AUSTRALASIAN JOURNAL OF COMBINATORICS
Volume **53** (2012), Pages 97–108

A time-constrained variation of the Watchman's Walk Problem

Danny Dyer*

Department of Mathematics and Statistics
Memorial University of Newfoundland
St. John's
Canada

Rebecca Milley†

Department of Mathematics and Statistics
Dalhousie University
Halifax
Canada

Abstract

Given a graph and a single watchman, the objective of the Watchman's Walk Problem is to find a closed dominating walk of minimum length which the watchman can traverse to efficiently guard the graph. When multiple guards are available, one natural variation is to assume fixed time constraints on the monitoring of vertices and attempt to minimize the number of guards required. We find upper bounds on the number of guards required to monitor trees when no vertex is unobserved for more than t units of time.

1 Introduction

In graph theory, a *dominating set* of a graph G is a set of vertices $D \subseteq V(G)$ with the property that every vertex of G is either in D or adjacent to a vertex of D. The concept of graph domination is widely researched (see [6]), and many results are known about the *domination number* of a graph: that is, the size of a smallest dominating set, denoted $\gamma(G)$.

* Supported by NSERC
† Supported by NSERC

A common application of domination is to model some sort of network with a graph and then use dominating sets to efficiently monitor the network. For example, given a museum whose rooms must be monitored, we could represent each room as a vertex and connect two vertices with an edge if there is a hallway between the corresponding rooms. If we assume a guard can see from one room to an adjacent room by looking down the hallway, then the museum can be constantly monitored with guards stationed at the vertices of a dominating set in the constructed graph.

A variation on this method of network monitoring was introduced by Hartnell, Rall, and Whitehead in [4]. Rather than placing one guard in each room of a dominating set, they considered having a single guard (or 'watchman') walk around the museum, beginning and ending in the same room, in such a way that the visited rooms form a dominating set. Thus after one complete walk every room has been either visited by the guard or seen by the guard from an adjacent room. In a graph, such a *walk* (that is, an alternating sequence of vertices and edges) is said to be *dominating*, and the stipulation that the walk begin and end on the same vertex means that it is *closed*. The goal is to minimize either the amount of time that vertices (rooms) are left unobserved or the length of the watchman's walk. The latter has been called the **Watchman's Walk Problem**: given a graph G, find a *minimum closed dominating walk* (MCDW) in G. The length of a MCDW in a graph G is denoted by $w_1(G)$ (where the 1 refers to the number of watchmen on the graph.)

Results for the original watchman's walk problem can be found in [4] and [5]. In the present paper we consider the variant of the problem dealing with the maximum time any vertex is left unguarded, first explored in detail in [2]. Here we have more than one guard available, but we wish to ensure that no vertex remains unobserved for more than some fixed time t. Our goal becomes to respect this time constraint while using as few guards as possible.

2 Preliminaries

We consider only *simple* graphs with no loops and no multiple edges. Recall that a *spanning tree* of a graph G is a tree containing all vertices of G and a subset of the edges of G. We refer to vertices of degree 1 in a tree as *leaves* and to a leaf's single neighbouring vertex as a *stem*. If T is a tree then $L(T)$ denotes the set of leaves of T, and we use T_0 to denote the leaf-deleted subtree $T \setminus L(T)$.

We denote the closed neighbourhood of a vertex $v \in V(G)$ by $N[v]$. We say a vertex u is *unobserved* if no vertex in $N[u]$ is occupied by a guard. Hence for a given graph G and length of time t, we are interested in finding the minimum number of guards needed to dominate the graph such that no vertex is unobserved for more than t consecutive units of time. More formally, for fixed time $t \in \mathbf{N}$, a graph G can be *t-monitored* by a set S of guards if there exists a function $f : S \times \mathbf{N} \to V(G)$ such that

(i) for every guard $g \in S$ and at every time $\tau \in \mathbf{N}$, $f(g, \tau + 1) \in N[f(g, \tau)]$, and

(ii) for every vertex $v \in V(G)$ and every interval $I \subset \mathbf{N}$ of length $t+1$, there exists a guard $g \in S$ and a time $\tau \in I$ such that $f(g, \tau) \in N[v]$.

Note that $f(g, \tau)$ is the vertex occupied by the guard g at time τ. Essentially, condition (i) ensures that in one unit of time, guards may move from a vertex to one of its neighbours (i.e., no 'jumping' is allowed), and condition (ii) ensures that every vertex has a guard within its closed neighbourhood at least once every $t+1$ units of time.

Although there is no such stipulation in the definition of t-monitoring, we may assume that each guard traverses a closed walk. To see this, let G be a graph t-monitored by guards and suppose one or more of these guards traverse a walk W that is not closed. At any fixed point in time, label a vertex 0^* if it is currently occupied by a guard, label a vertex 0 if it is unoccupied but adjacent to a vertex with a guard, and label every other vertex with a positive integer (at most t) according to the length of time that has elapsed since the vertex was last observed. Since both t and $|V(G)|$ are finite, there are only finitely many such labellings, and so at some point a vertex labelling will be repeated. When this happens, we can truncate W and have it repeat the sequence of vertices and edges that followed the first occurrence of that labelling. The new walk is closed and does not disrupt the t-monitoring of G. This proves the following theorem.

Theorem 1. *If a graph G can be t-monitored by m guards then G can be t-monitored by m guards whose walks are closed.*

For a given graph G and length of time t, denote by $W_t(G)$ the minimum value of $|S|$, the number of guards needed to t-monitor a graph. Note that $W_0(G) = \gamma(G)$, since if vertices cannot be unobserved for even a single unit of time then the guards must dominate all vertices while remaining stationary. When $t > 0$, efficient strategies for multiple guards are not clear for graphs in general, or even for an arbitrary tree. For trees that are paths, however, the strategy is simple. First, place one guard at an end of the path, beginning at a stem. If this guard is to ensure that the adjacent leaf is observed at least once every $t+1$ units of time then his walk must have length at most $t+1$ (or at most t, when t is even, since closed walks are necessarily of even length). Such a walk includes $\frac{t+1}{2}$ (or $\frac{t}{2}$) edges and $\frac{t+1}{2}+1$ (or $\frac{t}{2}+1$) vertices, and additionally the guard can monitor the two vertices adjacent to the ends of his walk. A second guard can in the same way monitor the next $\frac{t+1}{2}+3$ ($\frac{t}{2}+3$) vertices on the path. If we divide the total number of vertices n by the number of vertices monitored by each guard, taking the ceiling of this value if necessary, then we find that the minimum number of guards required to t-monitor a path is as presented in Theorem 2 below.

Theorem 2. *If T is a path on n vertices then*

$$W_t(T) = \begin{cases} \left\lceil \frac{2n}{t+7} \right\rceil & \text{when } t \text{ is odd,} \\ \left\lceil \frac{2n}{t+6} \right\rceil & \text{when } t \text{ is even.} \end{cases}$$

For any graph G, if m guards can (minimally) monitor G such that no vertex is unobserved for more than t units of time, then with those m guards no vertex is being unobserved for more than $t+1$ units of time; that is, $W_t(G) \geq W_{t+1}(G)$, and we have the following lemma.

Lemma 3. *For any graph G, $W_0(G) \geq W_1(G) \geq W_2(G) \geq \ldots$.*

It was shown in [4] that cut vertices in a graph must be visited by any MCDW and that vertices of degree 1 are never visited in such a walk. Theorem 4 follows immediately. For a given tree T, we denote by $2T_0$ the multigraph formed by doubling every edge of T_0; note that such a multigraph is necessarily Eulerian.

Theorem 4. [4] *If T is a tree then $w_1(T) = 2|E(T_0)|$, and an Eulerian circuit in the tree $2T_0$ is a MCDW for T.*

This extends to the following result for multiple guards.

Corollary 5. *For any tree T,*

$$W_t(T) \leq \left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil.$$

Proof. From Theorem 4 we know a MCDW in T is obtained by doubling every edge of T_0 and walking an Eulerian circuit W in the multigraph $2T_0$. Place guards at most distance $t+1$ apart on W and have them follow one another around the Eulerian circuit; this ensures no vertex is unobserved for more than t consecutive units of time. Since the total length of the circuit is $2|E(T_0)|$, the number of guards required to place one at least at every $(t+1)$th position is

$$\left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil.$$

Since T can be t-monitored by $\left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil$ guards, the minimum number of guards required is at most this number. □

It will be useful to know when a single guard is enough to t-monitor an arbitrary tree. Since every tree has at least two leaves, $|E(T_0)| \leq (n-1)-2$ for any tree T. Using Corollary 5, $\left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil \leq \left\lceil \frac{2(n-3)}{t+1} \right\rceil = 1$ if $2n-6 \leq t+1$. This gives the following corollary.

Corollary 6. *If T is a tree on $\frac{t+7}{2}$ or fewer vertices then T can be t-monitored with one guard.*

In these last two corollaries, our strategy is to first find a single minimum closed dominating walk and then have our guards 'share' this walk by spacing them along it as equally as possibly. Contrast this with our strategy for t-monitoring paths; in some sense paths form the opposite extreme, as each guard has his own closed walk which is disjoint from all others. For a general tree the problem of determining a strategy for multiple guards really amounts to determining which parts of the tree should be monitored by shared walks and which should be 'split' into disjoint walks. We must deal with this question repeatedly in the proof of our main result.

In [2], the authors prove the following bounds on $W_t(T)$, which we generalize in the next section.

Theorem 7. [2] *For any tree T with $n \geq 3$, $W_1(T) \leq \left\lfloor \frac{n-1}{2} \right\rfloor$, $W_2(T) \leq \left\lfloor \frac{2n}{5} \right\rfloor$, and $W_3(T) \leq \left\lfloor \frac{n}{3} \right\rfloor$.*

In each case there exist trees or families of trees attaining the upper bounds listed above. Those trees for which $W_1(T) = \lfloor \frac{n-1}{2} \rfloor$ are completely characterized in [2].

3 Main result

In this section we present generalized bounds for even and odd time t. The bulk of the work is done in Theorem 9; however, we first require the following lemma.

Lemma 8. *Suppose a tree T is dominated by guards sharing a single closed walk (that is, an Eulerian circuit in $2T_0$). If j vertices are attached to T in such a way that the resulting graph is still a tree then the existing guards can dominate the new vertices by adding at most $2j$ edges to their closed walk.*

Proof. We prove the result by induction on j. If only one vertex u is to be added then attach u to T at v. Either v is a non-leaf in T and must already be included in the original walk, in which case u is seen without any modification, or v is a leaf in T and the original walk must include the stem of v; in this case u can be seen by adding the edge from the stem to v, once in each direction. Here $j = 1$ vertex was added and at most $2 \cdot j = 2$ additional edges were required.

Assume the lemma holds for all k where $1 \leq k < j$, and suppose we add j vertices to T. Remove one of these, a leaf ℓ; by the induction hypothesis the guards can observe the $j-1$ additional vertices by increasing the length of their shared walk by at most $2(j-1)$. This walk must include either the stem s of ℓ or a neighbour of s. If s is already on the walk then no extra edges need to be traversed to monitor ℓ; if only a neighbour of s is on the walk then one additional edge must be traversed (twice) to get to s and back, thereby increasing the length of the walk by 2 edges. In total all j additional vertices can be dominated by adding at most $2(j-1)+2 = 2j$ edges to the shared closed walk. ☐

An upper bound for $W_t(T)$ with odd t will follow as a consequence of Theorem 9 below. Notice that the induction hypothesis includes two conditions on the structure of guards' walks; these conditions allow us to more easily obtain the proposed bound during the inductive step.

Theorem 9. *If t is odd then any tree T on $n \geq 3$ vertices can be t-monitored by $\left\lfloor \frac{2n+t-3}{t+3} \right\rfloor$ guards such that*
(1) the closed walks of any two guards are either identical or edge-disjoint, and
(2) a closed walk shared by $p \geq 1$ guards has length at most $p(t+1)$.

Proof. Let $k = \frac{t+3}{2}$, so that $\left\lfloor \frac{2n+t-3}{t+3} \right\rfloor = \left\lfloor \frac{n+k-3}{k} \right\rfloor$. It is easy to verify the theorem when $3 \leq n \leq k+2$; in this case

$$1 = \frac{3+k-3}{k} \leq \frac{n+k-3}{k} \leq \frac{(k+2)+k-3}{k} = \frac{2k-1}{k} < 2,$$

so $\left\lfloor \frac{n+k-3}{k} \right\rfloor = 1$ and we must show any tree T on n vertices can be t-monitored with *one* guard satisfying properties (1) and (2). Since T has at most $k+2 = \frac{t+3}{2} + 2$ vertices, one guard can t-monitor T by following an Eulerian circuit through T_0 with doubled edges, by Corollary 6. This also demonstrates property (2), and since only one guard is involved, property (1) follows trivially.

Assume inductively that any tree on m vertices, $3 \leq m \leq n-1$, can be t-monitored by $\left\lfloor \frac{m-3+k}{k} \right\rfloor$ guards whose walks satisfy properties (1) and (2). Let T be an arbitrary tree on n vertices. We will find k suitable vertices to remove from T, forming a subtree T' that by the induction hypothesis can be t-monitored by $\left\lfloor \frac{(n-k)+k-3}{k} \right\rfloor = \left\lfloor \frac{n-3}{k} \right\rfloor$ guards. If we can show that including the k vertices requires only one additional guard, whose walk preserves properties (1) and (2), then T can be t-monitored by $\left\lfloor \frac{n-3}{k} \right\rfloor + 1 = \left\lfloor \frac{n+k-3}{k} \right\rfloor$ guards and the theorem will be proved by induction. We select the k vertices as follows.

Find a non-leaf vertex v_0 such that $T \setminus v_0$ has at least one component with more than k vertices, and let S_1 be one such component. Let v_1 be the vertex in S_1 that is adjacent to v_0 in T, and root S_1 at v_1. If all branches of v_1 in S_1 have less than k vertices, relabel v_1 as v; otherwise, choose a branch with k or more vertices and call it S_2. Root S_2 at v_2, the vertex adjacent to v_1. If all branches of v_2 in S_2 have less than k vertices, relabel v_2 as v; otherwise, choose a branch with k or more vertices and call it S_3. We can repeat this procedure until eventually a vertex $v = v_i$ is found whose branches are all of size less than k. Furthermore, the subtree S_i (containing v_i and these branches) has at least k vertices, since S_i was chosen from the branches of v_{i-1} with precisely that property. Select k vertices from S_i beginning in one branch of v, ensuring that after each selection the unselected vertices are connected, and selecting from a second branch only after the first has been entirely selected, selecting from a third branch only if the second has been entirely selected, and so on.

If S_i has exactly k vertices then we select the entire component, including v, and let the tree T' be T with these k vertices removed. Then S_i is a tree with $k = \frac{t+3}{2}$ vertices, so by Corollary 6 it can be t-monitored by one guard. By the induction hypothesis the tree T' on $n-k$ vertices can be t-monitored by $\left\lfloor \frac{n-3}{k} \right\rfloor$ guards, whose walks satisfy properties (1) and (2). The one additional guard required for S_i gives a total of $\left\lfloor \frac{n-3}{k} \right\rfloor + 1 = \left\lfloor \frac{n+k-3}{k} \right\rfloor$ guards, and since the new guard does not enter T' and walks at most $2(k-2) = t-1$ edges, his walk does not violate properties (1) and (2).

If S_i has more than k vertices then we will not select the vertex v. At least one branch of S_i is entirely selected, since no single branch contains k or more vertices, and at most one branch is partially selected, since we are choosing the vertices one branch at a time. Let S be the subtree of T containing v and all completely selected branches. There are two cases.

Case 1: There is no partially selected branch; that is, S contains all k selected vertices. Since v is not selected, S is a subtree with $k+1$ vertices, at least two of which are leaves in T (S cannot have only a single branch because each branch of v has less than k vertices). Hence by Corollary 6, S can be t-monitored by one guard who traverses at most $2(k-2) = t-1$ edges. By the same reasoning as used above when $|V(S_i)| = k$, the theorem holds in this case.

Case 2: There is a partially-selected branch; call this branch B. Let T' be T with the k selected vertices removed. Since only some of the vertices of B are selected, part of this branch will be in the tree T'. By the induction hypothesis, T' on $n-k$ vertices can be t-monitored by $\lfloor \frac{n-3}{k} \rfloor$ guards whose walks are identical or edge-disjoint, where a closed walk shared by p guards has length at most $p(t+1)$. In the following two sub-cases, let B' be the branch B of v contained in T'.

Case 2a: The edges of B' belong to multiple edge-disjoint walks. At most one of these walks includes edges outside of B', because a single edge joins B' to the rest of T'. So at least one walk has edges only in B', and consequently at least one guard never leaves B in T'. Let one guard remain in B' and let any other guards whose walks were entirely in B' sit stationary at the vertex v. If there was a walk in B' involving edges outside of B', truncate this walk by excluding all edges of B'. The guard(s) on this walk can remain at v instead of entering B', and can resume the remainder of the walk at the appropriate time. These alterations do not affect the monitoring of any vertex in $T' \setminus B'$, and properties (1) and (2) are clearly not violated. The single guard remaining in B' can t-monitor all of the branch B in T, by Corollary 6, because B has less than $k = \frac{t+3}{2}$ vertices. This walk is edge-disjoint from all others and has length at most $2[(k-2)-1] = t-3$ (because B has at most $k-2$ edges and at least one vertex which is a leaf in T).

We now need to dominate the remainder of the k selected vertices, which are in S. Because S has at most $k = \frac{t+3}{2}$ vertices, one new guard can t-monitor S with a walk which is edge-disjoint from all others, because S and T' have only the vertex v in common. Thus property (1) is preserved. The walk has one guard and length at most $2[(k-1)-1] = t-1$ (S has at most $k-1$ edges and at least one leaf), so property (2) is also preserved, and we see that the theorem holds in this case.

Case 2b: The edges of B' belong to a single walk. If this walk does not include edges outside of B' then its guard(s) can monitor all of B in T, as in case 2a, so the result follows as above. Otherwise, we have a single closed walk W which necessarily visits the vertex v. Suppose p guards share W, so that W has length at most $p(t+1)$, and suppose $|B \setminus B'| = j$ (i.e., j vertices of B were selected). By Lemma 8, at most $2j$ edges must be added to the walk W in order for the guards on W to also dominate the j vertices of B that are not in T'.

Now, if j of the k selected vertices are in B then S has $k-j+1$ vertices, including v. At least one of these is a leaf of T, so at most $2(k-j-1) = t+1-2j$ edges must be traversed to visit all non-leaf vertices of S, including v. If we also add these edges to the walk W, which visits v, then in total we have a closed walk of length at most $p(t+1)+(2j)+(t+1-2j) = (p+1)(t+1)$. We can therefore place one additional guard on the expanded walk, so that $p+1$ guards are now spaced along it as equally as possible. Every vertex on or adjacent to this walk, including each of the k selected vertices, is then seen at least once every t units of time. Thus with one new guard and the described additions to W, the entire tree T can be t-monitored with $\left\lfloor \frac{n+k-3}{k} \right\rfloor$ guards; since the new guard is joining a walk of length at most $(p+1)(t+1)$ shared by $p+1$ guards, properties (1) and (2) are preserved. Thus the theorem holds in all cases. □

We now have as an immediate corollary the following upper bound for odd t.

Corollary 10. *If T is a tree of order n and $t > 0$ is an odd integer then*

$$W_t(T) \leq \left\lfloor \frac{2n+t-3}{t+3} \right\rfloor = \left\lfloor \frac{n+k-3}{k} \right\rfloor,$$

for $k = \frac{t+3}{2}$.

Recall from Lemma 3 that $W_t(G) \leq W_{t-1}(G)$ for any time t and any graph G. If t is even then from this inequality and Corollary 10 we have $W_t(G) \leq W_{t-1}(G) \leq \left\lfloor \frac{2n+(t-1)-3}{(t-1)+3} \right\rfloor$. The resulting upper bound for even t, presented below in Corollary 11, is notably weaker than the bound for odd t.

Corollary 11. *If G is a connected graph of order n and $t > 0$ is an even integer then*

$$W_t(G) \leq \left\lfloor \frac{2n+t-4}{t+2} \right\rfloor.$$

4 Analysis of the bound

In this section we explore the utility of Corollary 10 and construct a family of trees for which this upper bound is attained. Note that for $t = 1$ and $t = 3$, the upper bound presented here matches those found by [2].

We compare the bound of Corollary 5 with that of Corollary 10. The former has the disadvantage of requiring some specific knowledge of the tree, besides its order: namely, we need to know the number of non-leaf edges in the tree. If we know nothing of the graph we could assume only that T has at least two leaves, thus giving $|E(T_0)| \leq n-3$. With only this assumption, how does the bound of $\left\lfloor \frac{2n+t-3}{t+3} \right\rfloor$ compare to $\left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil$?

It is straightforward to check that $\dfrac{2n+t-3}{t+3} < \dfrac{2(n-3)}{t+1}$ if and only if $t^2+4t+15 < 4n$. For $n \geq 10$ and $t \leq \sqrt{n}$ we see that $t^2+4t+15 < n+4\sqrt{n}+15 < n+\frac{3n}{2}+\frac{3n}{2} = 4n$.

Hence with these conditions, and without knowing $|L(T)|$, the bound of Corollary 10 is strictly better than the bound of Corollary 5. Of course, to assume T has only 2 leaves is rather strong. Suppose we allow for T having up to $\sqrt{n}$ leaves. Let $L = |L(T)|$; then $\frac{2n+t-3}{t+3} < \frac{2(n-1-L)}{t+1}$ if and only if $t^2+2tL+6L+3 < 4n$, which is true if $n > 45$ and $t, L \leq \sqrt{n}$. Thus for reasonable assumptions on n, $|L(T)|$, and t, the bound $W_t(T) \leq \left\lfloor \frac{2n+t-3}{t+3} \right\rfloor$ is an improvement upon the bound $W_t(T) \leq \left\lceil \frac{2|E(T_0)|}{t+1} \right\rceil$.

A final point for the strength of Corollary 10 is that the bound is sharp in its current form; specifically, if the denominator is kept as $k = \frac{t+3}{2}$ then the numerator cannot be reduced. Theorem 12 partly demonstrates this, by identifying an infinite family of trees for which the bound of Corollary 10 is attained. To prove the sharpness claim, however, we have to further demonstrate that there is a tree attaining the bound for given t which also satisfies $\left\lfloor \frac{n+k-4}{k} \right\rfloor < \left\lfloor \frac{n+k-3}{k} \right\rfloor$; the tree T in Figure 1 suffices, since as a tree in the family of Theorem 12 we know $W_3(T) = \left\lfloor \frac{n+k-3}{k} \right\rfloor = 3$, but with $n = 9$ and $k = t = 3$ we have $\left\lfloor \frac{n+k-4}{k} \right\rfloor = 2$.

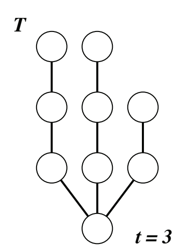

Figure 1: A tree T with $W_3(T) = 3 > \left\lfloor \frac{n+k-4}{k} \right\rfloor$.

Theorem 12. *Let t be an odd integer, let $k = \frac{t+3}{2}$, and let j be an integer satisfying $2 \leq j \leq k+1$. If T is formed by identifying one endpoint of any number of $(k+1)$-vertex paths and a single $(j+1)$-vertex path, then $W_t(T) = \left\lfloor \frac{2n+t-3}{t+3} \right\rfloor = \left\lfloor \frac{n+k-3}{k} \right\rfloor$.*

Proof. Let T be formed as described (see Figure 2), say with $m+1$ branches off a central vertex v. Fix a set of t-monitoring walks on T. With some interchanging of guards – but without any alteration in the sequence of guard-occupied vertices – we show that at least $m+1$ guards are involved in any t-monitoring of T.

Let B be a k-vertex branch with stem s. Note that the length of B (and all k-vertex branches) is such that it takes exactly $t+1$ units of time to get from the stem s to the central vertex v and back again to s. We claim that (possibly with some switching of guards) there is a guard who never leaves $B \cup \{v\}$. At some point in time, a guard g_1 visits s in order to monitor the neighbouring leaf. If the guard never leaves $B \cup \{v\}$ then we are done; otherwise, let g_2 be the next guard that visits

s, as there is not enough time for g_1 to walk from s to v, to a vertex off of B, and back again to s. For this same reason, the guards g_1 and g_2 meet at the latest at v, or else too much time passes between consecutive visits of s.

Now, at the meeting point of the two guards g_1 and g_2, have them swap places with one another: let each turn around and resume the other's walk from that time on. We do not change the timing or the sequence of edges traversed; indeed, at any given time, the set of occupied vertices of T is exactly as it was before, and so T is still t-monitored with the same number of guards. If we repeat this swapping strategy again when (if) a third guard enters B to visit s, and continue for each that follows, then the guard g_1 never leaves $B \cup \{v\}$. Applying the same argument to all other k-vertex branches, we find each branch has a unique guard who is restricted to that single branch (and possibly v). This accounts for at least m guards; but none of these enters the j-vertex branch, the leaf of which cannot be seen from v because $j \geq 2$. Thus there is at least one additional guard walking the tree.

We have shown that at least $m+1$ guards are required to t-monitor T. Since $j \leq k+1$, one guard is enough to monitor the entire j-vertex branch; so $m+1$ guards (one on each branch) are in fact sufficient to t-monitor T. That is, $W_t(T) = m+1$.

It remains to show that $\lfloor \frac{n+k-3}{k} \rfloor = m+1$ for this tree. We know that $|V(T)| = n = mk + j + 1$, so we have

$$\left\lfloor \frac{n+k-3}{k} \right\rfloor = \left\lfloor \frac{mk+k+j-2}{k} \right\rfloor = \left\lfloor \frac{(m+1)k}{k} + \frac{j-2}{k} \right\rfloor = m+1,$$

since $j \leq k+1 \;\Rightarrow\; \frac{j-2}{k} < 1$. Hence, these trees attain the bound of Corollary 10. □

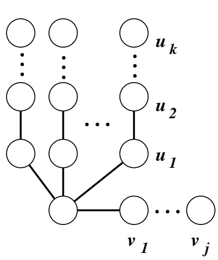

Figure 2: An infinite family of trees that attain the upper bound of Corollary 10, where t is odd, $k = \frac{t+3}{2}$, and $2 \leq j \leq k+1$.

5 Conclusions and open questions

Although our upper bounds are presented for trees, they are also valid for graphs in general. This is because the number of guards required for a general graph is always bounded above by the number of guards required for any spanning tree of the graph. This bound can be optimized by picking the "best" spanning tree – for our purposes, the spanning tree with as many leaves as possible, a problem which is known to be NP-hard [3].

The following question repeatedly arises when we are monitoring with multiple guards: if m guards can t-monitor a tree, can they do so with closed walks that are pairwise edge-disjoint or identical? We know the answer is yes for certain families of graphs (e.g., those described in Theorem 12), and we have never seen this restriction increase the required number of guards.

Another natural question is whether or not there exists a result analogous to Theorem 9 for *even* values of t. We have the following conjecture.

Conjecture 13. *If T is a tree of order n and $t > 0$ is an even integer then*

$$W_t(G) \leq \left\lfloor \frac{2n + t - 2}{t + 3} \right\rfloor .$$

This has been shown to hold for $t = 2$ [2] and $t = 4$ [7], using induction as in the proof of Theorem 9. However, when we analyze cases as we did for that theorem, there are certain instances when a general argument is not apparent; the proofs for $t = 2, 4$ rely instead on exhaustive checking of subcases. We note two reasons for the problems that arise when t is even. Firstly, our approach of removing $k = \frac{t+3}{2}$ vertices in the inductive step is no longer possible, since this k is not an integer. There are various ways to avoid this, but even the most practical — namely, setting $k = \frac{t+4}{2}$ and proceeding as with odd t — has the disadvantage that a relatively larger number of vertices are removed, making it harder for only one additional guard to monitor them.

Secondly, and more generally, a subtle problem arises when t is even that is unrelated to our method of proof. As mentioned in the discussion leading to Theorem 2, a guard who is solely responsible for the vertices he dominates must have a walk of length at most $t + 1$; but if t is even then this value is odd, and since any closed walk on a tree has even length, the maximum length is actually only t. In certain cases of our attempted proof, this loss means that a single guard is unable to monitor the desired number of vertices.

The authors are pleased to hear that as of the writing of this paper, Conjecture 13 has been verified [1].

References

[1] I. Beaton, R. Begin, S. Finbow and C.M. van Bommell, Even time constraints on the watchman's walk, Manuscript (2011), St. Francis Xavier University.

[2] C. Davies, S. Finbow, B. Hartnell, Q. Li and K. Schmeisser, The watchman problem on trees, *Bull. ICA* 37 (2003), 35–43.

[3] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco (1979).

[4] B.L. Hartnell, D.F. Rall and C.A. Whitehead, The watchman's walk problem: An introduction, *Congr. Numer.* 130 (1998), 149–155.

[5] B.L. Hartnell and C.A Whitehead, Minimum dominating walks in Cartesian product graphs, *Util. Math.* 65 (2004), 73–82.

[6] T.W. Haynes, S.T. Hedetniemi and P.J. Slater, *Fundamentals of domination in graphs*, Marcel Dekker, New York (1998).

[7] R. Keeping, *The watchman's walk problem and its variations*, M.Sc. Thesis, Memorial University of Newfoundland (2010). (Written under author's maiden name; now R. Milley.)

(Received 16 June 2011; revised 14 Feb 2012)