware program called AutoPics and made available by the company *Mydesktophelp* (at <http://www.mydesktophelp.com>). The GD graphics library (version gd-2.0.34-win32) has also been used; this library is available at <http://www.libgd.org/>.

Finally it must be stressed that the current version, `gsearchGUI v.0.9`, is still *beta* and requires further development. However we make it publicly available in the hope that it will prove useful to the graph search community.

References

- [1] B. Alspach, “Searching and sweeping graphs: a brief survey”, *Le Matematiche*, vol. 59, pp. 5–37, 2006.
- [2] L. Barriere, P. Flocchini, P. Fraigniaud, and N. Santoro, “Capture of an intruder by mobile agents”, in *Proc. of the 14th annual ACM symposium on Parallel algorithms and architectures (SPAA 2002)*, pp. 200–209, 2002.
- [3] L. Barriere, P. Fraigniaud, N. Santoro, and D. M. Thilikos, “Searching is not jumping”, in *Proc. of the 29th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2003)*, vol. 2880 of Springer LNCS, pp. 34–45, 2003.
- [4] D. Bienstock, “Graph searching, path-width, tree-width and related problems (a survey)”, in *Reliability of computer and communication networks*, vol. 5 of DIMACS Ser. Discrete Math. Theoret. Comput. Sci., pp. 33–49, 1991.
- [5] D. Bienstock and P. Seymour, “Monotonicity in graph searching”, *J. Algorithms*, vol.12, pp. 239–245, 1991.
- [6] R. Breisch, “An intuitive approach to speleotopology”, *Southwestern Cavers*, vol.6, pp. 72–78, 1967.
- [7] M. Chavira. *Beyond Treewidth in Probabilistic Inference*. Ph.D. dissertation, UCLA Dept. of Computer Science, 2007.
- [8] N. D. Dendris, L. M. Kirousis, and D. M. Thilikos, “Fugitive-search games on graphs and related parameters”, *Theoret. Comput. Sci.*, vol. 172, pp. 233–254, 1997.
- [9] J. A. Ellis, I. H. Sudborough, and J. S. Turner, “The vertex separation and search number of a graph”, *Inform. and Comput.*, vol. 113, pp. 50–79, 1994.
- [10] F.V. Fomin, P. Heggernes and R. Mihai. “Mixed search number and linear-width of interval and split graphs”. *Springer LNCS*, vol. 4769, pp.304-315, 2007
- [11] F. V. Fomin, D. M. Thilikos, and I. Todinca, “Connected graph searching in outerplanar graphs”, *Electronic Notes in Discrete Mathematics*, vol. 22, pp. 213–216, 2005.
- [12] F. V. Fomin, D. M. Thilikos, “An annotated bibliography on guaranteed graph searching”. *Theor. Comp. Sci.*, vol.399, pp.236-245, 2008.
- [13] P. Fraigniaud, F. V. Fomin, and D. M. Thilikos, “Connected branch decomposition and graph searching”, in *Proc. of SIAM Conference on Discrete Mathematics*, 2006.
- [14] P. Fraigniaud and N. Nisse, “Connected treewidth and connected graph searching”, in *Proc. of the 7th Latin American Symposium on Theoretical Informatics (LATIN 2006)*, vol. 3887 of Springer LNCS., pp. 479–490, 2006.
- [15] B.P. Gerkey, S. Thrun and G. Gordon. “Visibility-based pursuit-evasion with limited field of view”. *Proc. of the Natl. Conf. on Artificial Intelligence (AAAI 2004)*, pp. 20-27, 2004.
- [16] B.P. Gerkey, S. Thrun and G. Gordon. “Parallel stochastic hill-climbing with small teams”. In *Multi-Robot Systems: From Swarms to Intelligent Automata III*, pp. 65-77, 2005.

- [17] L. J. Guibas, J.-C. Latombe, S. M. Lavalle, D. Lin, and R. Motwani, “A visibility-based pursuit-evasion problem”, *Int. J. Comput. Geom. Appl.*, vol. 9, pp. 471–493, 1999.
- [18] J.M. Harris, J.L. Hirst and M.J. Mossinghoff. *Combinatorics and Graph Theory*. Springer, 2008.
- [19] G. Hollinger, Ath. Kehagias and S. Singh. “Probabilistic strategies for pursuit in cluttered environments with multiple robots”. In *Proc. ICRA 2007*, 2007.
- [20] G. Hollinger, Ath. Kehagias, S. Singh, D. Ferguson, and S. Srinivasa. “Anytime Guaranteed Search using Spanning Trees”. Tech. Report CMU-RI-TR-08-36, Robotics Institute, Carnegie Mellon University, 2008.
- [21] R. Jayakumar, K. Thulasiraman and M.N.S. Swamy. “MOD-CHAR: an implementation of Char’s spanning tree enumeration algorithm and its complexity analysis”. *IEEE Trans. on Circuits and Systems*, vol. 36, pp.219-228, 1989.
- [22] H Kaplan, R Shamir. “Pathwidth, bandwidth, and completion problems to proper interval graphs with small cliques”. *SIAM Journal on Computing*, vol. 25, pp.540-561, 1996.
- [23] Ath. Kehagias, G. Hollinger, and S. Singh. “A Graph Search Algorithm for Indoor Pursuit / Evasion”. Tech. Report CMU-RI-TR-08-38, Robotics Institute, Carnegie Mellon University, 2008.
- [24] N. G. Kinnersley, “The vertex separation number of a graph equals its pathwidth”, *Inf. Process. Lett.*, vol. 42, pp. 345–350, 1992.
- [25] L. M. Kirousis and C. H. Papadimitriou, “Searching and pebbling”, *Theoret. Comput. Sci.*, vol. 47, pp. 205–218, 1986.
- [26] A. S. LaPaugh, “Recontamination does not help to search a graph”, *J. Assoc. Comput. Mach.*, vol. 40, pp. 224–245, 1993.
- [27] S. M. Lavalle, B. H. Simov, and G. Slutzki, “An algorithm for searching a polygonal region with a flashlight”, *Int. J. Comput. Geom. Appl.*, vol.12, pp. 87–113, 2002.
- [28] S.M. LaValle. *Planning Algorithms*. Cambridge University Press, 2008.
- [29] N. Megiddo, S. L. Hakimi, M. R. Garey, D. S. Johnson, and C. H. Papadimitriou, “The complexity of searching a graph”, *J. Assoc. Comput. Mach.*, vol. 35, pp. 18–44. 16, 1988.
- [30] R. Nowakowski and P. Winkler, “Vertex-to-vertex pursuit in a graph”, *Discrete Math.*, vol. 43, pp. 235–239, 1983.
- [31] T. D. Parsons, “Pursuit-evasion in a graph”, in *Theory and applications of graphs*, Springer LNM, vol. 642, pp. 426– 441, 1978.
- [32] N. N. Petrov, A problem of pursuit in the absence of information on the pursued, *Differentsialnye Uravneniya*, vol. 18, pp. 1345–1352, 1982.
- [33] S. Singh and M. Wagh. “Robot path planning using intersecting convex shapes: Analysis and simulation”, *IEEE Journal of Robotics and Automation*, vol.?, pp.??-??, 1987.
- [34] J. D. H. Smith, “Minimal trees of given search number”, *Discrete Math.*, vol. 66, pp. 191–202, 1987.
- [35] A. Takahashi, S. Ueno, and Y. Kajitani, Mixed searching and proper- path-width, *Theoret. Comput. Sci.*, 137 (1995), pp. 253–268.

- [36] D. M. Thilikos, “Algorithms and obstructions for linear-width and related search parameters”, *Discrete Appl. Math.*, vol. 105, pp. 239–271, 2000.
- [37] D.B. Wilson, “Generating random spanning trees more quickly than the cover time”, *Proc. of the 28th annual ACM symposium on Theory of Computation*, pp.296-303, 1996.
- [38] B. Yang, “Strong-mixed searching and pathwidth”, *J. Comb. Optim.*, vol. 13, pp. 47–59, 2007.
- [39] B. Yang, D. Dyer, and B. Alspach, “Sweeping graphs with large clique number”, in *Proc. of the 15th International Symposium on Algorithms and Computation (ISAAC 2004)*, Springer LNCS, vol. 3341, pp. 908–920, 2004.
- [40] S Zilberstein. “Using anytime algorithms in intelligent systems”. *AI magazine*, vol. 17, pp. 73-83, 1996.