Dynamic Dominating Sets: the Eviction Model for Eternal Domination

William F. Klostermeyer
School of Computing
University of North Florida
Jacksonville, FL 32224-2669
wkloster@unf.edu

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

Gary MacGillivray
Dept. of Mathematics and Statistics
University of Victoria
Victoria, Canada
gmacgill@uvic.ca

Abstract

We consider a discrete-time dynamic problem in graphs in which the goal is to maintain a dominating set over an infinite sequence of time steps. At each time step, a specified vertex in the current dominating set must be replaced by a neighbor. In one version of the problem, the only change to the current dominating set is replacement of the specified vertex. In another version of the problem, other vertices in the dominating set can also be replaced by neighbors. A variety of results are presented relating these new parameters to the eternal domination number, domination number, and independence number of a graph.

1   Introduction

Consider the problem of maintaining a dominating set in a network where single nodes may fail or be shut down for a time unit. If the node that fails is not in the current dominating set, there is nothing to do. If it is, then it should be replaced by a neighboring vertex not in the current dominating set so that the modified collection of vertices is a dominating set that can also handle such node failures. This can be regarded as a variation of the eternal domination problem, which has received considerable recent attention [1, 3, 4, 5, 7, 8, 9, 10, 11]. A survey on graph protection using mobile guards, including eternal domination, can be found in [12]. In the remainder of this section, we review some basic terminology and definitions from the study of domination in graphs, as well as from eternal domination, before moving on to the variation that we shall study. In subsequent sections, we present results on complexity and bounds. The paper concludes with some suggestions for future research.

Denote the open and closed neighborhoods of a vertex $x \in V$ by $N(x)$ and $N[x]$, respectively. That is, $N(x) = \{v|xv \in E\}$ and $N[x] = N(x) \cup \{x\}$. Further, for $S \subseteq V$, let $N(S) = \bigcup_{x \in S} N(x)$. For any $X \subseteq V$ and $x \in X$, we say that $v \in V - X$ is an *external private neighbor* of x *with respect to* X if v is adjacent to x but to no other vertex in X. The set of all such vertices v is the *external private neighborhood of* x *with respect to* X.

A *dominating set* of graph G is a set $D \subseteq V$ with the property that for each $u \in V - D$, there exists $x \in D$ adjacent to u. A dominating set D is a *connected dominating set* if the subgraph $G[D]$ induced by D is connected. The minimum cardinality amongst all dominating sets of G is the *domination number* $\gamma(G)$, while the minimum cardinality amongst all connected dominating sets is the *connected domination number* $\gamma_c(G)$. A thorough background on domination can be found in [6].

For a graph G, we denote the independence number by $\beta(G)$, and the clique covering number by $\theta(G)$. Of course, $\theta(G) = \chi(\overline{G})$. We shall say the *size* of a clique is the number of vertices it contains.

1.1   Eternal Domination

Eternal domination problems can be modeled as two-player games between a *defender* and an *attacker*. To start the game, the defender chooses a set D_1 of vertices to be occupied by guards. The two sides then move alternately, with the attacker choosing a vertex r_i, $i \geq 1$ to attack, followed

by the defender choosing the next set D_{i+1} of vertices to be occupied by guards. Note that the location of an attack can be chosen by the attacker depending on the location of the guards. Each attack is handled by the defender by choosing the next guard locations subject to the rules of the particular game. The defender wins the game if they can successfully defend any series of attacks; the attacker wins otherwise.

In the *eternal dominating set problem*, each $D_i, i \geq 1$, is required to be a dominating set, each $r_i \in V$ (assume without loss of generality $r_i \notin D_i$), and D_{i+1} is obtained from D_i by moving one guard to r_i from a vertex $v \in D_i, v \in N(r_i)$. The size of a smallest eternal dominating set for G is denoted $\gamma^\infty(G)$. This problem was first studied in [3].

In the *m-eternal dominating set problem*, each $D_i, i \geq 1$, is required to be a dominating set, each $r_i \in V$ (assume without loss of generality $r_i \notin D_i$), and D_{i+1} is obtained from D_i by moving guards to neighboring vertices. Subject to the condition that $r_i \in D_{i+1}$ – some guard moves to r_i – each guard in D_i can either move to an adjacent vertex or stay at the vertex where he is located. The size of a smallest m-eternal dominating set for G is denoted $\gamma_m^\infty(G)$. This "all-guards move" version of the problem was introduced in [4].

We say that a vertex is *protected* if there is a guard on the vertex or on an adjacent vertex. We say that a vertex is "empty" or "unoccupied" if it contains no guard and "occupied" if it contains a guard.

1.2   Eviction Model

In a distributed computer system or a computer network, it may be desirable to maintain copies of a file throughout the network so that each computer can have fast access to a close copy of the file. However, having too many copies of a file in the network means it can be costly to propagate updates to the file, maintain consistency among copies, and keep track of the files' locations. Hence it is desirable to find a balance between these two competing interests.

In computer networks, it is also the case that servers storing file copies sometimes undergo maintenance or upgrades. At these times, it may be necessary to migrate the files on that machine to another in the network to maintain good access for all users of the network. This problem motivates the following variation of the eternal domination problem, which we term the *eviction model*. For consistency in terminology, we shall use the term "guard" synonymously with "file", i.e., we shall talk about moving guards

around the network, instead of files.

In the eternal domination set eviction problem, each $D_i, i \geq 1$, is required to be a dominating set. An attack is a vertex $r_i \in D_i$. If there exists at least one $v \in N(r_i)$ with $v \notin D_i$, then the next guard configuration D_{i+1} is obtained from D_i by moving the guard from r_i to a vertex $v \in N(r_i), v \notin D_i$ (i.e., this is the one-guard moves model). If no such v exists, then $D_{i+1} = D_i$. (The rationale for taking $D_{i+1} = D_i$ when every neighbor of r_i holds a guard is that the guards on vertices in $D - r_i$ form a dominating set.)

By an *eternal dominating set in the eviction model* we mean any dominating set starting from which it possible to defend all possible sequences of attacks. For any graph G, the vertex set $V(G)$ has this property, hence there is a least integer, $e^\infty(G)$, for which there is an eternal dominating set in the eviction model.

The eternal dominating set eviction problem can be regarded as maintaining a dynamic dominating set in a network model where single nodes can fail for a unit length interval of time, and where vertices in the dynamic dominating set that fail must be replaced by neighboring vertices, if there is a neighbor not in the dominating set.

Proposition 1 *In the one-guard moves eviction model, if $k < |V(G)|$ guards can defend an arbitrarily long sequence of attacks in G, then so can $k + 1$ guards.*

Proof: Initially place k guards on the k vertices of an eternal dominating set in the eviction model, and place the $(k + 1)$-st guard anywhere. Suppose vertex v is attacked. If the vertices holding the other k guards form an eternal dominating set in the eviction model, then move the guard at v arbitrarily. Otherwise, either the guard at v can be moved to a vertex in an eternal dominating set of size k in the eviction model, or the $(k + 1)$-st guard is on the vertex to which v should move. In the former case move the guard at v to maintain the eternal dominating set in the eviction model, and in the latter case move the guard at v arbitrarily (if the guard can move at all) as the eternal dominating set in the eviction model is guaranteed to be maintained. □

We shall also consider an "all-guards move" model of the eviction problem. This model requires that an attacked vertex be unoccupied after the attack is handled, until the next attack occurs. For example, suppose there is an edge uv with guards at u and v. If u is attacked, we do not allow the

guard at v to move to u in response to the attack at u. In other words, if u has a guard and is attacked, then the guard at u must leave and no guard can return to u until the next attack occurs (or later). An eternal dominating set in this model is defined as in the "one guard moves model", except that any guard on an unattacked vertex in the dominating set is allowed to move to a neighbor when there is an attack (so long as no guard moves to the attacked vertex). The size of a smallest eternal dominating set in the all-guards move eviction model for G is denoted $e_m^\infty(G)$.

Some simple examples may help to develop a feel for these problems. It is clear that $e^\infty(K_n) = e_m^\infty(K_n) = 1$. We have $e^\infty(K_{1,m}) = e_m^\infty(K_{1,m}) = m$. On the other hand, $e^\infty(C_6) = 3$, and $e_m^\infty(C_6) = 2$. Similarly, $e^\infty(P_5) = 3$ and $e_m^\infty(P_5) = 2$.

1.3   Previous Results

A *neo-colonization* is a partition $\{V_1, V_2, \ldots, V_t\}$ of the vertex set of a graph G such that each $G[V_i]$ is a connected graph. A part V_i is assigned a weight $\omega(V_i) = 1$ if it induces a clique and $\omega(V_i) = 1 + \gamma_c(G[V_i])$ otherwise. The parameter $\theta_c(G)$ is used to denote the minimum total weight of any neo-colonization of G and is called the *clique-connected cover number* of G. Goddard et al. [4] defined this parameter and proved that $\gamma_m^\infty(G) \le \theta_c(G)$.

Theorem 2 [8] *For any tree T, $\theta_c(T) = \gamma_m^\infty(T)$.*

It is easy to see that $\gamma_m^\infty(G) \le \gamma_c(G) + 1$.

The next theorem describes the fundamental inequality chain for the traditional eternal domination problems. Each of the inequalities have been shown to be sharp for certain graphs.

Theorem 3 [3, 4] *Let G be a graph. Then $\gamma(G) \le \gamma_m^\infty(G) \le \beta(G) \le \gamma^\infty(G) \le \theta(G)$.*

2   One Guard Moves

Let G and H be graphs. The *wreath product* (or lexicographic product) of G and H is the graph $G\,wr\,H$ with vertex set $V(G\,wr\,H) = V(G) \times V(H)$,

and edge set

$$E(G \, wr \, H) = \{(g,h)(g',h') : gg' \in E(G), \text{ or } g = g' \text{ and } hh' \in E(H)\}.$$

Informally, $G \, wr \, H$ is the graph obtained by replacing each vertex of G with a copy of H and adding all possible edges between copies of H that replaced adjacent vertices of G. Note that any subgraph induced by a set of $|V(G)|$ vertices, one from each copy of H, is isomorphic to G.

Proposition 4 *Let G be a graph. For any integer $n \geq 2$ we have*

$$\gamma(G) = \gamma(G \, wr \, K_n) = e^\infty(G \, wr \, K_n).$$

Proof: Let D be a dominating set of G. For any vertex $x \in K_n$, the set $D' = \{(d,x) : d \in D\}$ is a dominating set of $G \, wr \, K_n$. If one guard is placed on each copy of K_n corresponding to a vertex in D, the guard can always relocate, when necessary, to remain within that clique. The set of guards therefore always forms a dominating set. $\square$

Because the configuration of guards must always form a dominating set, $e^\infty(G) \geq \gamma(G)$ for any graph G. Proposition 4 provides infinitely many examples where equality occurs.

We consider the following decision problem:

EVICTION
INSTANCE: A graph G and an integer $k > 0$.
QUESTION: Is $e^\infty(G) \leq k$?

It is not clear that EVICTION belongs to NP, as it is not clear how to confirm in polynomial time that a given initial configuration of guards can defend all possible sequences of attacks. We describe a straightforward algorithm that decides whether a given configuration of k guards "works" (or whether there exists a configuration of k guards that "works") in time exponential in k. Given G, construct an arc-colored digraph D as follows. The vertex set of D is the set of k-vertex dominating sets of G. The set of colors is $V(G)$. There is an arc from X to Y of color v when $X - Y = \{v\}$, and v is adjacent in G to the unique vertex $w \in Y - X$. Now, delete any vertex X which is not the origin of an arc colored x for some $x \in X$. Repeat this step until no further vertices can be deleted. Call the resulting

digraph D'. The vertices of D' are the guard configurations from which any sequence of attacks can be defended.

Corollary 5 *EVICTION is NP-hard, even when the input is restricted to interval graphs.*

Proof: The transformation is from the NP-complete problem of deciding whether a given interval graph has a dominating set of size at most the given integer k [2]. Suppose an interval graph G and an integer k are given. The transformed instance of EVICTION is $G\, wr\, K_2$ and the same integer k. The result now follows from Proposition 4 on noting that $G\, wr\, K_2$ is an interval graph (make a second copy of each interval). $\square$

Corollary 6 *EVICTION is NP-hard, even when the input is restricted to chordal graphs.*

Corollary 5 motivates searching for bounds on e^∞. The following general upper bound holds for all graphs.

Theorem 7 *Let G be a graph. Then $e^\infty(G) \leq \theta(G)$.*

Proof: We will maintain the invariant that there is a partition into $\theta(G)$ cliques with exactly one guard on each of them. Let C be a minimum clique covering of G. For each clique $c \in C$, place a guard on some vertex of c. Hence the invariant holds initially.

Suppose the guard located in c is located on a vertex that is attacked. Assume that the guard can move, otherwise the guard trivially remains in its clique and the invariant holds. If c has at least two vertices, relocate the guard to another vertex of c, and the invariant holds. Otherwise, c has exactly one vertex. Relocate the guard to any available adjacent vertex, say x, belonging to a clique $c' \in C$. Since C is a minimum clique covering, c' has at least two vertices. The invariant now holds for the revised minimum clique covering $C' = (C - \{c, c'\}) \cup \{c + x, c' - x\}$. $\square$

There are many examples where the bound in Theorem 7 is not sharp (see Theorem 11). The example on the smallest number of vertices is $K_4 - e$.

3 Independence Number

In this section, we explore the relationship between $e^\infty(G)$ and the independence number of G. A general lower bound that holds for triangle-free graphs is shown first. As stated above, it is known that $\gamma^\infty(G) \geq \beta(G)$, see [4].

Theorem 8 *Let G be a triangle-free graph. Then $e^\infty(G) \geq \beta(G)$.*

Proof: Let I be a maximum independent set of G. Let D be an eternal dominating set in the eviction model D. Place guards on the vertices of D. We show that it is possible to force a guard to be located on every vertex of I, from which the result follows.

Suppose there is no guard on $v \in I$. Since the configuration of guards forms a dominating set, there is a guard on a neighbor u of v. Note that $u \notin I$. Attack u. Since G is triangle-free, the guard on u must either move to v, or to a vertex non-adjacent to v. In the latter case, there must be a guard on a neighbor of v and fewer neighbors of v must hold a guard. Repeating the procedure, since v must be dominated, eventually a guard must be located on v. $\square$

Proposition 9 $e^\infty(C_3) = 1$, $e^\infty(C_5) = 2$, and for $k \geq 3$, $e^\infty(C_{2k+1}) = k + 1$.

Proof: It is easy to see that $e^\infty(C_3) = 1$, and $e^\infty(C_5) = 2$. Let $k \geq 3$ and consider C_{2k+1}. By Theorems 7 and 8, $k \leq e^\infty \leq k + 1$. We show that k guards do not suffice.

Suppose C_{2k+1} is the cycle $v_1, v_2, \ldots, v_{2k+1}, v_1$. Following the proof of Theorem 8, it can be assumed that the k guards are located on the vertices in $\{v_1, v_3, \ldots, v_{2k-1}\}$. First attack the guard at v_3. If he relocates to v_4, then attacking the guard at v_1 necessarily results in a configuration that is not a dominating set; hence, to maintain a dominating set, he must relocate to v_2. Now attack the guard at v_5. To maintain a dominating set, he must relocate to v_4. For the same reason, attacking the guard at v_7 forces him to move to v_6. Continuing in this way, attacking the guard at v_{2k-1} necessarily forces a configuration that is not a dominating set. It follows that $e^\infty = k + 1$. $\square$

It follows that odd cycles with at least seven vertices are examples of graphs where equality holds in Theorem 7 and inequality holds in Theorem 8.

We now use Proposition 9 to construct connected triangle-free graphs with $\theta \geq e^\infty \geq \beta$ such that the differences $\theta - e^\infty$, $e^\infty - \beta$ are any two non-negative integer a and b, respectively. Suppose a and b are given. Begin with a path P on $2b + 2a$ vertices, $2b$ copies of C_7 and $2a$ copies of C_5. For $i = 1, 2, \ldots, 2b$, add an edge joining the i-th vertex of P to a vertex in the i-th copy of C_7. For $j = 1, 2, \ldots, 2a$, identify a vertex of the j-th copy of C_5 with vertex $2b + j$ of P. Call the resulting graph $H_{a,b}$. The clique covering number $\theta(H_{a,b}) = 4 \cdot 2b + 2 \cdot 2a + a$, as there are a cliques of size two that cover one vertex from each C_5. The independence number $\beta(H_{a,b}) = 3 \cdot 2b + b + 2 \cdot 2a$. We claim that $e^\infty(H_{a,b}) = 4 \cdot 2b + 2 \cdot 2a$. There must be four guards on the subgraph induced by a C_7 and its unique neighbor on P, otherwise attacking the guards on the C_7 as in the proof of Proposition 9 leads to a configuration that is not a dominating set. Furthermore, since any such subgraph can be covered by four cliques of size two, the guards can never be evicted out of the subgraph. Hence $4 \cdot 2b$ guards are necessary and sufficient to protect the subgraph induced by the first $2b$ vertices of P and the C_7's and can do so without a guard ever being forced to leave it. In a similar way, two guards are necessary and sufficient to protect each C_5 and can do so without ever being forced to leave it. Thus $e^\infty = 4 \cdot 2b + 2 \cdot 2a$. We now have $\theta - e^\infty = a$ and $e^\infty - \beta = b$, as desired.

We now show that equality holds in Theorem 8 for all bipartite graphs.

Lemma 10 *Let G be a bipartite graph. Then $e^\infty(G) \leq \beta(G)$.*

Proof: We will maintain the invariant that there is a maximum matching with a guard on some end of each of its edges, and a guard on each vertex not incident with an edge in the matching. Such a configuration of guards is easily seen to form a dominating set. By König's Theorem, any such configuration has β guards.

Let M be a maximum matching. A minimum vertex cover C (of size $|M|$) can be chosen so that each edge in M has exactly one end in C. Let W be the set of vertices of G which are not an end of an edge in M. Since W is a subset of $V - C$, and the complement of a minimum vertex cover is an independent set, the set W is independent. The neighborhood of any vertex in W is a subset of C.

Place a guard on each vertex of $C \cup W$, so that the invariant holds initially. If a guard in C is evicted, he relocates to the other end of the edge in M, and the invariant still holds. Suppose that a guard on $w \in W$ is evicted and relocates to $c \in C$. For this to be possible, the guard initially on c must have been evicted and moved to c' such that $cc' \in M$. Replace M by the maximum matching $M - cc' + cw$, and the invariant holds. Finally,

255

if a guard on an end of an edge in M which is not in C is evicted, then he relocates to the other end of the edge in the matching and the invariant holds again. □

The following is obtained by combining Theorem 8 and Lemma 10.

Theorem 11 *Let G be a bipartite graph. Then $e^{\infty}(G) = \beta(G)$.*

We now show that equality holds in Theorem 8 for many non-bipartite graphs. Let $G(n, k)$ be the graph with vertex set equal to the set of all k-subsets of an n-set in which two vertices are adjacent if and only if their intersection is nonempty So $G(n, k)$ is the complement of a Kneser graph and is sometimes called a Johnson graph. Johnson graphs were used in [5] to show that the eternal domination number may be much larger than the independence number of a graph. It is known that $\beta(G(n,k)) = \lfloor \frac{n}{k} \rfloor$, see for example [5]. We refer to the elements of the n-set as *symbols* and use the integers $1, 2, \ldots, n$ to represent these symbols.

Theorem 12 *For all $n \geq k \geq 1$, $e^{\infty}(G(n, k)) = \beta(G(n, k))$.*

Proof: By definition of adjacency in $G(n, k)$, a vertex X is dominated if and only if there is a guard on a vertex Y such that $X \cap Y \neq \emptyset$. We will say that a guard *covers* a symbol u if it is located on vertex Y such that $u \in Y$.

Suppose that there are fewer than $\beta = \lfloor \frac{n}{k} \rfloor$ guards. Then the maximum number of symbols covered is $k \cdot \left(\lfloor \frac{n}{k} \rfloor - 1 \right) \leq n - k$. Hence there is a set X of k symbols that are not covered. This vertex is not dominated by the configuration of guards.

Now suppose there are β guards. We will maintain the invariant that there are at most $k - 1$ symbols that are not covered. For any such configuration, the vertices that hold the guards form a dominating set. This holds for the initial placement of guards on the vertices $\{1, 2, \ldots, k\}$, $\{k+1, k+2, \ldots, 2k\}, \ldots, \{(\lfloor \frac{n}{k} \rfloor - 1)k + 1, (\lfloor \frac{n}{k} \rfloor - 1)k + 2, \ldots, \lfloor \frac{n}{k} \rfloor k\}$. Suppose it holds for the current configuration, and that the guard at $\{x_1, x_2, \ldots, x_k\}$ is attacked. This guard has a move to a vertex that covers all currently uncovered symbols, at least one of $x_1, x_2, \ldots, x_k$, and possibly some other symbols. The invariant holds after this move is made. □

A graph with $\beta = 1$ is a complete graph and has $e^{\infty} = 1$. The graphs with $\beta = 2$ can also be eternally dominated by a small number of guards. A *dominating vertex* in a graph G with n vertices is a vertex of degree $n - 1$.

Theorem 13 *Let G be a graph with $\beta(G) = 2$. If G has two dominating vertices, then $e^\infty(G) = 1$. Otherwise, $e^\infty(G) = 2$.*

Proof: If G has dominating vertices x and y, then a single guard can relocate back and forth between them and maintain a dominating set.

Finally, suppose G has at most one dominating vertex. Then G is the complement of a triangle-free graph with at most one isolated vertex. Initially locate the guards on any dominating set of size two, say $\{u, v\}$. Suppose the guard on u is attacked. If v has a non-neighbor $w \neq u$, then whether or not u and v are adjacent, the guard at u guard can relocate to w and the resulting configuration is a dominating set. If no such vertex w exists, the guard at u can relocate to any vertex z and the resulting configuration of guards is a dominating set. $\square$

A graph G is a *split graph* if $V(G)$ can be partitioned into I and C, where I is an independent set and the subgraph induced by C is a clique. There can be more than one such partition. When G is a split graph and I is a maximum independent set, every vertex in C has at least one neighbor in I (otherwise I is not maximum).

Theorem 14 *Let G be a connected split graph. Then $e^\infty(G) \leq \beta(G)$.*

Proof: Let I be a maximum independent set of G. Then each vertex in C has at least one neighbor in I. Initially, place one guard on each vertex of I. This is a dominating set. When a guard is evicted from a vertex in I to a neighboring vertex in C, the resulting configuration remains a dominating set because a guard in C dominates all vertices of C. If a guard in C is evicted, he relocates to his initial position in I, and the resulting configuration remains a dominating set. $\square$

3.1   The case $\beta = 3$

In the following, we show in several steps that any graph with independence number three has an eternal dominating set in the eviction model with at most five vertices. This is perhaps the main result of the paper and is motivated by results in [7] which, among other things, shows that any graph G with independence number three has $\gamma^\infty(G) \leq 6$; this bound is sharp as shown in [5]. The first few results eliminate some easy cases and provide a bit of structure with which to work.

Lemma 15 *Let G be a graph with $\beta = 3$. If G is disconnected or has a cut vertex, then $e^\infty(G) \leq 3$.*

Proof: If G is disconnected, then it consists of at most three connected components, at most one of which has $\beta = 2$ (the others having $\beta = 1$). The result then follows from Theorem 13 and the fact that $e^\infty(K_n) = 1$.

Suppose that G is connected and has a cut-vertex, x. Since $\beta = 3$, the graph $G - x$ has at most three components. If it has exactly three, then x must be adjacent to all the vertices in at least one of the components, else $\beta > 3$. Then each component is complete (and x union one of the components is complete) and in this case, G can be guarded by three guards. Suppose, then, that $G - x$ has two components, C_1 and C_2. Since $\beta = 3$, either C_1 or C_2 is complete, or both are complete.

Suppose C_1 and C_2 are both complete. Then each has at least two vertices, otherwise $\beta(G) < 3$. Let C_1 be guarded by one guard who is never forced to leave it. If the subgraph induced by $V(C_2) \cup \{x\}$ is complete (and thus can be guarded by one guard), then either $\beta(G) < 3$ (if x is adjacent to all vertices in C_1) or else the subgraph induced by $V(C_2) \cup \{x\}$ can be guarded by a guard who moves back and forth between two vertices in C_2. If the subgraph induced by $V(C_2) \cup \{x\}$ is not complete, then C_2 can be guarded by one guard that moves between vertices in C_2. Vertex x is guarded by a guard initially on x and which moves to a vertex in C_1 or C_2 when evicted (and then back to x when evicted again).

Suppose that C_1 is complete and C_2 is not complete. Then C_2 has at least three vertices. Let us first assume that x is not adjacent to all the vertices in C_1. If C_1 has at least two vertices, it can be guarded by a single guard who is never forced to leave it, while the subgraph induced by $V(C_2) \cup \{x\}$ can be guarded by two guards who are never forced to leave it. Suppose, then, that C_1 has only one vertex. If C_2 has at least four vertices, then it can be guarded by two guards who are never forced to leave it, while a third guard moves back and forth between the vertex in C_1 and x. The last case is when $C_2 \cong K_{1,2}$. Then G has only five vertices, and can be guarded by three guards.

On the other hand, if x is adjacent to all the vertices in C_1, then combine x and C_1 into a single clique and protect that clique with one guard that never leaves. The $\beta(C_2) = 2$ and can be guarded with two guards by Theorem 13. □

Lemma 16 *Let G be a 2-connected graph with $\beta = 3$. Then either G has disjoint independent sets I_1 and I_2 such that $|I_1 \cup I_2| = 5$, or G is a split*

graph.

Proof: If no such independent sets exist, then the vertices not in the independent set of size three must induce a clique. □

Corollary 17 *Let G be a 2-connected graph with $\beta = 3$. If G does not have disjoint independent sets I_1 and I_2 such that $|I_1 \cup I_2| = 5$, then $e^\infty(G) \leq 3$.*

Proof: By Lemma 16, G is a split graph with $\beta = 3$. The statement follows from Theorem 14. □

By the above results, it remains to consider 2-connected graphs which have disjoint independent sets of size 2 and 3, respectively. We define a *dominating configuration* to be a set of five vertices which is a dominating set and contains either an independent set of size three, or two disjoint independent sets of size two. Since a dominating configuration is a dominating set, guards placed on the vertices of such a set can defend any single attack. If this can always be done so that the resulting set of guards' locations is a dominating configuration, then we have an eternal dominating set in the eviction model. We next examine situations where this is not possible, and obtain structural information which will be useful in obtaining a strategy for defending G.

Lemma 18 *Let G be a 2-connected graph with $\beta = 3$. Suppose guards are located at the vertices of a dominating configuration. If any guard is attacked, then he can relocate so that the resulting set of guards' locations is a dominating set.*

Suppose guards are located at the vertices of a dominating configuration, D, and that the guard at v is attacked. Note that $D - \{v\}$ contains an independent set I of size two. If v has an external private neighbor w and the guard at v relocates to w, then $I \cup \{w\}$ is a maximum independent set and $(D - \{v\}) \cup \{w\}$ is dominating. Otherwise, v has no external private neighbors and no matter where the guard at v relocates (or if he can not move), the resulting set of guards' locations is a dominating set. □

Lemma 19 *Let G be a 2-connected graph with $\beta = 3$. Suppose guards are located at the vertices of a dominating configuration, $D = \{a, b, c, d, g\}$, but when the guard at g is attacked it is not possible to maintain a dominating configuration. Then*

259

(i) g has a neighbor $w \notin D$;

(ii) the subgraph induced by $D - \{g\}$ must have an independent set $\{a, b\}$ disjoint from an edge cd;

(iii) g has no external private neighbors;

(iv) every vertex in $N(g) - D$ is adjacent to both c and d;

(v) one of a and b is adjacent to both c and d; without loss of generality, $ac, ad \in E$.

(vi) the graph $H = G - (\{c, d\} \cup N[g])$ has independence number at most two.

Proof: Since it is not possible to maintain a dominating configuration, it must be that the guard at g can move when evicted. Hence g has a neighbor w not in the set of vertices occupied by the other guards. This proves (i).

The subgraph induced by $D - \{g\}$ must have an independent set of size two that is disjoint from another edge, otherwise we have a dominating configuration no matter to which vertex the guard at g relocates. Without loss of generality the vertices a and b are nonadjacent, and the vertices c and d are adjacent. Statement (ii) follows.

To see (iii), note that the vertices a, b and any private neighbor of g would form an independent set of size three. Since $\beta = 3$, it would be possible to maintain a dominating configuration, contrary to our hypothesis.

From (iii), we know that g has no external private neighbors. Furthermore, if g has a neighbor $w \notin D$ which is not adjacent to c, say, then w and c are an independent set of size two and a dominating configuration can be obtained by moving a guard from g to w. Hence (iv) is proved.

Since $\{a, b, c, d\}$ contains no independent set of size three, there must be at least two edges with one end in $\{a, b\}$ and the other end in $\{c, d\}$. If neither a nor b is adjacent to both c and d, then we can assume $ac, bd \in E$. But then $(D - \{g\}) \cup \{w\}$ is a dominating configuration with independent sets $\{a, d\}, \{b, c\}$. Statement (v) follows.

Finally, since, for any independent set X in $H = G - (\{c, d\} \cup N[g])$, the set $X \cup \{g\}$ is independent in G (as all neighbors of g are not in H, by the definition of H), we have $\beta(H) \leq 2$. This proves (vi). $\square$

We now use the structural information given above to describe a strategy for defending graphs G where a dominating configuration can not be maintained. Since there is perfect information, and G is known in advance, it can be assumed that this is known.

Lemma 20 *Let G and D be as in Lemma 19. If g is adjacent to both a and b, then $e^{\infty}(G) \leq 5$.*

Proof: Let $H = G - (\{c, d\} \cup N[g])$. By Lemma 19 (vi), $\beta(H) \leq 2$. By our hypothesis, $D \cap V(H) = \emptyset$.

If H is connected and has at least two vertices, then by Theorem 13, it can be defended by either one or two guards who are never forced to move to a vertex in $G - H$ (only one guard is needed when H has two vertices). These guards, together with a guard on g who will move back and forth between g and b, and a guard on c that will move back and forth between c and d, are an eternal dominating set in the eviction model.

Suppose H has only one vertex, z. Since G is 2-connected, z has at least two neighbors. If z is adjacent to both c and d, then guards placed on g and c as above suffice to defend G. If z is not adjacent to both c and d, then place a guard on z, a guard on c and a guard on g. If z has a neighbor $x \notin \{b, c, d\}$ then G is defended if the guard on z moves back and forth to x, the guard on c moves back and forth to d, and the guard on g moves back and forth to b. Otherwise, since z has degree at least two, $zb \in E(G)$. Moving the guard on z back and forth to b, the guard on c back and forth to d, and the guard on g back and forth to w (from Lemma 19 (i)) successfully defends G.

Finally, suppose H is disconnected. Then it consists of two disjoint cliques, one or both of which might have only one vertex. If both have at least two vertices, then placing one guard on each of them, plus a guard that moves between c and d, and a guard that moves between a and g, defends G. The case where only one of the cliques is a singleton can be handled by placing a guard on the other clique, and then proceeding as above when H was a singleton.

Hence assume that the two disjoint cliques of H each consist of single vertices, say z_1 and z_2. Since G is 2-connected, each vertex has degree at least two.

If at least one of z_1 and z_2 are adjacent to c and d, then we are done as above. Thus it remains to consider the situation where neither of these vertices is adjacent to both c and d. By definition of H, neither of them is

261

adjacent to g. Observe that if w is the only neighbor of g not in $\{a, b, c, d\}$, then G has eight vertices and one can verify that five guards suffice. Otherwise, let $w' \neq w$ be another neighbor of g. The vertex w' has the same adjacency properties of w in terms of being adjacent to g, c, and d.

If z_1, say, has a neighbor $w_1 \notin \{b, c, d\}$, then G can be defended similarly to when H was a singleton. Hence assume both z_1 and z_2 are adjacent to b, and one of c and d.

Suppose first that both z_1 and z_2 are only adjacent to b and c. Observe that ww' must be an edge; otherwise $\{w, w', z_1, z_2\}$ is an independent set of size four. For the same reason, $N(g) - \{a, b, c, d\}$ must induce a clique. Our strategy is as follows. One guard moves back and forth along gw, one guard moves back and forth along ad, one guard moves back and forth along bz_2, and one guard moves back and forth along cz_1.

Suppose both z_1 and z_2 are adjacent to b, z_1 is adjacent to c (and not d), and z_2 is adjacent to d (and not c). The same strategy as immediately above defends G. $\square$

Lemma 21 *Let G and D be as in Lemma 19. If g is adjacent to exactly one of a and b, then $e^\infty(G) \leq 5$.*

Proof: As before, let $H = G - (\{c, d\} \cup N[g])$. By Lemma 19 (vi), $\beta(H) \leq 2$.

We first consider the situation where g is adjacent to b and not a. In this situation, $a \in V(H)$ and $b \notin V(H)$. The argument is almost exactly the same as in the proof of Lemma 20. In fact, it is a bit simpler. If H consists of just the vertex a, then it is protected by the guard moving back and forth between c and d. And if H consists of two disjoint cliques, either one of the cliques is just a or at least one of the cliques contains at least two vertices.

We now consider the situation where g is adjacent to a and not b. In this situation $b \in V(H)$ and $a \notin V(H)$. The argument is as in the proof of Lemma 20 if H is connected and has at least two vertices. If H is disconnected then it has exactly two components, each of which is a clique. If both of these have at least two vertices the argument is also as before. We must therefore consider the situations where H consists of: (i) only the vertex b; (ii) a clique of size at least two not containing b and another component consisting of the vertex b. (ii) a clique of size at least two containing b and another component consisting of a vertex z; (iv) non-adjacent vertices b and z. These are considered in turn.

Suppose H consists only of the vertex b. If b is adjacent to a vertex $w \in N(g) - D$ (from Lemma 19 (i)), then G is defended by three guards moving back and forth between b and w, c and d, and g and a, respectively. If b is not adjacent to any such w, then since G is 2-connected it must be adjacent to c and d; by hypothesis b is not adjacent to a or z. The graph G is then defended by two guards moving back and forth between c and d, and g and a, respectively.

Suppose H consists of a clique of size at least two not containing b and another component consisting of the vertex b. A single guard can defend the clique of size at least two not containing b without ever leaving it, and the rest of G can be defended as when H consists of only b.

Suppose H consists of a clique of size at least two containing b and another component consisting of a vertex z. A single guard can defend the clique of size at least two containing b without ever leaving it. If z is adjacent to both c and d, then G is defended by two more guards, one moving between c and d, and the other moving between g and a. Hence assume z is adjacent to at most one of c and d. Since G is 2-connected, the vertex z has degree at least two and is therefore adjacent to a vertex $w \in N(g) - D$ (recall that at least one such w exists), or to a. In the former situation, G is defended by three more guards moving between c and d, g and a, and z and w, respectively. In the latter situation it is defended by guards moving between c and d, g and w, and z and a, respectively.

Suppose H consists of non-adjacent vertices b and z. In each situation that follows, we will assume the existence of a guard moving back and forth between c and d. The possible neighbors of b are c, d and vertices $w \in N(g) - \{a, c, d\}$. The possible neighbors of z are c, d and vertices in $N(g)$ (including a). Since the guard at g can not relocate so that the vertices holding guards form a dominating configuration, for every vertex $w \in N(g) - \{a, c, d\}$ either $bw \in E$ or $zw \in E(G)$.

Consider the scenario where b is adjacent to c and d. Then b is guarded. If z is also adjacent to c and d, then G is defended by the guard moving between c and d and a guard moving between g and a. If z is adjacent to a, then G is guarded by the guard moving between c and d and guards moving between g and $w \in N(g) - \{a, c, d\}$, and z and a, respectively. And if z is adjacent to $w \in N(g) - \{a, c, d\}$, then G is defended by the guard moving between c and d and guards moving between z and w, and g and a, respectively.

It remains to consider the scenario where b is adjacent to at most one of c and d. Hence it is adjacent to at least one vertex $w \in N(g) - \{a, c, d\}$. If z is adjacent to both c and d, then G is guarded as when H consists only of

b. If z is adjacent to a, then either G has at most seven vertices and we are done as before, or there exists a vertex $w' \in N(g) - \{a, c, d\}$. In the latter situation, G is defended by the guard moving between c and d and guards moving between g and w', z and a, and b and w, respectively. If z is not adjacent to a, then it has at least one neighbor in $N(g) - \{a, c, d\}$. If w is the only vertex in $N(g) - \{a, c, d\}$, then G has at most seven vertices and we are done as before. Otherwise, since every vertex in $N(g) - \{a, c, d\}$ is adjacent to b or z and both of these vertices have a neighbor in this set, there exist different vertices $x, x' \in N(g) - \{a, c, d\}$ such that $bx, zx' \in E(G)$ and G is defended by the guard moving between c and d, and guards moving between g and a, b and x, and z and x', respectively. $\square$

Lemma 22 *Let G and D be as in Lemma 19. If g is independent of a and b, then $e^\infty(G) \leq 5$.*

Proof: As before, let $H = G - (\{c, d\} \cup N[g])$. By Lemma 19 (vi), $\beta(H) \leq 2$. In the situations that follow, there is a guard that moves back and forth between c and d. Note that $a, b \in V(H)$.

First suppose H is connected. Since $a, b \in V(H)$ and $ab \notin E$, the graph H has at least three vertices. If H has at least four vertices, then all vertices in the subgraph H can be defended by two guards who are never forced to move to a vertex not in $V(H)$. If H has three vertices, then it is a path on three vertices with ends a and b, say a, f, b. The vertex a is defended by the guard moving between c and d. Put a guard on b. When evicted, it moves back and forth on the edge bf.

Now suppose H is disconnected. Then it consists of two disjoint cliques, one or both of which might have only one vertex. Note that a and b are in different cliques, say L_a and L_b, respectively.

If L_b has at least two vertices, then a guard initially located on b can defend L_b without ever leaving it. If L_a also has at least two vertices then one guard initially on a can defend it. Otherwise, $V(L_a) = \{a\}$, which is defended by the guard moving between c and d.

Hence assume L_b has only one vertex, b. If b is adjacent to c and d, then L_b is defended by the guard moving between c and d. Since L_a can be defended as above, we may further assume $bd \notin E(G)$. Since G is 2-connected, $N(b) - \{c\} \neq \emptyset$. Hence b is adjacent to a vertex $w \in N(g) - \{c, d\}$. If there exists a vertex $w' \in N(g) \backslash \{c, d, w\}$, then G is defended by the guard moving between c and d, and guards moving between b and w, and g and w', respectively. Otherwise, G has only six vertices and, as before, it can be seen directly that $e^\infty \leq 5$. $\square$

Combining the results above yields the following theorem.

Theorem 23 *Let graph G have $\beta(G) = 3$. Then $e^\infty(G) \leq 5$.*

4 All-Guards Move

4.1 Bounds

Obviously $e^\infty m(G) \leq e^\infty(G)$, for all graphs G.

Theorem 24 *Let G be a connected graph. Then $e^\infty m(G) \leq \beta(G)$.*

Proof: Let I be a maximum independent set of G. Let D denote the configuration of guards in G, initially $D = I$. Suppose an attack occurs at vertex $v \in D$. If v has an external private neighbor u, then move the guard from v to u and observe the new configuration of guards is a maximum independent set (and thus a dominating set). So let us assume v has no external private neighbor. Let w be a neighbor of v. Move the guard from v to w; this configuration of guards is still a dominating set. Upon the next attack, move the guard from w back to v. Thus there is at most one guard outside a maximum independent set at any one time and the configuration of guards is a dominating set at all times. $\square$

4.2 Trees

We describe a partitioning scheme for trees that will determine e_m^∞. This scheme is inspired by the neo-colonization concept devised in [4] for the m-eternal domination number (see Section 1.3).

Let T be a tree. A *stem* of a tree T is a vertex of degree at least two that is adjacent to a leaf. We shall assume that stems have degree greater than one, otherwise T is a K_2. A vertex of T that is not a leaf is called an *internal* vertex.

A tree is a *star* if it is isomorphic to $K_{1,m}, m \geq 0$. The *center* vertex of a star with $m > 1$ is the vertex of degree m.

We partition the internal vertices of T into *loners*, *weak stems* and *strong stems* depending on whether they are adjacent to zero, exactly one or at least two leaves.

We give two simple examples to illustrate some of the nuances of the problem: consider the tree T_2 consisting of two $K_{1,2}$'s with their center vertices joined to a common vertex v. Note that this graph can be protected with four guards and no guard ever needs to occupy v. As another example, it is easy to see that $e_m^\infty(P_5) = 2$.

A *star partitioning* P is a partitioning of the vertex set of a tree T into parts that each induce a star, and such that no two K_1 parts are adjacent. It is easy to see that every tree has a star partitioning. We begin by establishing some properties of star partitionings.

Proposition 25 *Let T be a tree with at least two vertices. Then T has a star partitioning P such that (i) each weak stem of T and its adjacent leaf form a K_2 part of P, and (ii) any K_1 part is adjacent to at least two parts that are not K_1's.*

Proof: If statement (i) is false, then some leaf ℓ adjacent to a weak stem w is a K_1 part. By the definition, $\{w\}$ is not a part. If w is the center vertex of a non-trivial star, then ℓ can be absorbed into that star. Otherwise, P can be reconfigured so that it has the same number of parts and $\{w, \ell\}$ is a part. Statement (ii) follows from (i) and the definition. □

Call a star partitioning of T *special* if it satisfies the conditions in Proposition 25. We assign a star S in a special star partitioning weight $k - 1$, where k is the number of vertices of S. For a special star partitioning P of a tree T, define $wt(P)$, the *weight of P*, to be the sum of the weights of the stars in P.

Corollary 26 *Let P be a special star partitioning of a tree T with at least two vertices. Then no leaf of T is a part in P.*

Proposition 27 *Let P be a special star partitioning of a tree T with at least two vertices. Then T has an m-eternal dominating set in the eviction model with size $wt(P)$.*

Proof: Define $D \subseteq V(T)$ inductively as follows. Let $T_0 = T$, $P_0 = P$, and $D_0 = \emptyset$. Suppose D_i, P_i and T_i have been defined, that T_i contains at least two vertices, that D_i dominates $V(T) - V(T_i)$, and that P_i is a special star partitioning of T_i. If $V(T_i) = \emptyset$, then set $D = D_i$. Otherwise, define D_{i+1}, T_{i+1} and T_{i+1} as follows. Let S_i be a star in P_i that contains a stem adjacent to an end of a longest path in T_i. Then S_i contains a leaf

266

ℓ_i of T_i. Let L_i be the set of leaves of S_i. Let $D_{i+1} = D_i \cup (L_i - \{\ell_i\})$. By choice of S_i, at most one K_1 part of P_i is adjacent to a vertex of S_i. Let R_i be any such part; $V(R_i)$ is either empty or a singleton. Let $T_{i+1} = T_i - (V(S_i) \cup V(R_i))$. If $V(R_i) = \emptyset$, set $P_{i+1} = P_i - \{V(S_i)\}$, and if $V(R_i) \neq \emptyset$, then set $P_{i+1} = P_i - \{V(S_i), V(R_i)\}$. By choice of S_i and Corollary 26, we have that T_{i+1} is a tree with at least two vertices, that D_{i+1} dominates $V(T) - V(T_{i+1})$, and that P_{i+1} is a special star partitioning of T_{i+1}.

By construction, the set D is a dominating set of T containing $wt(P)$ vertices. We now show it is an eternal dominating set in the eviction model. Place guards on all vertices of D. Recall that attacks occur only at vertices with guards. We will maintain the invariant that guards will remain within their respective stars (from the partitioning P) at all times; hence no attack will ever occur at a vertex which forms a K_1 part.

Since guards remain on their original stars, each vertex in a non-trivial star from P will be protected at all times. Suppose there is an attack at a vertex $v \in D$ that is protecting a vertex w such that w is a K_1 part. Since the guard at v moves to another vertex in its star, it will not protect w. However, as w is adjacent to another star S of size at least two, a guard in S can move (if necessary) to a vertex within S that is adjacent to w. This may cause a "ripple effect" since this move may leave another K_1 part unprotected, but we can address that in a similar fashion (relocating a guard within star S') and since T is a tree, this sequence of guard moves will terminate with every vertex protected, since a K_1 part can not be a leaf of T. $\square$

Let $s(T)$ denote the minimum weight of any special star partitioning of T.

Proposition 28 *For any tree T with at least two vertices, $s(T) = e_m^\infty(T)$.*

Proof: We have $s(T) \geq e_m^\infty(T)$ by Proposition 27. We must show that $s(T) \leq e_m^\infty(T)$. The proof is by induction on the number of vertices of T. When T has two vertices, the proposition is trivial. Suppose T has more than two vertices. If T has a weak stem w with adjacent leaf u such that $T - \{u, w\}$ is connected, then any m-eternal dominating set must keep a guard on w or u at all times. Since $\{u, v\}$ must be a K_2 part (of weight one) in any minimum-weight special star partitioning, the result follows by induction. Suppose there is no such weak stem. Then there is a strong stem v with adjacent leaves $v_1, v_2, \ldots v_k$, $k \geq 2$. Then, in any m-eternal dominating set, there must be at least k guards on these $k+1$ vertices at all

times. The result follows by induction on observing that if v has a non-leaf neighbor z that is a loner, then z is either in a K_2 or a K_1 part in any minimum-weight star partitioning. $\square$

This result tells us, for example, that $e^\infty m(P_n) = \lceil \frac{n+1}{3} \rceil$, as every third vertex in a path can be a K_1 part.

5 Questions for Future Research

We state some questions in this section. Question 1 concerns bounds on e^∞ and Question 2 asks for characterizations. Question 1(i) seems the most fundamental.

Question 1 *(i) Is $e^\infty(G) \leq \gamma^\infty(G)$ for all graphs G?*
(ii) Is there a constant c such that $e^\infty(G) \leq c\beta(G)$, for all graphs G?

Question 2 *(i) Can we characterize the graphs with e^∞ equal to γ?*
(ii) Can we characterize the graphs with e^∞ less than or equal to β?
(iii) Can we describe additional graphs with e^∞ less than or equal to θ?

Recall that the *upper domination number* of G, denoted $\Gamma(G)$ is the size of a largest minimal dominating set.

Question 3 *Is there a constant c · (possibly $c = 1$) such that $e^\infty(G) \leq c\Gamma(G)$ for all connected graphs G?*

Question 4 *(i) Can the graphs with e_m^∞ equal to γ be characterized?*
(ii) Can the graphs with e_m^∞ equal to β be characterized?

Acknowledgments

References

[1] M. Anderson, C. Barrientos, R. Brigham, J. Carrington, R. Vitray, J. Yellen, Maximum demand graphs for eternal security, *J. Combin. Math. Combin. Comput.* **61** (2007), 111–128.

[2] K. S. Booth, and J. H. Johnson, Dominating Sets in Chordal Graphs, *SIAM J. Comput.* **11** (1982), 191–199.

[3] A.P. Burger, E.J. Cockayne, W.R. Gründlingh, C.M. Mynhardt, J.H. van Vuuren, W. Winterbach, Infinite order domination in graphs, *J. Combin. Math. Combin. Comput.* **50** (2004), 179–194.

[4] W. Goddard, S.M. Hedetniemi, S.T. Hedetniemi, Eternal security in graphs, *J. Combin. Math. Combin. Comput.* **52** (2005), 169–180.

[5] J. Goldwasser, W.F. Klostermeyer, Tight bounds for eternal dominating sets in graphs, *Discrete Math.* **308** (2008), 2589–2593.

[6] T. W. Haynes, S. T. Hedetniemi, P. J. Slater, *Fundamentals of Domination in Graphs*. Marcel Dekker, New York, 1998.

[7] W.F. Klostermeyer, G. MacGillivray, Eternal security in graphs of fixed independence number, *J. Combin. Math. Combin. Comput.* **63** (2007), 97–101.

[8] W.F. Klostermeyer, G. MacGillivray, Eternal dominating sets in graphs, *J. Combin. Math. Combin. Comput.* **68** (2009), 97–111.

[9] W.F. Klostermeyer, G. MacGillivray, Eternal dominating sets in trees, to appear in *J. Combin. Math. Combin. Comput.*

[10] W.F. Klostermeyer and C.M. Mynhardt, Graphs with Equal Eternal Vertex Cover and Eternal Domination Numbers, *Discrete Math.* **311** (2011), 1371–1379.

[11] W.F. Klostermeyer and C.M. Mynhardt, Vertex Covers and Eternal Dominating Sets, *Discrete Applied Mathematics* **160** (2012), pp. 1183–1190.

[12] W.F. Klostermeyer and C.M. Mynhardt, Protecting a Graph with Mobile Guards, arXiv:1407.5228 [math.CO]