Eternal Dominating Sets in Graphs

William F. Klostermeyer† and Gary MacGillivray‡*

† School of Computing
University of North Florida
Jacksonville, FL 32224-2669

‡ Dept. of Mathematics and Statistics
University of Victoria
Victoria, Canada

Abstract

Results are presented on the *eternal domination* problem: defending a graph from an infinite sequence of attacks, so that each attack is defended by a guard at most distance one from the attack. We first consider the model where at most one guard moves to defend an attack. Our focus is on the relationship between the number of guards and the independence and clique covering numbers of the graph. We establish results concerning which triples of these parameters can be attained by some graph, and determine the exact value of the number of guards for graphs in certain classes. We then turn our attention to the variant of the problem in which every guard can relocate to an adjacent vertex in defence of an attack. We give a linear algorithm to determine the minimum number of guards necessary to defend a tree, and use it to answer another question about defending trees.

1 Introduction

Several recent papers have studied the problem of protecting the vertices in a graph from a series of one or more attacks [1, 2, 5]. In such a problem, guards are located at vertices, can protect the vertices at which they are

*Research supported by NSERC.

located, and can move to a neighboring vertex to defend an attack there. Under this set of rules, a guard located at each vertex of a dominating set can defend a graph against a single attack. Several variations of this problem have been proposed including Roman Domination [3], Weak Roman Domination [4] and k-secure sets/eternal secure sets [1, 2, 5, 7].

Let $R = r_1, r_2, \ldots$ be a sequence of vertices of a graph. The elements of R are the locations of a sequence of consecutive attacks (or service requests) at vertices, each of which must be defended (or attended) by a guard. We consider the model where at most one guard is allowed to move to defend each attack, except in Section 4 where we consider the model in which all guards can move in response to an attack. In order to avoid confusion, we state only the definitions for the "one guard moves" model here and defer the definitions for the other model to Section 4.

A set D is an *eternal secure set* if, for all possible sequences of attacks $R = r_1, r_2, \ldots$ there exists a sequence $D = D_0, D_1, \ldots$ of dominating sets such that $D_i = (D_{i-1} \setminus \{v\}) \cup \{r_i\}$, where $v \in D_{i-1}$ and $r_i \in N[v] = N(v) \cup \{v\}$ (note that $v = r_i$ is possible). The set D_i is the set of locations of the guards after the attack at r_i is defended. If $v \neq r_i$, we say that the guard at v has *moved* to r_i. The size of a smallest eternal secure set in G is the *eternal domination number* (also known as *eternal security number*), and is denoted by $\gamma_\infty(G)$ or simply γ_∞ [2].

For a graph G, the relationship between $\gamma_\infty(G)$, and its domination number $\gamma(G)$, independence number $\alpha(G)$, and clique covering number $\theta(G)$ (i.e., the chromatic number of $\overline{G}$) is explored in Sections 2 and 3 of this paper. The main result in Section 2 describes triples of integers $(\alpha, \gamma_\infty, \theta)$ that are achieved by some graph. Section 3 determines the eternal domination number of graphs in certain classes. In Section 4, we consider the variation of the problem in which every guard is allowed to move at each step. We solve some problems stated by Goddard et al. [5], including giving a linear-time algorithm to defend the vertices of a tree.

2 Results for Graphs

In this section and the next, at most one guard is allowed to move to defend an attack.

2.1 Structural Results

Goddard et al. [5] noted that, for all graphs G,

$$\alpha(G) \leq \gamma_{\infty}(G) \leq \theta(G).$$

Another upper bound on γ_{∞} is the following.

Theorem 1 [8] *For any graph G with independence number $\alpha(G) \geq 1$,*

$$\gamma_{\infty}(G) \leq \binom{\alpha(G)+1}{2}.$$

First, a construction which is useful for obtaining graphs with given values for these parameters is noted. Recall that the *join* of two disjoint graphs G_1 and G_2 is the graph $G_1 + G_2$ obtained from $G_1 \cup G_2$ by adding all possible edges joining vertices of G_1 to vertices of G_2. The *star product* of $n \geq 1$ disjoint graphs $G_1, G_2, \ldots, G_n$ is the graph

$$\mathcal{S}(G_1, G_2, \ldots, G_n) = K_1 + (G_1 \cup G_2 \cup \cdots \cup G_n).$$

That is, $\mathcal{S}(G_1, G_2, \ldots, G_n)$ is obtained from the disjoint union $G_1 \cup G_2 \cup \cdots \cup G_n$ by adding a new vertex and joining it to all of the other vertices.

The proof of the following proposition is similar to that of Theorem 2 in [8].

Proposition 2 *Given $n \geq 1$ disjoint graphs $G_1, G_2, \ldots, G_n$, for the graph $\mathcal{S} = \mathcal{S}(G_1, G_2, \ldots, G_n)$ we have*

- $\alpha(\mathcal{S}) = \alpha(G_1) + \alpha(G_2) + \cdots + \alpha(G_n)$,
- $\gamma_{\infty}(\mathcal{S}) = \gamma_{\infty}(G_1) + \gamma_{\infty}(G_2) + \cdots + \gamma_{\infty}(G_n)$, *and*
- $\theta(\mathcal{S}) = \theta(G_1) + \theta(G_2) + \cdots + \theta(G_n)$.

Proof: The statements regarding $\alpha(S)$ and $\theta(S)$ are obvious. We argue that $\gamma_{\infty}(S) = \gamma_{\infty}(G_1) + \gamma_{\infty}(G_2) + \cdots + \gamma_{\infty}(G_n)$. Clearly this number of guards suffices to defend $\mathcal{S}$ from any sequence of attacks. On the other hand, for $i = 1, 2, \ldots, n$ there exists a sequence of attacks at vertices of G_i which can only be defended if $\gamma_{\infty}(G_i)$ guards are located on G_i. Since, if $i \neq j$, no single move can relocate a guard at a vertex of G_i to a vertex of G_j, it follows that $\gamma_{\infty}(G_1) + \gamma_{\infty}(G_2) + \cdots + \gamma_{\infty}(G_n)$ are necessary to defend $\mathcal{S}$. □

A simple lower bound is given.

Corollary 3 [8] *For every positive integer n there exists a connected graph G_n with independence number $2n$ and eternal security number $3n$. That is, for every positive integer n there exists a graph G_n such that*

$$\frac{\gamma_\infty(G_n)}{\alpha(G_n)} \geq \frac{3}{2}.$$

A triple of positive integers (a, g, t) is called *realizable* if there exists a graph with $\alpha = a$, $\gamma_\infty = g$ and $\theta = t$. The above corollary shows that for all integers $n \geq 1$ the triple $(2n, 3n, 3n)$ is realizable. Theorem 1 shows that no triple with $g > \binom{a+1}{2}$ is realizable. Goddard et al. [5] gave an example to show that $(2, 6, 10)$ is realizable. This was the first example of a realizable triple with g strictly between a and t. It is natural to ask: *Which triples are realizable?*

Before providing a partial solution to this question, we observe that Proposition 2 implies that the set of realizable triples has a nice algebraic structure: it is closed with respect to addition.

Proposition 4 *If the $n \geq 1$ triples $(a_1, g_1, t_1), (a_2, g_2, t_2), \ldots, (a_n, g_n, t_n)$ are all realizable, then so is $(a_1 + a_2 + \cdots + a_n, g_1 + g_2 + \cdots + g_n, t_1 + t_2 + \cdots + t_n)$.*

Proof: Suppose that, for $i = 1, 2, \ldots, n$, the triple (a_i, g_i, t_i) is realized by the graph G_i. Then, by Proposition 2, the triple $(a_1 + a_2 + \cdots + a_n, g_1 + g_2 + \cdots + g_n, t_1 + t_2 + \cdots + t_n)$ is realized by $\mathcal{S}(G_1, G_2, \ldots, G_n)$. □

Taken together, the sequence of results below shows:

Theorem 5

1. *The only realizable triple with $a = 1$ is $(1, 1, 1)$.*
2. *The only realizable triples with $a = 2$ are $(2, 2, 2)$ and $(2, 3, t)$, where $t \geq 3$.*
3. *For all integers a, g and t with $3 \leq a \leq g \leq \frac{3}{2}a$ and $g \leq t$, the triple (a, g, t) is realizable.*

Statement 1 is clear. The next few results determine the realizable triples with $a = 2$.

Proposition 6 [2] *If $\theta(G) \leq 3$, then $\theta(G) = \gamma_\infty(G)$. That is, if (a, g, t) is a realizable triple with $t \leq 3$, then $g = t$.*

Corollary 7 [7] *If $\alpha(G) \leq 2$ and $\alpha(G) = \gamma_\infty(G)$, then $\alpha(G) = \theta(G)$. That is, the only realizable triple with $a = g = 2$ is $(2,2,2)$, which is realized by C_4.*

Corollary 8 *For any $t \geq 3$, the triple $(2,3,t)$ is realizable.*

Proof: The complement of a triangle-free graph of chromatic number t has $\alpha = 2$ and $\theta = t$. By Theorem 1 and Corollary 7, such a graph has $\gamma_\infty = 3$. □.

We now turn our attention to Statement 3 of Theorem 5.

Proposition 9 [2] *For any integer $a \geq 1$, the triple (a,a,a) is realized by a perfect graph with independence number a.*

Theorem 10 [7] *For any integers $a \geq 3$ and $t > a$, there exists a graph G with $\alpha(G) = \gamma_\infty(G) = a$, and $\theta = t$. That is, for any integers $a \geq 3$ and $t > a$, the triple (a,a,t) is realizable.*

Proposition 11 *Let a and g be positive integers with $a \geq 3$ and $a \leq g \leq \frac{3}{2}a$. Then, for any integer $t \geq g$ the triple (a,g,t) is realizable.*

Proof: The proof is by induction on $a+g$. The base case, $a+g=6$, is covered by Proposition 9 and Theorem 10. Suppose that the statement is true for all triples with $6 \leq a+g < n$, for some integer $n \geq 7$. Consider a triple (a,g,t) satisfying the hypotheses and having $a+g=n \geq 7$. Then $g \geq 4$. We consider several cases. If $a = g$ then the result follows from Theorem 10. The only remaining case when $a = 3$ is $g = 4$. Here, by Corollary 8, the triple (a,g,t) is the sum of the realizable triples $(2,3,t-1)$ and $(1,1,1)$. When $a = 4$, the cases $g = 5$ and $g = 6$ must be considered. If $a = 4$ and $g = 5$, then (a,g,t) is the sum of the realizable triples $(2,2,2)$ and $(2,3,t-2)$. If $a = 4$ and $g = 6$ then (a,g,t) is the sum of the realizable triples $(2,3,3)$ and $(2,3,t-3)$. The final case is $a \geq 5$ and $a < g \leq \frac{3}{2}a$. Here $(a-2)+(g-3) \geq 6$ so that, by the induction hypothesis, (a,g,t) is the sum of the realizable triples $(2,3,3)$ and $(a-2,g-3,t-3)$. The result now follows by induction. □

This completes the proof of Theorem 5. A consequence of Theorem 5 is that there can be arbitrarily large gaps between the independence, eternal security, and clique-covering numbers.

Corollary 12 *Let c and d be positive integers. Then there exists a connected graph with $\alpha + c < \gamma_\infty$ and $\gamma_\infty + d < \theta$.*

Proof: Let $\alpha = 2c + 2 \geq 3$. By Statement 3 of Theorem 5, the triple $(2c+2, 3c+3, 3c+d+4) = (\alpha, \alpha+c+1, \alpha+c+d+2)$ is realizable. □

We now turn our attention to the special situation when the domination number of a graph equals its eternal domination number. The proof of the following is straightforward and is omitted.

Lemma 13 (**Cloning Lemma**) *Let G be a graph and v a vertex of G. Let G' be the graph obtained from G by adding a new vertex v' such that $v'w \in E(G')$ for each $w \in N[v]$. Then $\gamma(G) = \gamma(G'), \gamma_\infty(G) = \gamma_\infty(G')$, and $\theta(G) = \theta(G')$.*

Theorem 14 *For any graph G, $\gamma(G) = \gamma_\infty(G)$ if and only if $\gamma(G) = \theta(G)$.*

Proof: If $\gamma(G) = \theta(G)$, then since $\gamma(G) \leq \alpha(G) \leq \gamma_\infty(G) \leq \theta(G)$ for all graphs G, we have that $\gamma(G) = \gamma_\infty(G)$.

Now suppose $\gamma(G) = \gamma_\infty(G)$. Obviously, $\gamma(G) = \alpha(G)$ and every maximal independent set is a minimum dominating set and thus also an eternal secure set. Hence every maximal independent set of G is also a maximum independent set. Let D be a maximal independent set. From the Cloning Lemma, we can assume that each vertex in D has at least one *private neighbor* (i.e., a vertex of $V - D$ that is adjacent to no other vertex of D). Otherwise, modify G accordingly so the Cloning Lemma applies. Since D is a dominating set and a maximum independent set, the private neighbors of each vertex $v \in D$ induce a clique.

Assume we have $\gamma(G)$ guards in G located at the vertices of D. Consider an attack at vertex $w \notin D$, where w is not a private neighbor of any vertex in D (if no such w exists, then we are done). Then w is adjacent to some vertex $v \in D$. Send the guard from v to w. Since $\gamma_\infty(G) = |D|$ and since no private neighbor of v is adjacent to any vertex in D except v, it must be that w is adjacent to each private neighbor of v. It follows that $\theta(G) = |D|$. □

2.2 Complexity

When $\alpha(G) = 2$, $\gamma_\infty(G) \leq 3$ [5], and one can determine the eternal security number of such graphs in polynomial time, due to Theorem 5 of [2]. We can say more than this, however.

Theorem 15 *The eternal security number of graphs of bounded independence number can be computed in polynomial time.*

Proof: Suppose $\gamma_\infty(G)$ is bounded by the constant $B \leq \binom{\alpha(G)+1}{2}$. Let $n = |V(G)|$.

Let i be an integer such that $\alpha \leq i \leq B$. Construct an arc-labelled digraph X_i as follows: The vertices of X_i are the dominating subsets of $V(G)$ of size i. There is an arc from W to Z labelled r if whenever the guards are at W and the attack is at r then the attack can be defended by the guards moving to Z. Note that each vertex may be incident with several arcs of the same label, and that loops are allowed. The digraph X_i can be constructed in time polynomial in n, as it has at most n^i vertices (since B is a constant).

Since the digraph X_i captures every possible guards' strategy, it is clear that the eternal security number of G is at most i if and only if X_i has an induced subgraph in which each vertex has an arc of each label going out from it.

It can be determined iteratively whether X_i has a subgraph of the required type. Repeatedly delete any vertex with no incoming arcs, or with no outgoing arc of some label. If the process terminates with an empty digraph, then the answer is no. If the process terminates with a non-empty digraph, then by definition of the reduction this digraph has the required property. Since at least one vertex is removed at each step, and the number of vertices is polynomial in n, the procedure takes polynomial time.

In order to determine the eternal security number, we need to find the smallest i such that $\alpha \leq i \leq B$ and applying the above reduction to X_i does not result in an empty digraph. This can be done by trying the values $i = \alpha, i = \alpha + 1$, etc. in turn until either a positive answer is obtained, or $i = B - 1$ has been tried. In the latter case, $\gamma_\infty = B$. Since B is a constant, this procedure takes time polynomial in n. □

3 Graph Classes

In this section, we focus on certain classes of graphs, as suggested in [2]. Since perfect graphs have $\alpha = \theta$, one should consider classes that include graphs that are not perfect, such as the following.

Let C_n^k be the k^{th} power of the cycle on n vertices. We assume that $2k+1 < n$. It was shown in [2] that $\gamma_\infty(C_n) = \lceil \frac{n}{2} \rceil = \theta(C_n)$. Observe that $\gamma(C_n^k) = \lceil \frac{n}{2k+1} \rceil$, $\alpha(C_n^k) = \lfloor \frac{n}{k+1} \rfloor$, and $\theta(C_n^k) = \lceil \frac{n}{k+1} \rceil$.

Theorem 16 *Let C_n^k be the k^{th} power of the cycle on n vertices. Then*

$\gamma_\infty(C_n^k) = \theta(C_n^k)$, *for all* $k \geq 1, n \geq 3$.

Proof: From the information above, the only case we need to consider is when $\alpha(C_n^k) + 1 = \theta(C_n^k)$, which occurs when $k+1$ does not divide n. Number the vertices $v_0, v_1, v_2 \ldots, v_{n-1}$ clockwise around the cycle in the obvious way. Suppose there are $\alpha(C_n^k)$ guards at independent vertices $v_0, v_{k+1}, v_{2(k+1)}, \ldots$. Let p be the maximum integer such that a guard resides at v_p (so p is the maximum integer such that $p = q(k+1) < n-k$, for some $q \geq 1$). Note that no guard is on $v_{n-1}, v_{n-2}, \ldots, v_{n-k-1}$ (i.e., $k+1$ consecutive vertices). Consider an attack at v_1. If defended by the guard at v_0, we proceed as follows. Have an attack at v_{p-1}. If defended by the guard at v_p, we are done (now, no guard is adjacent to v_{n-k}), so this attack must be defended by the guard at $v_{p-(k+1)}$. Continuing in this way, consider a sequence of similar attacks so that each guard starting with the closest counterclockwise guard to $v_{p-(k+1)}$ and proceeding counterclockwise must be forced to move, eventually with the guard at v_1 moving to v_k. At this point, no guard will be adjacent to v_{n-2}.

On the other hand, suppose the first attack is defended by the guard at v_{k+1}. But now we have a similar configuration to one above, with two guards at consecutively numbered vertices and a "gap" of $k+2$ consecutively numbered vertices (starting just after the two consecutively numbered vertices with guards), none of which have a guard. □

We now consider the complements of powers of cycles, denoted $\overline{C_n^k}$, with $2k+1 < n$. These are, in fact, the *circular cliques* which are central to the theory of circular colorings of graphs. It is well known, and not difficult to prove, that $\overline{C_n^k}$ has independence number $k+1$. We will also make use of the easily verified observation that $\theta(\overline{C_n^k}) = k+1$ if $k+1|n$ and $k+2$ otherwise.

Theorem 17 *Let* $\overline{C_n^k}$ *be complement of the* k^{th} *power of the cycle on* n *vertices. Then* $\gamma_\infty(\overline{C_n^k}) = \theta(\overline{C_n^k})$, *for all* $k \geq 1, n \geq 3$.

Proof: Number the vertices of $\overline{C_n^k}$ $v_0, v_2, \ldots, v_{n-1}$ "around the cycle" (i.e., so $v_i v_{i+1}$ is an edge in the graph's complement, C_n^k). Assume $k+1$ does not divide n (else the theorem is obvious) and suppose by way of contradiction that $\gamma_\infty(\overline{C_n^k}) = k+1 < \theta(\overline{C_n^k})$. Assume there are attacks so that the $k+1$ guards are located at the $k+1$ independent vertices $v_1, v_2, \ldots, v_{k+1}$. Now consider an attack at v_{k+3}, which can be defended by the guard at either v_1 or v_2. If this attack is defended by the guard at v_1, then after the guard defends the attack, there is no guard adjacent to v_{k+2}. On the other hand, suppose the attack at v_{k+3} is defended by the guard at v_2. After the guard at v_2 moves to v_{k+3}, consider $k-1$ subsequent attacks at $v_{k+4}, v_{k+5}, \ldots$. Each of these attack can only be defended by a unique guard, the guard at $v_3, v_4, \ldots$, respectively, and once these attacks are defended, there is no guard adjacent to v_n. □

Lemma 18 *If G is an induced subgraph of H then $\gamma_\infty(H) \geq \gamma_\infty(G)$.*

Proof: After a sequence of attacks on vertices of G that requires $\gamma_\infty(G)$ guards to defend there must be $\gamma_\infty(G)$ guards located on vertices of G. □

A *cactus* graph is a graph in which each edge is contained in at most one cycle. Every block of such a graph is either K_2 or a cycle, and any induced subgraph of a cactus graph is also a cactus graph.

Theorem 19 *Let G be a cactus graph. Then $\gamma_\infty(G) = \theta(G)$.*

Proof: The proof is by induction on the number of vertices. The theorem is obviously true for trees (which are perfect) and cycles. Partition the edge set of G into two cactus graphs, H_1, H_2 having one vertex, v, in common; so v is a cut-vertex. By choosing H_1 to correspond to an end-block in the block cutpoint tree, this can be done so that H_1 is either a K_2 or a cycle.

Suppose first that H_1 is a copy of K_2 with vertex set $\{v, y\}$. There are two cases to consider, depending on the relationship of $\theta(G)$ to $\theta(H_2)$.

If $\theta(G) = \theta(H_2)$, then by Lemma 18 we have $\gamma_\infty(G) \geq \gamma_\infty(H_2) = \theta(H_2) = \theta(G)$, so that $\gamma_\infty(G) = \theta(G)$. Hence, suppose that $\theta(G) = \theta(H_2)+1$. First, attack at y so that a guard is located there. Now it follows that the subgraph, H, induced by $V(G) - \{v, y\}$ has $\theta(H) = \theta(H_2)$, and consequently $\gamma_\infty(H) = \theta(H) = \theta(H_2)$. Since a guard located at y can not defend a vertex of H, at least $\theta(H)$ more guards are required in order to defend G. Therefore, $\gamma_\infty(G) \geq 1 + \theta(H) = 1 + \theta(H_2) = \theta(G)$, so that $\gamma_\infty(G) = \theta(G)$.

Now suppose that H_1 is a cycle. Let P be the path induced by $V(H_1) - N[v]$. There are two cases to consider, depending on the relationship of $\theta(H_2)$ to $\theta(H_2 - v)$.

Suppose $\theta(H_2 - v) = \theta(H_2)$. Then $\gamma_\infty(H_2 - v) = \theta(H_2 - v) = \theta(H_2)$. Thus there is a sequence of attacks that results on $\theta(H_2)$ guards being located at vertices of $H_2 - v$. None of these guards can defend an attack at a vertex of P. It follows that at least $\theta(P)$ more guards are required to defend G. Therefore $\gamma_\infty(G) \geq \theta(H_2)+\theta(P) = \theta(G)$, so that $\gamma_\infty(G) = \theta(G)$.

On the other hand, suppose $\theta(H_2 - v) = \theta(H_2) - 1$. We consider two sub-cases depending on the parity H_1.

Suppose that H_1 is an even cycle We have $\gamma_\infty(H_2 - v) = \theta(H_2 - v) = \theta(H_2) - 1$, and can assume that there has been a sequence of attacks so that $\theta(H_2 - v)$ guards are located in $H_2 - v$. None of these guards can defend an attack at a vertex of P. It follows that at least $\theta(P)$ more guards are required in order to defend G. Therefore, $\gamma_\infty(G) \geq \theta(H_2 - v) + \theta(P) = \theta(H_2) - 1 + \theta(P) = \theta(H_2) - 1 + \theta(H_1) = \theta(G)$, so that $\gamma_\infty(G) = \theta(G)$.

Finally, suppose that H_1 is the odd cycle with vertex sequence $v = v_1$, $v_2, \ldots, v_{2m+1}, v_1$. Suppose there has been a sequence of attacks that results in $\gamma_\infty(H_2) = \theta(H_2)$ guards being located on vertices of H_2 and no guard being located at v. Then, since no guard on a vertex of $H_2 - v$ can defend an attack at a vertex of H_1, at least $\gamma_\infty(H_1 - v) = \theta(H_1 - v) = \theta(H_1) - 1$ more guards are required to defend G. In this case, $\gamma_\infty(G) \geq \theta(H_2) + \theta(H_1) - 1 = \theta(G)$, and so $\gamma_\infty(G) = \theta(G)$ Hence, consider a sequence of attacks that results in $\theta(H_2)$ guards being located at vertices of H_2, with a guard being located at v. Consider the sequence of attacks $v_3, v_5, \ldots, v_{2m-3}$. None of these can be defended by the guard at v, and two of these can be defended by the same guard. Hence there are $\theta(H_1) - 1$ guards in H_1, located at $v_1, v_3, v_5, \ldots, v_{2m-3}$. Suppose that these are all of the guards in H_1, so that the total number of guards is $\theta(H_2) - 1 + \theta(H_1) - 1 = \theta(G) - 1$. We shall obtain a contradiction. Consider an attack at v_2. It must be defended by the guard at v_3, otherwise v_{2m-1} is not dominated. Now, the sequence of attacks $v_4, v_6, \ldots, v_{2m-4}$ must be defended by the guards at $v_5, v_7, \ldots, v_{2m-3}$, respectively. However, when the guard at v_{2m-3} moves to v_{2m-4}, the vertex v_{2m-2} is no longer dominated. Therefore $\theta(G) - 1$ guards are insufficient to defend G, and it follows that $\gamma_\infty(G) = \theta(G)$. □

We note that all known graphs G with $\theta(G) = \alpha(G) + 1$ have $\gamma_\infty G = \theta(G)$.

4 All Guards Move

In this section we allow that each guard can move when an attack occurs. First, the relevant definitions are stated.

Let $D = \{x_1, x_2, \ldots, x_k\}$ and $D' = \{x'_1, x'_2, \ldots, x'_k\}$ be dominating sets of a graph G such that for $i = 1, 2, \ldots, k$ either $x_i = x'_i$ or x_i is adjacent to x'_i. Then D' is said to be obtained from D by a *guards move.* When $x_i \neq x'_i$, we say that the guard at x_i has *moved* to x'_i.

A dominating set D of a graph G is an *eternal m-secure set* if, for all possible sequences of attacks $R = r_1, r_2, \dots$ there exists a sequence $D = D_0, D_1, \dots$ of dominating sets such that $r_i \in D_i$ and D_i is obtained from D_{i-1} by a guards move. As before, the set D_i is the set of locations of the guards after the attack at r_i is defended. The size of a smallest eternal m-secure set in G is the *eternal m-security number*), and is denoted by $\gamma_m(G)$ or simply γ_m.

It is an open question whether, in this model, there is any advantage in allowing two guards to occupy the same vertex (there is no advantage in the "one guard moves" model [2]). Note, however, that our definition of guard move has implicitly disallowed the possibility of two guards being on the same vertex at the same time.

4.1 Vertex-Transitive Graphs

Goddard et al. [5] determine $\gamma_m(G)$ exactly for complete graphs, paths, cycles, and complete bipartite graphs. Further, they show that the lower bound $\gamma_m(G) = \gamma(G)$ holds for all Cayley graphs and conjecture this holds for all vertex-transitive graphs. We now disprove this conjecture.

Let P be the Petersen graph. Suppose it consists of the "outer" 5-cycle $1, 2, 3, 4, 5, 1$, "inner" 5-cycle $1', 3', 5', 2', 4', 1'$ and the matching $11', 22', \dots, 55'$. The graph P has domination number three. Up to symmetry, there is only one dominating set of size three: $\{1, 3', 4'\}$. Suppose guards are located these vertices. Consider an attack at vertex $1'$. If the guard at vertex 1 moves to $1'$, there is no way for the other guards to move and maintain a dominating set. By symmetry, assume the guard at $4'$ moves to $1'$. Some guard must move so as to dominate vertex 4. If the guard at 1 moves to 5 then there is no move for the guard at $3'$ that will maintain a dominating set. Since there is no dominating set of size 3 that contains 1 and $1'$, it follows that the guard at 1 moves to 2. But, again, there is no move for the guard at $3'$ that will maintain a dominating set.

4.2 A Linear Algorithm for the Eternal Security Number of a Tree

Let T be a tree. If T is isomorphic to K_2, then $\gamma_m = 1$, and if T is isomorphic to $K_{1,m}, m \geq 2$, then $\gamma_m = 2$.

The algorithm consists of repeatedly applying the two reductions below. We first describe the reductions, and then establish correctness and time complexity of the algorithm. We note that the proofs of the reductions give a strategy under which the tree can be defended by γ_m guards.

R1: Let x be a vertex of T incident to $\ell \geq 2$ leaves and to exactly one vertex of degree at least two. Delete all leaves incident to x.

R2: Let x be a vertex of degree two in T such that x is adjacent to exactly one leaf, y. Delete both x and y.

Lemma 20 *It T' is the result of applying reduction R1 to the tree T, then T' is a tree and $\gamma_m(T) = 1 + \gamma_m(T')$.*

Proof. Let $\ell_1, \ell_2, \ldots, \ell_k$, $k \geq 2$, be the leaves adjacent to x, and let w be the unique neighbor of x with degree at least two. It is clear that T' is a tree.

In order to defend against attacks at $\ell_1, \ell_2, \ldots, \ell_k$, there must always be at least two guards in $N[x]$, and two guards suffice. Thus, in a minimum eternal m-secure set, there are two guards that defend these leaves. In our algorithm, one of these two guards must always be on x, and the other may be on w or one of $\ell_1, \ell_2, \ldots, \ell_k$; these two guards will always move so as to be on adjacent vertices. The same strategy eternally defends T', after removing one of the two guards. Therefore, $\gamma_m(T') \leq \gamma_m(T) - 1$, or $\gamma_m(T) \geq \gamma_m(T') + 1$.

On the other hand, in any strategy that eternally defends T' there must be a guard at x or w in order to defend against attacks at x. Adding an initial guard at ℓ_1 and adapting the strategy for T' so that there is always a guard at x and so the two guards in $N[x]$ are always at adjacent vertices eternally defends T. Therefore, $\gamma_m(T) \leq 1 + \gamma_m(T')$. This completes the proof. □

Lemma 21 *It T' is the result of applying reduction R2 to the tree T, then T' is a tree and $\gamma_m(T) = 1 + \gamma_m(T')$.*

Proof. It is clear that T' is a tree. In order to defend against attacks at x there must always be a guard on x or y. This guard defends the vertices in $N[x] = \{x, y\}$, and the remaining $\gamma_m(T) - 1$ guards defend the rest of the vertices of T. Therefore, $\gamma_m(T') \leq \gamma_m(T) - 1$, or $\gamma_m(T) \geq \gamma_m(T') + 1$.

Conversely, any strategy that eternally defends T' can be extended to a strategy that eternally defends T by adding a guard to defend $N[x] = \{x, y\}$. Therefore, $\gamma_m(T) \leq 1 + \gamma_m(T')$. This completes the proof. □

Let T be a tree that is isomorphic to neither K_2 nor $K_{1,m}, m \geq 2$. Let x be a vertex of maximum eccentricity among those which are adjacent to one of more leaves. If $deg(x) \geq 3$ then R1 can be applied at x, otherwise R2 can be applied at x. Thus, if T is isomorphic to neither K_2 nor $K_{1,m}, m \geq 2$, one of the reductions R1 and R2 can be applied. A list of vertices all of whose neighbors except one are leaves can also be maintained and updated as the reductions are applied. The degrees of the vertices of T can be computed in linear time, and these can be updated in constant time as the reductions are applied. The total number of operations is at most $|V(T)|$. An application of reduction R1 can be thought of as $deg(x) - 1$ single vertex operations that each take constant time, and similarly for reduction R2. Thus, the algorithm can be implemented to run in linear time. □

We describe an infinite collection of graphs for which certain attacks result in every guard having to move in order to maintain an eternal m-secure set. Let P_n be a path with vertex sequence $x_1, x_2, \ldots, x_n$. Let G_n be the graph obtained from P_n and the disjoint sets of new vertices $\{u_1, u_2, \ldots, u_n\}$ and $\{v_1, v_2, \ldots, v_n\}$ by joining u_i and v_i to x_i for $i = 1, 2, \ldots, n$. Then $\gamma_m(G_n) = n + 1$ and any minimum eternal m-secure set in G_n must have a guard at each vertex of P. For any feasible initial configuration of the guards, each attack in the sequence $u_1, u_n, u_1, u_n, \ldots$, except possibly the first, causes every guard to move.

4.3 A Vertex Partition Theorem for Eternal Security in Trees

Define a *neo-colonization* as a partition $\{V_1, V_2, \ldots, V_t\}$ of graph G such that each V_i induces a connected graph. A part V_i is assigned a weight of one if it induces a clique and $1+\gamma_c(G[V_i])$, otherwise, where $\gamma_c(G[V_i])$ is the size of the smallest connected dominating set in the subgraph induced by V_i. [A connected dominating set of G is a set D such that D is a dominating set and $G[D]$ is connected]. Then $\theta_c(G)$ is the minimum weight of any neo-colonization of G. Goddard et al. proved that $\gamma_m(G) \leq \theta_c(G) \leq \gamma_c(G) + 1$. We can now prove the following theorem, which was conjectured in [5].

Theorem 22 *Let T be a tree. Then $\theta_c(T) = \gamma_m(T)$.*

Proof. The proof is by induction on the number of vertices in T. It is obviously true when T has one or two vertices. Consider a tree T with n vertices.

First suppose that T is such that reduction R2 can be applied, yielding T'. By the inductive hypothesis, $\theta_c(T') = \gamma_m(T')$ and it follows from Lemma 2 that $\gamma_m(T') + 1 = \gamma_m(T)$. Let P be a minimum weight neo-colonization partition of T'. In a minimum weight neo-colonization partition of T, either (1) x and y are a clique, in which case it is obvious that $\theta_c(T) = \gamma_m(T) = \theta_c(T') + 1$ or (2) y is a clique and x is either a clique or added to some part of P (i.e., the part containing x in the neo-colonization of T is a proper superset of a part in the neo-colonization of T'), in which case $\theta_c(T) \geq \theta_c(T') + 1$, or (3) x and y are added to some part of P, which implies in the weighting of this part in T, it must be weighted as a connected dominating set, so we have $\theta_c(T) \geq \theta_c(T') + 1$.

On the other hand, suppose T is such that reduction R1 can be applied, yielding T'. By the inductive hypothesis, $\theta_c(T') = \gamma_m(T')$ and it follows from Lemma 21 that $\gamma_m(T') + 1 = \gamma_m(T)$. Let P be a minimum weight neo-colonization partitioning of T'. In an minimum weight neo-colonization partitioning of T, either (1) x is a clique, in which case the leaves adjacent to x must each be cliques in this partition of T, and so obviously $\theta_c(T) \geq \theta_c(T') + 2$ or (2) x and the leaves adjacent to it form a part by themselves and thus $\theta_c(T) \geq \theta_c(T') + 2$ or (3) x, and because of the minimality of the partitioning, its adjacent leaves, are added to some part of P. If this part was a clique in P, then the weight of the part is two greater in the neo-colonization of T, since it must be weighted as a connected dominating set in the neo-colonization of T. If this part was not a clique in P, then the weight if the part is one greater in the neo-colonization of T than in that of T'. □

Note added in press: Goldwasser and Klostermeyer have recently proved [6] that for all $t \geq 3$, certain complements of Kneser graphs satisfy $\alpha = t$ and $\gamma_\infty = \binom{t+1}{2}$. That is, the bound in Theorem 1 is tight. It is therefore natural to ask how far the realizability results above can be extended. Given enough base cases, it is possible to use the "star product" method of this paper to show realizability of any triple (a, g, t) with $g \leq ca$ for a fixed constant c. On the other hand, at this time it is unknown whether the triples $(3, 5, t)$, $t \geq 5$ are realizable. The inductive proof of Proposition 11 appears not to carry through with a quadratic bound on g, since $\binom{a+1}{2}$ is greater than any sum $\binom{x}{2}+\binom{y}{2}$ where $x+y = a$: the cases $\binom{a}{2}+1 < g < \binom{a+1}{2}$ seem to need to be handled using a different method.

References

[1] A. Burger, E. Cockayne, W. Gründlingh, C. Mynhardt, J. van Vuuren, and W. Winterbach (2004), Finite Order Domination in Graphs (2004), *J. Comb. Math. Comb. Comput.*, vol. 49, pp. 159-176

[2] A. Burger, E. Cockayne, W. Gründlingh, C. Mynhardt, J. van Vuuren, and W. Winterbach (2004), Infinite Order Domination in Graphs (2004), *J. Comb. Math. Comb. Comput.*, vol. 50, pp. 179-194

[3] E. Cockayne, P. Dreyer, S. M. Hedetniemi, and S. T. Hedetniemi (2004), Roman domination in graphs *Discrete Math.*, vol. 278, pp. 11-22

[4] E. Cockayne, O. Favaron, and C. Mynhardt (2003), Secure domination, weak roman domination and forbidden subgraphs, *Bull. Inst. Combin. Appl.*, vol. 39, pp. 87-100

[5] W. Goddard, S. M. Hedetniemi, and S. T. Hedetniemi (2005), Eternal Security in Graphs, *J. Comb. Math. Comb. Comput.*, vol. 52, pp. 160-180

[6] J. Goldwasser and W. Klostermeyer (2007), Tight Bounds on Eternal Dominating Sets, manuscript

[7] W. Klostermeyer and G. MacGillivray (2005), Eternally Secure Sets, Independence Sets, and Cliques, *AKCE International Journal of Graphs and Combinatorics*, vol. 2 (2005), pp.119-122

[8] W. Klostermeyer and G. MacGillivray (2006), Eternal Security in Graphs of Fixed Independence Number, to appear in *J. Comb. Math. Comb. Comput.*