A Characterization of Spartan Graphs and New Lower Bounds for Eternal Vertex Cover

Neeldhara Misra
IIT Gandhinagar, India

Saraswati Girish Nanoti
Indian Institute of Science, Bangalore, India

Abstract

The eternal vertex cover game is played between an attacker and a defender on an undirected graph G. The defender identifies k vertices to position guards initially. The attacker, on their turn, attacks an edge e, and the defender must move a guard along e to defend the attack. The defender may move other guards as well, under the constraint that every guard moves at most once and to a neighboring vertex. The smallest number of guards required to defend attacks forever is called the eternal vertex cover number of G, denoted $\mathsf{evc}(G)$.

For any graph G, $\mathsf{evc}(G)$ is at least $\mathsf{mvc}(G)$ (the vertex cover number of G). A graph is Spartan if $\mathsf{evc}(G) = \mathsf{mvc}(G)$. It is known that a bipartite graph is Spartan if and only if every edge belongs to a perfect matching. We show that the only König graphs that are Spartan are the bipartite Spartan graphs. We also give new lower bounds for $\mathsf{evc}(G)$, generalizing a known lower bound based on cut vertices. We finally show a new matching-based characterization of all Spartan graphs.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Eternal Vertex Cover, Vertex Cover, König Graphs, Spartan Graphs, Matchings

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2025.45

Related Version *ArXiV*: https://arxiv.org/abs/2504.06832

Funding *Neeldhara Misra*: Supported by the SERB Early Career Researcher Grant ECR/2018/002967.
Saraswati Girish Nanoti: Supported by CSIR.

Acknowledgements This work was done when the second author was a student at IIT Gandhinagar. The second author thanks CSIR for the support.

1 Introduction

Recall that a subset of vertices S is called a *vertex cover* if, for every edge $e = \{u, v\}$, at least one of u or v belongs to S. The eternal vertex cover game, introduced by Klostermeyer and Mynhardt [4], is a turn-based two-player game played on a graph between players typically referred to as the attacker and the defender. The defender is tasked with placing guards on the vertices of the graph to protect against attacks on any of its edges, while the attacker selects edges to attack. The defender's goal is to ensure that every attack can be countered by moving a guard along the attacked edge, while optionally moving additional guards, under the constraint that guards move at most once and to a neighboring vertex. When the defender is able to defend attacks forever, the positions of the guards at every stage of the game form a vertex cover.

The minimum number of guards required for the defender to successfully defend against an infinite sequence of attacks is known as the *eternal vertex cover number*, denoted as $\mathsf{evc}(G)$. We also use $\mathsf{mvc}(G)$ to denote the *minimum* vertex cover number of G, which is the size of a smallest vertex cover of G. It is well known that [4]:

$$\underbrace{\mathsf{mvc}(G) \leqslant \mathsf{evc}(G)}_{\text{The placement of guardsis a vertex cover.}} \qquad \text{and} \qquad \underbrace{\mathsf{evc}(G) \leqslant 2 \cdot \mathsf{mvc}(G)}_{\text{Guard both endpoints of a maximum matching.}}$$

A natural question is to identify the graphs for which the lower bound is tight, i.e., when $\mathsf{evc}(G) = \mathsf{mvc}(G)$. Such graphs are called *Spartan* – here, the defender can maintain a vertex cover with the minimum number of guards, making them an optimal solution in terms of resource usage. We are interested in the question of what such graphs look like from a structural perspective.

In [2], Babu et al. obtain a characterization for a class of graphs that includes chordal graphs and internally triangulated planar graphs. We briefly recall this result. Let the graph class $\mathcal{F}$ denote the class of all connected graphs G for which each minimum vertex cover of G that contains all the cut vertices of G induces a connected subgraph in G. (A cut vertex is a vertex whose removal disconnects the graph.) Let $G(V, E)$ be a graph that belongs to $\mathcal{F}$, with at least two vertices, and $X \subset V$ be the set of cut vertices of G. Then it is shown [2] that G is Spartan if and only if for every vertex $v \in V \backslash X$, there exists a minimum vertex cover S_v of G such that $X \cup \{v\} \subset S_v$. It is also known that bipartite graphs are Spartan if and only if every edge belongs to a perfect matching (essentially elementary) [6]. Our first result involves expanding the scope of this characterization to König graphs, which are graphs where the minimum vertex cover equals the maximum matching, and hence a natural generalization of bipartite graphs. We show that the only König graphs that are Spartan are those that are also bipartite and satisfy the conditions for being Spartan within bipartite graphs:

▶ **Theorem 1.** *A König-Egerváry graph* G *is Spartan if and only if it is bipartite and essentially elementary.*

We also develop a series of increasingly stronger pre-requisites – i.e., necessary conditions – for a graph to be Spartan. To begin with, note that for a graph G which has more than one vertex, if $\mathsf{evc}(G) = \mathsf{mvc}(G)$, then every vertex v of G must belong to a minimum-sized vertex cover S_v. Indeed, if not, then from any configuration, attacking an edge incident to v will result in the attacker winning. A result in [2] takes this further to show that if $\mathsf{evc}(G) = \mathsf{mvc}(G)$, then each vertex must belong to a minimum-sized vertex cover which contains all the cut vertices.

We generalize this idea further in the following way. Let S be a vertex cover of G and let $I = V(G) \setminus S$ be an independent set of G. A "bad situation" for the defender who has their guards positioned on the vertices of S is the following: there is a $T \subseteq I = V(G) \setminus S$ such that one of the connected components C of $G \setminus T$ is such that $|V(C) \cap S| = \mathsf{mvc}(C)$. This is because the attacker can attack an edge that has one endpoint in C and the other in T, forcing a guard out of C, and it is easy to check that the shortfall in $G[C]$ "cannot be fixed", creating a vulnerable edge that will lead to a win for the attacker (see Figure 1 where the defender has one guard on the vertex q with respect to the attack qb).

A vertex cover S is *weakly good* if there is no subset T of $V(G) \setminus S$ for which the bad situation described here occurs. We show that a graph G is Spartan only if every vertex is contained in a minimum-sized and weakly good vertex cover.

From the defender's perspective, sometimes it is not enough to just avoid the bad situation described above with the vertex covers that they work with. Indeed, suppose we have a component C as shown on the right side of Figure 1. Here, $|V(C) \cap S| > \mathsf{mvc}(C)$, but note

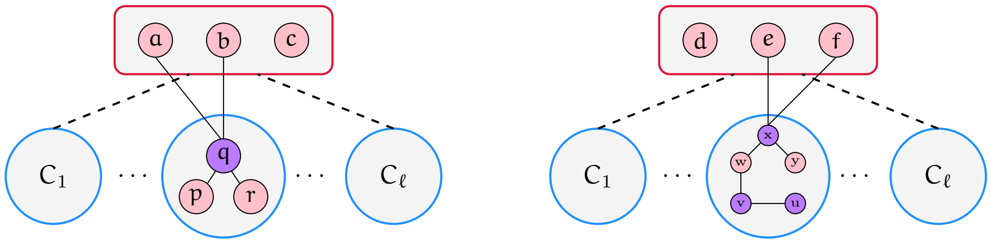

Figure 1 Examples of bad situations. Left: an example of a bad scenario with $T = \{a, b, c\}$. Note that the minimum vertex cover number of the component with the vertices $\{p, q, r\}$ is one and the size of its intersection with the vertex cover of G, i.e., $\{q\}$ is also one. Right: an example of a bad scenario with $T = \{d, e, f\}$. Here, the minimum vertex cover number of the component with the vertices $\{u, v, w, x, y\}$ is two and the intersection of this component with the vertex cover of G is $\{x, u, v\}$; however the set $\{u, v\}$ is not compatible with any minimum-sized vertex cover of the component.

that if the attacker attacks xe, then the defender is stuck with guards on v and u: while in principle $\mathsf{mvc}(C) = 2$, these two guards cannot be moved in a way that will defend all edges in the component C, and as before there is no pathway for guards from outside C to enter C.

While not as stark as before, this is a more nuanced bad scenario. Informally speaking, vertex covers that avoid even this scenario are called *strongly good* vertex covers. We show that a graph is Spartan only if every vertex is contained in a minimum-sized and strongly good vertex cover.

Our final and main result is a characterization of Spartan graphs at large. Let us begin by making a defender's strategy structurally explicit. To this end, consider a graph where we are able to find a non-empty family $\mathcal{F}$ of minimum-sized vertex covers such that the following holds:

For every $S \in \mathcal{F}$ *and for every edge* $uv \in G$ *such that* $u \in S$ *and* $v \notin S$*, there exists* $T \in \mathcal{F}$ *such that* $v \in T$ *and* T *is reachable from* S *via guard movement while defending* uv*: this is to say that if the defender had guards positioned on* S*, then they will be able to move them in such a way that the guards finally end up on* T *and one of the guards moves along the edge* uv*.*

Notice that such a family $\mathcal{F}$ naturally translates to a strategy for the defender: the defender can start with guards positioned on S for any $S \in \mathcal{F}$, and if any edge uv is attacked, the defender can either exchange guards (if both u and v are in S) or move guards to the T guaranteed by the property above. This would work indefinitely since every vertex cover in $\mathcal{F}$ offers the same opportunities to the defender.

This brings us to the question of capturing movements between minimum-sized vertex covers S and T. We refer to the vertices in $V \setminus (S \cup T)$ as the *dead zone*. The defender hopes to move guards currently positioned at S to a new target vertex cover T while also defending an arbitrary but fixed edge uv with (say) $u \in S \setminus T$ and $v \in T \setminus S$. Notice that the defender can transform the configuration of guards from S to T while defending uv if there are $|S \setminus T| - 1$ vertex-disjoint paths between $(S \setminus T) \setminus \{u\}$ and $(T \setminus S) \setminus \{v\}$ that do not use any vertices from the dead zone. We say that S and T are *mutually reachable* if such paths exist.

It is intuitive that the existence of a family of mutually reachable minimum vertex covers is a necessary and sufficient condition for a graph to be Spartan: indeed, given such a family, the defender has a strategy, and if the defender has a strategy, then the guard positions naturally correspond to such a family. Our main contribution is to capture the requirement of mutual reachability between vertex covers S and T in terms of matchings in an appropriately

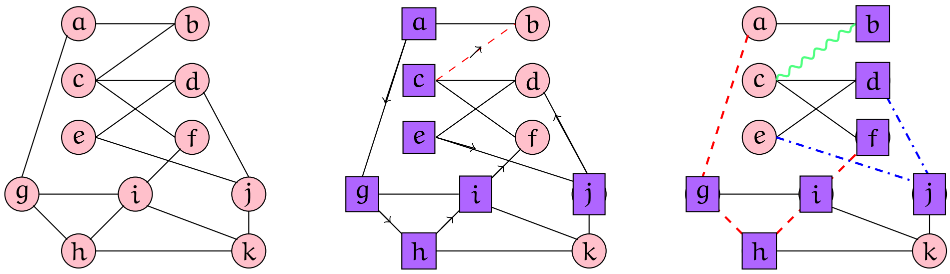

Figure 2 Demonstrating reachability between vertex covers $S = \{a, c, e, g, h, i, j\}$ and $T = \{b, d, f, g, h, i, j\}$ with the vertex $\{k\}$ in the dead zone. Here $S \setminus T = \{a, c, e\}$ and $T \setminus S = \{b, d, f\}$. There are three vertex-disjoint paths $P_1 = aghif$, $P_2 = ejd$ and $P_3 = cb$ (the attacked edge), from $S \setminus T$ to $T \setminus S$ with internal vertices from $S \cap T$. Thus, the guard on i moves to f, the guard on h moves to i, the guard on g moves to h and the guard on a moves to g. Notice that this movement happens along the path P_1. Similar movements happen along P_2 and P_3 and thus the guards reconfigure from S to T. Note that no vertex of P_1, P_2 and P_3 is common, if there was a common vertex then the guard which just arrived on the vertex could not have moved again before the next turn. Hence it is important that the paths P_1, P_2 and P_3 are vertex-disjoint when we want to use them to reconfigure from S to T.

defined auxiliary graph $\mathfrak{h}_G(S, T)$ based on S and T. The notion of the auxiliary graph required to translate the condition of mutual reachability in terms of matchings is somewhat technical and we omit the specifics from the introduction. The non-trivial technical aspect is to ensure that the reachability guaranteed by disjoint paths as described above is in fact captured by the existence of a matching in $\mathfrak{h}_G(S, T)$.

While it is not clear that our characterization is efficiently testable, we propose it with the goal of offering structural insight into the class of general Spartan graphs, and hope that this lays the groundwork for future work on algorithmic aspects of recognizing Spartan graphs.

2 Preliminaries

For any natural number $n \in \mathbb{N}$, the set $[1, n]$ denotes all the natural numbers between 1 and n (both included). Let $G = (V, E)$ be a simple, finite, undirected graph with $n > 1$ vertices and m edges (unless mentioned otherwise). Also unless mentioned otherwise, we will assume that the graph G is connected (and in particular, has no isolated vertices). We use the standard notation from [3]. If S is a subset of V, $G[S]$ denotes the subgraph of G induced by the vertices in S. The set of all vertices v such that $uv \in E(G)$ is denoted by $N(u)$ and is said to be the *open neighbourhood* of u. The set $N(u) \cup \{u\}$ is said to be the *closed neighbourhood* of u.

A *path* on n vertices is a graph P with vertices $\{v_1, v_2, \ldots, v_n\}$ such that each $(v_i, v_{i+1}) \in E(P)$ for $i \in [1, n-1]$. A *cycle* on n vertices is a graph C with vertices $\{v_1, v_2, \ldots, v_n\}$ such that each $(v_i, v_{i+1}) \in E(C)$ for $i \in [1, n-1]$ and $(v_n, v_1) \in E(C)$. A graph $G(V, E)$ is said to be *connected*, if for every pair of distinct vertices $u, v \in V$, there exists a path between u and v.

A subset S of $V(G)$ is said to be a *vertex cover* of G if for every edge $(u, v) \in E$, either $u \in S$ or $v \in S$ (or both). The size of a smallest vertex cover of the graph G is called the *minimum vertex cover number* of G and denoted by $\mathsf{mvc}(G)$. The graphs with $\mathsf{evc}(G) = \mathsf{mvc}(G)$ are known as *Spartan graphs*.

A subset S of $V(G)$ is said to be an *independent set* of G if no edge of $E(G)$ has both its endpoints in S. It can be seen that the complement of a vertex cover will be an independent set and vice versa.

A *matching* is a set of edges with no common endpoint. A *perfect matching* is a matching which contains one edge adjacent to each vertex of a graph. A graph may not always have a perfect matching. A matching of the largest cardinality in a graph is called a *maximum matching* and the size of this matching is denoted by $mm(G)$.

A *bipartite graph* $G = (V, E)$ is a graph whose vertex set can be partitioned into two independent sets, say $V = (A \cup B)$ such that every edge is between a vertex in A and one in B. Clearly, both A and B are vertex covers of G. If a bipartite graph $G = (A \cup B, E)$ is connected and its *only* minimum-sized vertex covers are A and B, then we say that G is *elementary*. If every connected component of a bipartite graph is elementary, then we call it *essentially elementary*.

In Section 4, we need to use the following result about bipartite graphs which is given in [5] and can also be obtained by a slight modification of a result in [6]. We give an explicit proof for completeness.

▶ **Lemma 2** ([6]). *Let G be a connected bipartite graph such that $V(G) = A \cup B$ where A and B are non-empty independent sets. If G has at least one perfect matching and G is non-elementary, then there exist non-empty $S \subsetneq A$ and $T \subsetneq B$ such that $|N(S)| = |S|$ and $|N(T)| = |T|$.*

Proof. Let G be a non-elementary connected bipartite graph with $V(G) = A \cup B$, where A and B are non-empty independent sets, and a perfect matching. Suppose that for every non-empty subset S of A, $S \subsetneq A$ we have $|N(S)| > |S|$. Let $ab \in E(G)$. We show that there exists a perfect matching between $A \setminus \{a\}$ and $B \setminus \{b\}$. Consider any non-empty $X \subseteq A \setminus \{a\}$. If $|N(X) \cap (B \setminus \{b\})| < |X|$, then $|N(X) \cap B| \leqslant |X|$. That is, $X \subsetneq A$ such that $|N(X)| \leqslant |X|$, which is contrary to our assumption that for every subset S of A, $S \subsetneq A$ we have $|N(S)| > |S|$. Thus, we cannot have any $X \subseteq A \setminus \{a\}$ such that $|N(X) \cap (B \setminus \{b\})| < |X|$. Therefore, by Hall's theorem, there exists a perfect matching M_1 between $A \setminus \{a\}$ and $B \setminus \{b\}$. Thus, $M = M_1 \cup \{ab\}$ is a perfect matching of G containing ab. Since ab was an arbitrary edge of G, every edge of G belongs to some perfect matching, and G is connected. This implies that G is elementary which is a contradiction. Therefore, there must exist a non-empty subset S of A such that $S \subsetneq A$ and $|N(S)| \leqslant |S|$ but $|N(S)|$ cannot be less than $|S|$ because G has a perfect matching. Thus $|N(S)| = |S|$. With a symmetric argument, we get that there exists a non-empty $T \subsetneq B$ such that $|N(T)| = |T|$. ◀

3 New Lower Bounds

In this section we give some necessary conditions which a graph G must satisfy, in order to have $\mathsf{evc}(G) = \mathsf{mvc}(G)$. One such condition is stated in [2] which is as follows:

▶ **Lemma 3** ([2]). *For a graph G which has more than one vertex, if $\mathsf{evc}(G) = \mathsf{mvc}(G)$ then every vertex v of G must belong to a minimum-sized vertex cover S_v.*

For understanding the nature of Spartan graphs, it is sufficient to only look at connected components. The lemma below which was proven in [6] justifies this.

▶ **Lemma 4** ([6]). *Let G be a Spartan graph with connected components $C_1, \ldots, C_\ell$. G is Spartan if and only if $G[C_i]$ is Spartan for all $1 \leqslant i \leqslant [\ell]$.*

Thus, we now look at only connected graphs.

▶ **Lemma 5.** *If a graph* G *with more than one vertex is Spartan and* I *is a maximum independent set of* G*, then there is a matching of* I *to* $V(G) \setminus I$*, which saturates* I*.*

Proof. Consider a Spartan graph G (with more than one vertex), and let I be a maximum independent set of G and $C = V(G) \setminus I$ be a minimum-sized vertex cover. Suppose there is no matching from I to S which saturates I, then by Hall's theorem, there exists $X \subseteq I$ such that $|N(X)| < |X|$. (Here, the neighborhood of any subset of I in the entire graph is the same as the neighborhood of that subset in C because I is an independent set.)

Consider an inclusion-wise minimal such subset X of I. Then for every $x \in X$, there exists a perfect matching from $X \setminus \{x\}$ to $N(X)$. If not, then there exists a $Y \subseteq X \setminus \{x\}$ such that $|N(Y)| < |Y|$ and as $Y \subseteq X \setminus \{x\}$ which means that $Y \subset X$, this contradicts the inclusion-wise minimality of X. Hence, there exists a perfect matching from $X \setminus \{x\}$ to $N(X)$ and this also means that $|X| = |N(X)| + 1$. Next, we make a claim which demonstrates how any vertex cover of G intersects the vertices in $X \cup N(X)$ and its proof is given in the Appendix.

▷ Claim 6. Any minimum-sized vertex cover of G cannot contain more than $|N(X)|$ vertices from $X \cup N(X)$.

Now consider a vertex $x \in X$. Since G is Spartan, there is a minimum-sized vertex cover S_x of G such that $x \in S_x$ by Lemma 3. We have already shown that $X \setminus \{x\}$ has a perfect matching with $N(X)$. Thus, any vertex cover must contain $|N(X)|$ vertices from $X \setminus \{x\} \cup N(X)$. Thus, S_x contains x and $|N(X)|$ vertices from $X \setminus \{x\} \cup N(X)$. That is, S_x contains $|N(X)|$ vertices from $X \cup N(X)$. This contradicts the above claim. Thus, the existence of a Hall's violator set X is not possible and there exists a matching from I to $S = V(G) \setminus I$ which saturates I. ◀

This immediately gives us a simple corollary (proof in Appendix).

▶ **Corollary 7.** *If* G *is a Spartan graph with more than one vertex and* $|V(G)| = n$*, then* $\mathsf{mvc}(G) \geqslant n/2$*.*

▶ **Definition 8.** *A vertex cover* S *of a graph* G *is said to be a* weakly good *vertex cover if it satisfies the following: Let* $I = V(G) \setminus S$ *be the corresponding independent set. For every* $T \subseteq I$*, if* $C_1, C_2, \ldots C_\ell$ *are the connected components of* $G[V(G) \setminus T]$*, we must have* $|V(C_i) \cap S| > \mathsf{mvc}(C_i)$ *(for all* $i \in [1, \ell]$*).*

When we are dealing with more than $\mathsf{mvc}(G)$ guards, and we are in the model of the game where more than one guard is able to occupy a single vertex, then we will need to look at configurations of guards, not just vertex covers or vertex sets. A configuration of guards is a description of how many guards are present on each vertex of the graph G. A configuration can also be viewed as a function which maps each vertex to a non-negative integer such that the sum of the values of this function on all the vertices of G is equal to the number of guards. At some places, we will also view a configuration as a multi-set where a vertex appears as many times as the number of guards present on it. The manner in which we are using the word configuration, will be clear from context. When we say a configuration of guards $\mathcal{C}$ on a vertex cover S, we mean that the vertices having one or more guards in $\mathcal{C}$ (i.e., a non-zero value for the configuration function) are precisely the ones in S.

▶ **Definition 9.** *A configuration of guards* $\mathcal{C}$ *on a vertex cover* S *is said to be a* weakly good *configuration if it satisfies the following: Let* $I = V(G) \setminus S$ *be the corresponding independent set. For every* $T \subseteq I$*, if* $C_1, C_2, \ldots C_\ell$ *are the connected components of* $G[V(G) \setminus T]$*, the number of guards in* C_i *must be greater than* $\mathsf{mvc}(C_i)$ *(for all* $i \in [1, \ell]$*).*

If the number of guards on G is the same as $\mathsf{mvc}(\mathsf{G})$, or we are working in the model where only one guard per vertex is allowed, then the vertex cover formed by vertices occupied by the guards in a weakly good configuration is a weakly good vertex cover. It is easier to see the notion of a weakly good vertex cover (or configuration) by looking at a vertex cover (or configuration) which is not weakly good.

▶ **Definition 10.** *Let S be a vertex cover of a graph G such that S is not weakly good. There exists a non-empty $T \subseteq I = V(G) \setminus S$ such that if $C_1, C_2, \dots C_\ell$ are the connected components of $G[V(G) \setminus T]$, we must have $|V(C_i) \cap S| = \mathsf{mvc}(C_i)$ for some C_i (where $i \in [1, \ell]$). Then we will denote T as a* weakly bad set *corresponding to $\mathcal{C}$.*

▶ **Definition 11.** *Let S be a vertex cover of a graph G and let $\mathcal{C}$ be a configuration of guards on S such that $\mathcal{C}$ is not weakly good. There exists a non-empty $T \subseteq I = V(G) \setminus S$ such that if $C_1, C_2, \dots C_\ell$ are the connected components of $G[V(G) \setminus T]$, there must exist a C_i (where $i \in [1, \ell]$) such that the number of guards in C_i must be equal to $\mathsf{mvc}(C_i)$ (cannot be less than $\mathsf{mvc}(C_i)$ because S is a vertex cover). Then we will denote T as a* weakly bad set *corresponding to S.*

For every vertex cover S which is not weakly good, there must exist at least one corresponding weakly bad set $T \subset V(G) \setminus S$. Similarly for every configuration $\mathcal{C}$ which is not weakly good, there must exist at least one corresponding weakly bad set $T \subset V(G) \setminus S$. Conversely, if for a vertex cover S (or a configuration $\mathcal{C}$ occupying the vertices in a vertex cover S), there exists a subset T of $V(G) \setminus S$ which is weakly bad, then the vertex cover S (or the configuration $\mathcal{C}$) is not weakly good.

▶ **Lemma 12.** *The $\mathsf{evc}(G)$ is at least k, where $k \in \mathbb{N}$ such that every vertex of G belongs to a weakly good configuration of size at most k. This means that a graph G has $\mathsf{evc}(G) = k$ only if for each $v \in V(G)$, there exists a weakly good configuration $\mathcal{C}_v$ corresponding to a vertex cover S_v such that $\mathcal{C}_v$ has k guards and $v \in S_v$. In particular, a graph G is Spartan only if for each $v \in V(G)$, there exists a minimum-sized and weakly good vertex cover S_v such that $v \in S_v$.*

Proof. Suppose a graph G has $\mathsf{evc}(G) = k$, and there exists $v \in V(G)$ such that there does not exist a weakly good configuration with k guards with a guard on v. Without loss of generality, we can assume that the defender starts with a configuration with a guard on v, because the attacker can always attack an edge adjacent to v, and ensure that at least one guard comes to v. As there exists no weakly good configuration with k guards, which has a guard on v, the configuration $\mathcal{C}$ formed by the guards cannot be a weakly good configuration. If the guards do not occupy a vertex cover, some edge must be vulnerable and the attacker can attack that edge. So, we can assume that the set S of vertices occupied by the guards in the configuration $\mathcal{C}$ forms a vertex cover. Therefore, there exists a weakly bad subset T of the unoccupied vertices such that if $C_1, C_2, \dots, C_\ell$ are the connected components of $G[V(G) \setminus T]$, we must have a C_i (where $i \in [1, \ell]$), which contains $\mathsf{mvc}(C_i)$ guards. Now, since the components are not adjacent to each other, there exists an edge from a vertex $w \in V(C_i)$ to a vertex $u \in T$ because the graph G is connected. Since no vertex in T has a guard, the guard on w must come to u. Therefore, some edge in C_i will become vulnerable, as there were only $\mathsf{mvc}(C_i)$ many guards in $V(C_i)$, and one guard has moved out of C_i, while no guard from outside can come to a vertex of C_i in one step. Thus, the attacker wins, and this contradicts $\mathsf{evc}(G) = k$. ◀

▶ **Lemma 13.** *Let S be a minimum-sized vertex cover of a graph G, such that there exists a cut vertex x of G such that $x \notin S$. Then S cannot be a weakly good vertex cover.*

Proof. Let x be a cut vertex of a graph G, and let S be a vertex cover of G which does not contain x. Then we show that $\{x\} \subset V(G) \setminus S$ is weakly bad, which will imply that S is not a weakly good vertex cover. Since x is a cut vertex, there exist connected components $C_1, C_2, \ldots, C_\ell$ where $\ell \geqslant 2$ of $V(G) \setminus \{x\}$. If $|S \cap V(C_i)| > \mathsf{mvc}(C_i)$ for all $i \in [1, \ell]$, then $S' = \{x\} \cup S_1 \cup S_2 \cup \ldots \cup S_\ell$ will be a vertex cover of G of size $|S| - \ell + 1$ (where S_i is a vertex cover of C_i of size $\mathsf{mvc}(C_i)$). This is not possible as $|S'| < |S|$ and S is a minimum-sized vertex cover of G. Thus, there exists some $i \in [1, \ell]$ such that $|S \cap V(C_i)| = \mathsf{mvc}(C_i)$, which implies that $\{x\}$ is a weakly bad set. Hence, S is not a weakly good vertex cover. ◀

Thus, with Lemma 12 and Lemma 13, we get the result in [2] that states that if $\mathsf{evc}(G) = \mathsf{mvc}(G)$, then each vertex must belong to a minimum-sized vertex cover which contains all the cut vertices.

We next define a generalization of a bad set, which is used to define the notion of a strongly good vertex cover. We use this notion to give another necessary condition for a graph G to be Spartan. In order to define this, we define the notion of compatible vertex sets (or vertex covers) in a graph. We also give analogous definitions of compatible configurations and strongly good configurations. This gives us a new lower bound for $\mathsf{evc}(G)$.

▶ **Definition 14.** *Two vertex sets (vertex covers) S_1 and S_2 of a graph G are said to be* compatible*, if there exist $|S_1 \cap S_2|$–many vertex-disjoint paths between $S_1 \setminus S_2$ and $S_2 \setminus S_1$.*

▶ **Definition 15.** *Two configurations $\mathcal{C}_1$ and $\mathcal{C}_2$ are said to be* compatible*, if there exist $|\mathcal{C}_1 \cap \mathcal{C}_2|$–many vertex disjoint paths (counting multiplicity) between $\mathcal{C}_1 \setminus \mathcal{C}_2$ and $\mathcal{C}_2 \setminus \mathcal{C}_1$. Here we view each configuration as a multi-set over the set of vertices $V(G)$.*

Also, when we say $\mathcal{C}_i \setminus \mathcal{C}_j$, we account for multiplicity of each vertex. That is, if v occurs 5 times in $\mathcal{C}_1$ and 3 times in $\mathcal{C}_2$, v will occur 2 times in $\mathcal{C}_1 \setminus \mathcal{C}_2$. Also, when we say vertex-disjoint paths (counting multiplicity) between $\mathcal{C}_1 \setminus \mathcal{C}_2$ and $\mathcal{C}_2 \setminus \mathcal{C}_1$, we mean that if a vertex v occurs multiple times in $\mathcal{C}_1 \setminus \mathcal{C}_2$, each occurrence of v will be an endpoint of exactly one path from $\mathcal{C}_1 \setminus \mathcal{C}_2$. Similarly, if a vertex v occurs multiple times in $\mathcal{C}_2 \setminus \mathcal{C}_1$, each occurrence of v will be an endpoint of exactly one path to $\mathcal{C}_2 \setminus \mathcal{C}_1$. If we consider the union of all the multisets formed by the intermediate vertices (i.e., the vertices of $\mathcal{C}_1 \cap \mathcal{C}_2$) of each path, the number of times a vertex v occurs in the union, should be less than or equal to the number of times v occurs in $\mathcal{C}_1 \cap \mathcal{C}_2$.

It is clear that two configurations $\mathcal{C}_1$ and $\mathcal{C}_2$ are compatible, if and only if there is a possible movement of guards from $\mathcal{C}_1$ to $\mathcal{C}_2$, where each guard moves at most one step. The guards can rearrange themselves by moving one step from $\mathcal{C}_1 \setminus \mathcal{C}_2$ to $\mathcal{C}_2 \setminus \mathcal{C}_1$ if and only if there exist $|\mathcal{C}_1 \cap \mathcal{C}_2|$–many vertex disjoint paths (counting multiplicity) between $\mathcal{C}_1 \setminus \mathcal{C}_2$ and $\mathcal{C}_2 \setminus \mathcal{C}_1$. For more explanation on the equivalence between the guard movements and vertex-disjoint paths, please refer to Figure 2.

The following is a known result shown in [2].

▶ **Lemma 16** ([2])**.** *Two minimum-sized vertex covers of a graph G are always compatible.*

▶ **Definition 17.** *For a vertex cover S, a set $T \subset V(G) \setminus S$ is* strongly bad*, if there is some connected component C_i of $V(G) \setminus T$ and some $v \in N(T) \cap V(C_i)$, such that no vertex cover of C_i of size $|S \cap V(C_i)| - 1$ is compatible with $(S \cap V(C_i)) \setminus \{v\}$.*

▶ **Definition 18.** *A vertex cover S of a graph G is said to be* strongly good*, if $V(G) \setminus S$ does not have a strongly bad subset.*

Now we analogously define a strongly good configuration.

▶ **Definition 19.** *For a configuration $\mathcal{C}$ of guards on a vertex cover S, a set $T \subset V(G) \setminus S$ is* strongly bad*, if there is some connected component C_i of $V(G) \setminus T$ and some $v \in N(T) \cap V(C_i)$, such that no configuration on a vertex cover of C_i of size $|\mathcal{C} \cap V(C_i)| - 1$ is compatible with $(\mathcal{C} \cap V(C_i)) \setminus \{v\}$.*

▶ **Definition 20.** *A configuration $\mathcal{C}$ on a vertex cover S of a graph G is said to be* strongly good*, if $V(G) \setminus S$ does not have a strongly bad subset.*

Now we make an observation which relates the notions of weakly good and strongly good configurations (vertex covers).

▶ **Lemma 21.** *A strongly good configuration (vertex cover) is a weakly good configuration (vertex cover).*

Proof. We prove the contrapositive of this statement, i.e., a configuration on a vertex cover which is not a weakly good configuration, is also not a strongly good configuration. Let $\mathcal{C}$ be a configuration on a vertex cover S, such that $\mathcal{C}$ is not weakly good. Therefore, there exists a non-empty $T \subseteq I = V(G) \setminus S$, such that if $C_1, C_2, \ldots, C_\ell$ are the connected components of $G[V(G) \setminus T]$, there must exist a C_i (where $i \in [1, \ell]$) such that the number of guards in C_i (in $\mathcal{C}$) is equal to $\mathsf{mvc}(C_i)$. Since T is non-empty and the graph G is connected, there exists $v \in V(C_i)$ such that v is adjacent to T and since S is a vertex cover v is in S (and has exactly one guard in $\mathcal{C}$ as the number of guards in C_i (in $\mathcal{C}$) is equal to $\mathsf{mvc}(C_i)$). The configuration $\mathcal{C} \cap C_i \setminus \{v\}$ is not compatible with any vertex cover of C_i, as the number of guards in $\mathcal{C} \cap C_i \setminus \{v\}$ is less than $\mathsf{mvc}(C_i)$. Thus the set T is a strongly bad set, and $\mathcal{C}$ is not a strongly good configuration.

The same proof works for vertex covers, if we consider a vertex cover as a configuration with one guard per vertex. ◀

▶ **Lemma 22.** *The $\mathsf{evc}(G)$ is at least k, where $k \in \mathbb{N}$ such that every vertex of G belongs to a strongly good configuration of size at most k. This means that graph G has $\mathsf{evc}(G) = k$, only if for each $v \in V(G)$, there exists a $k-$sized strongly good configuration $\mathcal{C}_v$ on a vertex cover S_v, such that $v \in S_v$. In particular, a graph G is Spartan if for each $v \in V(G)$, there exists a minimum-sized and strongly good vertex cover S_v, such that $v \in S_v$.*

Proof. Suppose that there exists a graph G with $\mathsf{evc}(G) = k$ and a vertex v, such that there exists no $k-$sized strongly good configuration on a vertex cover containing v. Without loss of generality, we can assume that the initial configuration has at least one guard on v, because if not, the attacker can force a guard to come to v by attacking an edge adjacent to v. If the vertices occupied by the guards in the resulting configuration do not form a vertex cover, then there exists some edge which is vulnerable, and hence the attacker wins. Therefore, the resulting configuration must have guards on a vertex cover. Since there is no $k-$sized strongly good configuration $\mathcal{C}_v$ on a vertex cover S_v such that $v \in S_v$, for the configuration $\mathcal{C}$ formed by the guards on the vertex cover S, there exists a strongly bad subset $T \subset V(G) \setminus S$. Hence, there is some connected component C_i of $V(G) \setminus T$ and some $u \in N(T) \cap V(C_i)$ such that no configuration on a vertex cover of C_i of size $\mathcal{C} \cap V(C_i) - 1$ is compatible with $S \cap V(C_i) \setminus \{u\}$. Suppose the attacker attacks an edge joining u to a vertex in T, the guard on u is forced to move to T. No guard from a vertex outside C_i can come to C_i and the guards in C_i cannot rearrange themselves to form a vertex cover of C_i. Therefore, some edge will be vulnerable no matter how the guards rearrange themselves, which contradicts $\mathsf{evc}(G) = k$.

Therefore, a graph G has $\mathsf{evc}(G) = k$, only if for each $v \in V(G)$, there exists a $k-$sized strongly good configuration $\mathcal{C}_v$ on a vertex cover S_v where $v \in S_v$. ◀

4 König–Egerváry Graphs

A graph G is said to be a *König–Egerváry* graph, if the size of a smallest vertex cover of G is equal to the size of the largest matching of G, i.e., $\mathsf{mvc}(G) = \mathfrak{mm}(G)$. This class is a natural and strict[1] generalization of bipartite graphs. In this section, we derive the necessary and sufficient condition for a König-Egerváry graph to be Spartan. Before that, we have a corollary of Lemma 12, which gives a sense of how a Spartan graph will look like. The proof is given in the Appendix.

▶ **Corollary 23.** *If a graph G with more than one vertex is Spartan and I is an independent set of G which is not maximal, then $|N(I)| > |I|$.*

The main result of this section is the following:

▶ **Theorem 1.** *A König-Egerváry graph G is Spartan if and only if it is bipartite and essentially elementary.*

Proof. As described in Lemma 4, it suffices to look at connected graphs i.e., it will be sufficient to prove that a connected König-Egerváry graph G is Spartan if and only if it is bipartite and elementary. The reverse direction has already been proved in [1]. Thus we need to only show that if a connected König-Egerváry graph G is Spartan, then it is bipartite and elementary.

Since G is Spartan, let S be a minimum-sized vertex cover of G, which is an eternal vertex cover configuration, and let $I = V(G) \setminus S$ be the corresponding independent set. By Lemma 5, there exists a matching M from I to S, which saturates I. This means that $|S| \geqslant |I|$. Since S is a minimum-sized vertex cover, each edge of any matching, and hence a maximum matching of G must have at least one endpoint in S. Suppose some edge of a maximum matching has two endpoints in S, then $|S| > \mathfrak{mm}(G)$, which contradicts the assumption that G is a König-Egerváry graph. Therefore, every edge of a maximum matching of G has exactly one endpoint in S. Thus, $|S| \leqslant |V(G) \setminus S|$ i.e., $|S| \leqslant |I|$. Hence we obtain $|S| = |I| = {}^{V(G)}/_{2}$, and M is a perfect matching of G.

Now, consider the bipartite graph H with $V(H) = V(G)$, and $E(H) = E(G) \setminus E(G[S])$. That is, H has the same vertex set as that of G, and the edge set of H consists of edges of G going across from S to I. Now $V(H) = S \cup I$, such that S and I are independent sets in H, and M is a perfect matching of H. Suppose there exists $T \subsetneq I$, such that $|N(T)| = |T|$ in H. Since I is an independent set in G as well, we have $|N(T)| = |T|$ in G. Since $T \subsetneq I$, T is an independent set of G which is not maximal. Thus, by Corollary 23, we get that G is not Spartan, which is a contradiction. Therefore, H must be elementary, by Lemma 2.

Thus, we know that H can have only the two minimum-sized vertex covers S and I (using an alternate definition of elementary bipartite graph from [5]). Since $E(H) \subseteq E(G)$, any vertex cover of G must be a vertex cover of H. Thus, G cannot have any minimum-sized vertex cover other than S and I. If there is an edge of G with both its endpoints in S, then I is not a vertex cover of G and thus S is the only minimum-sized vertex cover of G. For $v \in I$, there does not exist a minimum-sized vertex cover of G which contains v. Thus, using Lemma 3, we get that G is not Spartan, which is a contradiction. Thus, there cannot be any edge with both the endpoints in S. Thus, $E(G) = E(H)$ and we already know that $V(G) = V(H)$ and H is bipartite and elementary. Therefore, G is bipartite and elementary. ◀

---

[1] There are König-Egerváry graphs which are non-bipartite as well. For example consider G with $V(G) = \{a_1, b_1, a_2, b_2\}$ and $E(G) = \{a_1a_2, a_1b_1, a_2b_2, a_1b_2, a_2b_1\}$. This graph has a perfect matching $\{a_1b_1, a_2b_2\}$ and thus $\mathfrak{mm}(G) = 2$ and thus $\mathsf{mvc}(G) \geqslant 2$. But $\{a_1, a_2\}$ forms a vertex cover of G of size 2 thus $\mathsf{mvc}(G) = \mathfrak{mm}(G)$ and G is not bipartite because it contains a triangle.

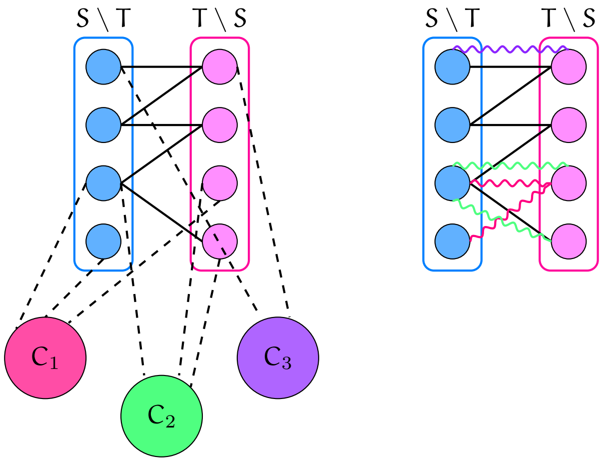

Figure 3 Depicting the construction of $\mathfrak{h}_G(S,T)$ where S and T are two minimum sized vertex covers of G. The connected components of $S \cap T$ are C_1 (red), C_2 (green) and C_3 (purple). The edges is G from $S \setminus T$ to $T \setminus S$ are also there in $\mathfrak{h}_G(S,T)$ as "real edges". For every u in $S \setminus T$ and v in $T \setminus S$ such that u and v are both adjacent to a vertex in C_1, there is a red wavy "helper" edge in $\mathfrak{h}_G(S,T)$, and similarly there are green and purple helper edges corresponding to C_2 and C_3.

5 General Graphs

Before stating the main result of this section, we need the definition of an auxiliary graph, defined for a pair of vertex covers of a graph G. Also, in this section, as we are dealing with Spartan-ness, we are dealing with configurations of guards of size $\mathsf{mvc}(G)$. If there is more than one guard on some vertex of a configuration, the set of vertices occupied by guards does not form a vertex cover, and this configuration can be attacked by the attacker immediately. Hence, we look at only those configurations with one guard per vertex.

▶ **Definition 24.** *Let S and T be two vertex covers of a graph G. Let $H_1, \ldots, H_k$ be the connected components of $G[S \cap T]$. For i, where $1 \leqslant i \leqslant k$, we say that two vertices $u \in S \setminus T$ and $v \in T \setminus S$ are* pseudo-adjacent via i *if both u and v are adjacent to some vertex in $V(H_i)$.*

Let $D = (S \setminus T) \cup (T \setminus S)$. We define two subsets of $D \times D$:

$$E_1 := \{(u,v) \mid u \in S \setminus T, v \in T \setminus S, uv \in E(G)\},$$

and

$$E_2 := \{(u,v) \mid u \in S \setminus T, v \in T \setminus S, uv \text{ are pseudo-adjacent via } i \text{ for some } i\}.$$

We use $\mathfrak{h}_G(S,T)$ to denote the graph $(D, E_1 \uplus E_2)$.

Note that $\mathfrak{h}_G(S,T)$ may have multi-edges. In the context of this graph, we refer to the edges in E_1 as *real* edges and the edges in E_2 as *helper* edges.

Note that the helper edges can be naturally partitioned into k sets $E_2^{(1)}, \ldots, E_2^{(k)}$, where $E_2^{(i)}$ consists of the helper edges that are pseudo-adjacent via i. For convenience, we refer to the edges in $E_2^{(i)}$ as being edges of color i.

The following can be easily observed:

▶ **Lemma 25.** *The graph $\mathfrak{h}_G(S,T)$ is bipartite for a graph G with at least one edge.*

Proof. Since $S \setminus T$ and $T \setminus S$ are subsets of independent sets $V(G) \setminus T$ and $V(G) \setminus S$ respectively, the graph $\mathfrak{h}_G(S,T)$ is bipartite. ◀

▶ **Theorem 26.** *A graph* G *is Spartan if and only if there exists a non-empty family* $\mathcal{F}$ *of minimum-sized vertex covers such that the following condition holds:*

For every $S \in \mathcal{F}$ *and for every edge* $uv \in G$ *for which* $u \in S$ *and* $v \notin S$*, there exists* $T \in \mathcal{F}$ *such that* $v \in T$ *and:*

(1) *either* $u \notin T$ *and there is a perfect matching in* $\mathfrak{h}_G(S, T)$ *which contains the edge* uv*,*

(2) *or* $u \in T$*, and there is a perfect matching in* $\mathfrak{h}_G(S, T)$*, where the matched partner of* v*, say* w*, has a neighbor – in* G *– among the vertices of* X*, where* X *is the connected component of* $G[S \cap T]$ *that contains* u*.*

Proof. First, assume that G is Spartan. Then, there is a strategy for indefinite defense of G with $\mathsf{mvc}(G)$ guards. Let $\mathcal{F}$ be the set of all vertex covers that are used in the strategy: note that they are all of minimum size by definition. Since G is Spartan, $\mathcal{F}$ is indeed non-empty.

Consider a vertex cover $S \in \mathcal{F}$ and an edge uv such that $u \in S$ and $v \notin S$.

As S occurs in the strategy of the defender, there is a way to defend an attack on an edge when guards occupy the vertices of S (one guard on each vertex). We attack the edge uv here and observe the promised defense: let us say that the guards are positioned on the vertex cover T after the defense is executed. Clearly, $T \in \mathcal{F}$, by definition. Depending on whether $u \in T$ or not, we can conclude that either (1) holds or (2) does (respectively), by tracing the movement of the guards from vertices in $S \setminus T$ to vertices in $T \setminus S$.

Suppose $u \notin T$, we show that (1) holds. Let $S' = S \setminus T$, i.e., the set of vertices which had a guard before the attack and do not have a guard after the defense is completed. Let $T' = T \setminus S$, i.e., the set of vertices which had a guard after the defense is completed and do not have a guard before the attack. Let $|S'| = |T'| = k$. Then there must be a collection $\mathcal{P}$ of k vertex disjoint paths from S' to T' (with the starting vertex of each path from S', end vertex in T', and each intermediate vertex from $S \cap T$), and one of these paths must be the edge uv. We show how the collection $\mathcal{P}$ is obtained. Each path is obtained by tracing the movement of each individual guard. The guard on u is forced to move to v by the attacker and hence the edge uv is traced by a guard. Thus $uv \in \mathcal{P}$. Similarly look at the movement of a guard on a vertex say $u_1 \neq u$ of S'. This guard moves to some vertex u_2. If $u_2 \in T$, $u_1u_2 \in \mathcal{P}$. Otherwise $u_2 \in S \cap T$ as the guard can only move for one step after each attack and hence u_2 must have a guard both before and after the attack. The guard which previously was on u_2 moves to some u_3 (as both S and T are vertex covers hence there must be only one guard per vertex after the reconfiguration has been done). Now again if $u_3 \in T$, $u_1u_2u_3 \in \mathcal{P}$. Otherwise $u_3 \in S \cap T$ and the guard which was previously on u_3 must move to some u_4. This process will only stop when a guard moves to a vertex which did not already have a guard, i.e., a vertex in T'. We obtain k paths by tracing the movement of each of the k guards in S' (including the guard moving from u to v). It is clear that a vertex v in S' or T' cannot belong to two paths because this will indicate that two guards started from the vertex v (if $v \in S'$) or two guards ended up at v (if $v \in T'$). A vertex $v \in S \cap T$ cannot belong to two paths in $\mathcal{P}$, because this will indicate that this vertex has at least two guards after the reconfiguration is done which is not possible. Also, each path in $\mathcal{P}$ which contains more than one edge must have all the intermediate vertices from the same connected component of $S \cap T$. This is because the intermediate vertices contain a guard both before and after the reconfiguration and hence lie in $S \cap T$. Each guard can move only to a neighboring vertex, the intermediate vertices of a path in $\mathcal{P}$ also form a path in $S \cap T$. Thus a path $u_1u_2 \ldots u_t \in \mathcal{P}$ where $t > 2$ implies $u_1 \in S'$, $u_t \in T'$ and $u_2, u_3, \ldots, u_{t-1} \in V(H_i)$, for some i. Therefore, in the graph $\mathfrak{h}_G(S, T)$, there exists a helper edge of color i between u_1 and u_t. Also an edge in $\mathcal{P}$ corresponds to a real edge in $\mathfrak{h}_G(S, T)$. Thus, each path in $\mathcal{P}$ gives an edge in $\mathfrak{h}_G(S, T)$. The edges in $\mathfrak{h}_G(S, T)$ corresponding to the paths in $\mathcal{P}$ form a perfect matching, because each vertex in S' is adjacent to exactly one edge corresponding to a path in $\mathcal{P}$. This is because there is only one guard on each vertex of S' before the attack and each vertex in S' has no guard after the attack. Thus, $\mathfrak{h}_G(S, T)$ contains a perfect matching containing the edge uv and condition (1) holds.

Now suppose $u \in T$, we show that (2) holds. Let S' and T' be the same as above, and similarly $k = |S'| = |T'|$. By the same reasoning as the previous case, we have k vertex disjoint paths in G, which correspond to a perfect matching M in $\mathfrak{h}_G(S, T)$. Now as $u \in S$ and $u \in T$, a guard is present on u before the attack and after the reconfiguration is complete. Since T is obtained from S by applying the winning strategy for a defense after the attacker attacks the edge uv, the guard on u must move to v while reconfiguring from S to T. Since a guard is present on u after the attack, there must exist a path $u_1u_2 \dots u_t$ in $\mathcal{P}$, where $t > 2$, such that $u_1 \in S'$, $u_{t-1} = u$ and $u_t = v$. Also, $u_2, u_3, \dots, u_{t-1}$ belong to the same connected component of H_i. This path corresponds to the following movement: A guard on a vertex u_1 of S' moves to a vertex u_2 of $S \cap T$. The guard on u_2 moves to u_3 of $S \cap T$, and so on, till the guard on u_{t-2} moves to u, and the guard on u moves to v. Clearly, the vertices $u_2, u_3, \dots, u_{t-1}$ must belong to the same connected component of $S \cap T$, because a guard can only move to a neighboring vertex. The edge u_1v is a helper edge in $\mathfrak{h}_G(S, T)$ and lies in M. The vertex u_1 has a neighbor u_2 in the connected component of $G[S \cap T]$ that contains u (because there is a path $u_2u_3 \dots u_{t-1} = u$ in $G[S \cap T]$). Hence, (2) holds.

Thus, we have shown that when G is Spartan, the family $\mathcal{F}$ of all vertex covers used in a strategy of the defender; satisfies the condition in Theorem 26.

Now we show the converse, i.e., if a graph G has a family of minimum-sized vertex covers $\mathcal{F}$ which satisfy the condition in Theorem 26, then G is Spartan. For this, we show that if the guards occupy a vertex cover $S \in \mathcal{F}$ and an arbitrary edge uv is attacked, the guards can reconfigure themselves, such that at least one guard moves across the attacked edge uv, and the final positions of the guards form a vertex cover $T \in \mathcal{F}$.

Since S is a vertex cover, there cannot be any edge with both the endpoints outside S. If an edge uv such that $u, v \in S$ is attacked, the guard on u moves to v, and the guard on v moves to u. The attack is defended and the configuration S is restored. Therefore, we need to consider the attack on an edge uv such that $u \in S$ and $v \notin S$.

By the given condition, there exists a vertex cover $T \in \mathcal{F}$ such that and $v \in T$ such that:

(1) either $u \notin T$ and there is a perfect matching in $\mathfrak{h}_G(S, T)$ which contains the edge uv,

(2) or $u \in T$, and there is a perfect matching in $\mathfrak{h}_G(S, T)$, where the matched partner of v, say w, has a neighbor – in G – among the vertices of X, where X is the connected component of $G[S \cap T]$ that contains u.

Suppose (1) holds. We show that it is possible to reconfigure the guards from S to T such that one guard moves across uv. Let $S' = S \setminus T$, $T' = T \setminus S$ and $|S'| = |T'| = k$. It is enough to show that there is a collection $\mathcal{P}$ containing uv of k vertex disjoint paths from S' to T' with all the intermediate vertices in $S \cap T$. A path $u_1u_2 \dots u_t$ in G where $u_1 \in S'$, $u_t \in T'$ and $u_2, u_3, \dots, u_{t-1} \in S \cap T$ will represent the movement of a guard from u_1 to u_2, the movement of the guard previously on u_2 to u_3, and so on up to the movement of the guard previously on u_{t-1} to u_t.

If there exists a perfect matching M in $\mathfrak{h}_G(S, T)$ such that it consists of possibly some real edges and at most one helper edge of each color, then there exists a collection $\mathcal{P}$ of k vertex disjoint paths from S' to T' with intermediate vertices in each path from $S \cap T$. We show that a path in $\mathcal{P}$ can be obtained from each edge of M. A real edge in $\mathfrak{h}_G(S, T)$ is also an edge in G. Thus, for each real edge e in M, add e to $\mathcal{P}$. A helper edge uv of color i implies the existence of a path $u_1(=u)u_2u_3\dots u_t(=v)$ in G where $t>1$ and $u_2,u_3,\dots,u_{t-1}\in V(H_i)$, where H_i is a connected component of $S\cap T$. For each helper edge $e=wz$ in M of color i, add a path $u_1(=w)u_2u_3\dots u_t(=z)$ in G where $t>1$ and $u_2,u_3,\dots,u_{t-1}\in V(H_i)$ to $\mathcal{P}$. Clearly, there are k paths in $\mathcal{P}$ (one path obtained from each edge of M). We show that the paths in $\mathcal{P}$ are vertex disjoint. Since M is a matching, the endpoints of each edge in M are distinct. Hence the endpoints of all the k paths are distinct. Since all the intermediate vertices of each path are distinct and each path of length at least 2 is obtained from a helper edge of different color (as there are no two helper edges of the same color in M), each path obtained from a helper edge in M has intermediate vertices from distinct components of $G[S\cap T]$. Thus all the intermediate vertices of each path in $\mathcal{P}$ are disjoint.

Now we show that there exists a perfect matching in $\mathfrak{h}_G[S,T]$ which contains possibly some real edges and at most one helper edge of each color. By Lemma 16, there exists a perfect matching from S' to T' in G. This means that there exists a perfect matching M_p in $\mathfrak{h}_G(S,T)$ which consists of only real edges. Let M be the perfect matching in $\mathfrak{h}_G(S,T)$ such that M contains uv (exists by condition (1)), and the size of $M\cap M_p$ is as large as possible. Now, consider the graph H in $\mathfrak{h}_G(S,T)$ such that $V(H)=V(\mathfrak{h}_G(S,T))$ and $E(H)=M\cup M_p$. Since M and M_p are both perfect matchings, H is a union of edges and even cycles.

▷ Claim 27. There can be at most one cycle in H.

Proof. If there are two distinct cycles C_1 and C_2 in H, at most one of them can contain the edge uv. Therefore, without loss of generality, we can assume that $uv\notin E(C_1)$. Suppose C_1 is a cycle and $C_1=u_1v_1u_2v_2\dots u_tv_tu_1$, where $M_1:=\{u_1v_1,u_2v_2,\dots,u_tv_t\}\subset M$ and $M_2:=\{v_1u_2,v_2u_3,\dots,v_{t-1}u_t,v_tu_1\}\subset M_p$ with $t>1$. The matching $M':=M\setminus M_1\cup M_2$ has strictly more intersection with M_p than M, and also contains the edge uv, which is a contradiction. ◁

If H has no cycle, then $M=M_p$, which means that all the edges of M are real, and we are done. Now, we consider the case where M has exactly one even cycle C. If $uv\notin E(C)$, then we get a contradiction by the same argument as the claim above. Therefore, let $C=u_1(=u)v_1(=v)u_2v_2\dots u_tv_tu_1$. Here $M_1:=\{u_1v_1,u_2v_2,\dots,u_tv_t\}\subset M$ and $M_2:=\{v_1u_2,v_2u_3,\dots,v_{t-1}u_t,v_tu_1\}\subset M_p$. Also, $\{u_1,u_2,\dots,u_t\}\subset S'$ and $\{v_1,v_2,\dots,v_t\}\subset T'$. This is because $u\in S'$ and $\mathfrak{h}_G(S,T)$ is bipartite (by Lemma 25). All the edges in M which are not in M_1, are also the edges of M_p, and hence are real. Next, we show that M cannot contain two (or more) helper edges of the same color.

▷ Claim 28. M can contain at most one helper edge of each color.

Proof. Suppose M contains two helper edges of the same color (say i). These two edges must lie in M_1, because the edges of M outside M_1 are real. Let these two edges be u_pv_p and u_qv_q, where $1<p<q\leqslant t$. (As uv is a real edge, we have $p\neq 1$.) Therefore, there must exist edge u_pv_q of color i. This is because u_p and v_p are both adjacent to some vertex in $V(H_i)$, and similarly u_q and v_q are both adjacent to some vertex in $V(H_i)$. Hence, u_p and v_q are adjacent to some vertex in $V(H_i)$. Also, $M_3:=\{v_pu_{p+1},v_{p+1}u_{p+2},\dots,v_{q-1}u_q\}\subset M_p$, hence the edges in M_3 are real edges of $\mathfrak{h}_G(S,T)$. Let $M_4=\{u_pv_p,u_{p+1}v_{p+1},\dots,u_qv_q\}$ and $M_5=(M_1\setminus M_4)\cup M_3\cup\{u_pv_q\}$. Consider the perfect matching $M'=(M\setminus M_1)\cup M_5$. It can be checked that M' is a perfect matching which contains uv, and has greater intersection with M_p than M, which is a contradiction. ◁

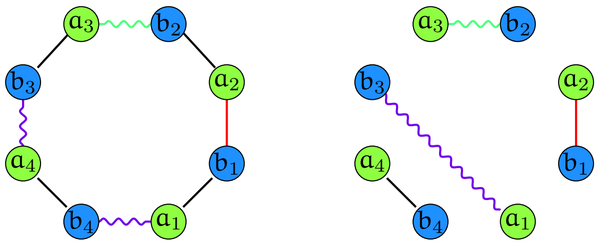

Figure 4 Demonstrating the exchange argument from the proof.

Thus, we have shown the existence of a perfect matching between S' and T' containing the edge uv in $\mathfrak{h}_G(S, T)$, which contains at most one helper edge of each color. This implies that there is a collection of $|S'|$ vertex-disjoint paths from S' to T', such that one path in this collection is the edge uv. Therefore, the guards can reconfigure from S to T (such that one guard moves across uv).

Now, suppose (2) holds. Since w has a neighbor in the same connected component $X = C_i$ (say) of $G[S \cap T]$ that contains u, this means that there exists a path $(w =)u_1u_2 \dots u_{t-1}(= u)u_t(= v)$, where $t > 1$, such that $u_2, u_3, \dots, u_{t-1}(= u)$ belong to $V(C_i)$, that is, wv is a helper edge of color i in $\mathfrak{h}_G(S, T)$. As we have seen above, a matching in $\mathfrak{h}_G(S, T)$, with at most one helper edge of each color, can be used to show the existence of $|S \setminus T|$ vertex-disjoint paths from $S' = S \setminus T$ to $T' = T \setminus S$, with the intermediate vertices from $S \cap T$. This shows that, it is possible to reconfigure the guards from S to T.

Now, if we have a perfect matching in $\mathfrak{h}_G(S, T)$, containing wv with at most one helper edge of each color, it can be seen that it is possible to reconfigure the guards from S to T such that at least one guard moves across uv. This is because, we have already seen that such a matching represents vertex-disjoint paths in G, and particularly, the edge wv represents the path $(w =)u_1u_2 \dots u_{t-1}(= u)u_t(= v)$. This means that, it is possible to move the guard on w to u_2, the guard previously on u_2 to u_3, and so on up to the guard previously on u_{t-2} to u, and the guard previously on u to v. Thus, the existence of a perfect matching in $\mathfrak{h}_G(S, T)$ containing wv with at most one helper edge of each color, is enough to show that the defender can defend an attack on the edge uv, and reconfigure the guards to a vertex cover T, while starting from a vertex cover S. It remains to show the existence of such a perfect matching in $\mathfrak{h}_G(S, T)$.

By Lemma 16, we know that there exists a perfect matching M_p between S' and T' in $\mathfrak{h}_G(S, T)$, which consists of real edges only. Let M be a perfect matching which contains a helper edge of color i adjacent to v (where $u \in V(C_i)$ for the component C_i of $G[S \cap T]$) such that $M \cap M_p$ is as large as possible. It can be seen by a similar argument to the case (1) that the graph H, with $V(H) = S' \cup T'$ and $E(H) = M \cup M_p$, can have at most one cycle (even length).

If H has no cycle, then $M = M_p$, which is not possible because M has only real edges and M_p has at least one helper edge of color i. Let the (exactly one) cycle in H be given by $C = u_1v_1u_2v_2 \dots u_tv_tu_1$, where $M_1 := \{u_1v_1, u_2v_2, \dots, u_tv_t\} \subset M$, and $M_2 := \{v_1u_2, v_2u_3, \dots, v_tu_1\} \subset M_p$. Again, if $wv \notin M_1$, then $M' = M \setminus M_1 \cup M_2$ is a perfect matching containing a helper edge wv (which is of color i and adjacent to v) and $M' \cap M_p$ has greater size than $M \cap M_p$, which is a contradiction. Therefore, $wv \in M_1$. Also, if M_1 contains any two helper edges of the same color (other than wv), then we can use an exchange argument similar to the case (1) shown above to get a contradiction.

It remains to show that C cannot have an edge of color i other than wv. Without loss of generality, let $w = u_1$ and $v = v_1$. Let $u_p v_p$ for some $p > 1$ be the other helper edge of color i. Also, recall that similarly to case (1), $\{u_1, u_2, \ldots, u_t\} \subset S'$ and $\{v_1, v_2, \ldots, v_t\} \subset T'$. This is because $u \in S'$, and $\mathfrak{h}_G(S,T)$ is bipartite (by Lemma 25). This means that $u_p \in S'$, and we already know that $v \in T'$. Since $u_p v_p$ is an edge of color i, u_p is adjacent to $V(C_i)$ in G. We already know that v is adjacent to $V(C_i)$ in G. Therefore, there exists a helper edge $u_p v$ of color i in $\mathfrak{h}_G(S,T)$. Let $M_3 := \{v_2u_1, v_3u_2, \ldots, v_pu_{p-1}, u_pv\}$, all these edges exist in $\mathfrak{h}_G(S,T)$ because the edges in M_3 other than u_pv are from M_p (as seen in the above paragraph) and hence are, in fact, real edges and we have already seen the existence of the edge u_pv which is a helper edge of color i. Let $M_4 := \{u_1v_1, u_2v_2, \ldots, u_pv_p\} \subset M_1$. Let $M_5 := (M_1 \setminus M_4) \cup M_3$. The perfect matching $M' = (M \setminus M_1) \cup M_5$ contains a helper edge u_pv of color i (adjacent to v) and has a greater intersection (in size) with M_p compared to M, which is a contradiction. Therefore, the guards on S can be reconfigured to T while defending the attack on the edge uv. Thus, we have shown that it is always possible for the guards to reconfigure between the vertex covers in $\mathcal{F}$ while defending every attack; thus the graph G is Spartan. ◀

6 Concluding Remarks

In this paper, we give a necessary and sufficient condition for a graph G to be Spartan, i.e., to satisfy $\mathsf{evc}(G) = \mathsf{mvc}(G)$. There are several directions to be pursued further. An important question is whether the complexity of checking whether a given graph G is Spartan is less than that of computing $\mathsf{evc}(G)$. In terms of checking whether $\mathsf{evc}(G) = \mathsf{mvc}(G)$, although we have a complete characterization, the question of finding a simpler characterization remains open. In particular, it would be interesting to know whether there exists a graph G, such that every vertex of G belongs to a strongly (weakly) good vertex cover, but $\mathsf{evc}(G) \neq \mathsf{mvc}(G)$. If not, it would be good to know the proof that this condition is indeed sufficient to guarantee $\mathsf{evc}(G) = \mathsf{mvc}(G)$. Also, we know that there exist weakly good vertex covers that are not strongly good and hence can be destroyed by the attacker (hence, not evc configurations). However, we still do not know whether the condition: "For each vertex v there exists a weakly good minimum vertex cover containing v" implies the following condition: "For each vertex v there exists a strongly good minimum vertex cover containing v".

References

1. Hisashi Araki, Toshihiro Fujito, and Shota Inoue. On the eternal vertex cover numbers of generalized trees. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, 98-A(6):1153–1160, 2015. doi:10.1587/TRANSFUN.E98.A.1153.
2. Jasine Babu, L. Sunil Chandran, Mathew C. Francis, Veena Prabhakaran, Deepak Rajendraprasad, and Nandini J Warrier. On graphs whose eternal vertex cover number and vertex cover number coincide. *Discrete Applied Mathematics (in press)*, 2021.
3. Reinhard Diestel. *Graph theory, Sixth Edition*. Springer, 2025.
4. William F. Klostermeyer and Christina M. Mynhardt. Edge protection in graphs. *Australas. J Comb.*, 45:235–250, 2009. URL: http://ajc.maths.uq.edu.au/pdf/45/ajc_v45_p235.pdf.
5. László Lovász and Michael D Plummer. *Matching theory*, volume 367. American Mathematical Soc., 2009.
6. Neeldhara Misra and Saraswati Girish Nanoti. Spartan bipartite graphs are essentially elementary. In *MFCS*, volume 272 of *LIPIcs*, pages 68:1–68:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.MFCS.2023.68.

A Appendix

Proof of Claim 6 in Lemma 5. Let T be a minimum-sized vertex cover of G, such that T contains more than $|\mathsf{N}(\mathsf{X})|$ vertices from $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$. Consider the set $\mathsf{T}' = \big(\mathsf{T} \setminus (\mathsf{X} \cup \mathsf{N}(\mathsf{X}))\big) \cup \mathsf{N}(\mathsf{X})$. Clearly, the size of T' is less than the size of T. We show that T' is also a vertex cover of G, which contradicts the fact that T is a minimum-sized vertex cover of G. Any edge with both endpoints outside $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$ is covered by T, and hence covered by T', because the vertices outside $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$ which belong to T also belong to T'. Any edge with both endpoints in $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$ is also covered by T', as $\mathsf{N}(\mathsf{X}) \subset \mathsf{T}'$ and no edge in G can have both endpoints in X, as X is an independent set. An edge with one endpoint in $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$ and the other endpoint outside $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$, must have an endpoint in $\mathsf{N}(\mathsf{X})$ because a vertex in X cannot be adjacent to a vertex outside $\mathsf{N}(\mathsf{X})$ (by the definition of $\mathsf{N}(\mathsf{X})$). Thus, such an edge is also covered by T', as $\mathsf{N}(\mathsf{X}) \subseteq \mathsf{T}'$. Thus ,every edge of G has at least one endpoint in T'. This is a contradiction, as G cannot have a vertex cover of size smaller than T. Thus, any minimum-sized vertex cover of G cannot contain more than $\mathsf{N}(\mathsf{X})$ vertices from $\mathsf{X} \cup \mathsf{N}(\mathsf{X})$. ◁

Proof of Corollary 7. Let G be a Spartan graph with more than one vertex, and $|\mathsf{V}(\mathsf{G})| = n$. Let I be a maximum independent set of G. Let $\mathsf{S} = \mathsf{V}(\mathsf{G}) \setminus I$. Then, $|\mathsf{N}(I)| \geqslant |I|$ by Lemma 5, and thus $|I| \leqslant |\mathsf{V}(\mathsf{G})|/2$, which means that $|\mathsf{S}| \geqslant |\mathsf{V}(\mathsf{G})|/2$, i.e., $\mathsf{mvc}(\mathsf{G}) \geqslant |\mathsf{V}(\mathsf{G})|/2$. ◀

Proof of Corollary 23. Let G be a graph with more than one vertex which is Spartan. Let I be an independent set of G which is not maximal, such that $|\mathsf{N}(I)| = |I|$. Denote $\mathsf{N}(I)$ by C. Consider an inclusion-wise minimal such set I. The graph H with vertex set $I \cup \mathsf{C}$ and edge set $\mathsf{E}(\mathsf{G}[I \cup \mathsf{C}]) \setminus \mathsf{E}(\mathsf{C})$ must be bipartite (because I is an independent set) and elementary (using the inclusion-wise minimality of I and Lemma 2). Therefore, using an alternate definition of elementary bipartite graph from [5], I and C are the only two minimum-sized vertex covers of H. Now, any minimum vertex cover S of G cannot contain more than $|I|$ many vertices from $\mathsf{G}[I \cup \mathsf{C}]$. Suppose not, $\mathsf{S} \setminus (I \cup \mathsf{C}) \cup \mathsf{C}$ is a strictly smaller vertex cover of G (this is a vertex cover, as there are no edges from I to $\mathsf{V}(\mathsf{G}) \setminus \mathsf{C}$, and hence $\mathsf{S} \setminus \mathsf{C}$).

Now, since G is Spartan, by Lemma 12, every vertex in I belongs to a minimum-sized weakly good vertex cover. Consider a minimum-sized weakly good vertex cover T, which contains some $v \in I$. Clearly, T can contain only $|I|$ vertices from $\mathsf{C} \cup I$. Also as T is also a vertex cover of H of size $|I|$ and H has only two minimum vertex covers, $I \subset \mathsf{T}$ and $\mathsf{C} \cap \mathsf{T} = \emptyset$. The graph $\mathsf{G}[\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C})]$ is non-empty as I is not maximal. If $\mathsf{T} \cap (\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C}))$ is more than $mvc(\mathsf{G}[\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C})])$, then a smaller vertex cover of G than T can be obtained from the union of C and a minimum vertex cover of $(\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C}))$. Thus $\mathsf{T} \cap (\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C}))$ is equal to the $mvc(\mathsf{G}[\mathsf{V}(\mathsf{G}) \setminus (I \cup \mathsf{C})])$, which makes C a weakly bad subset for T which is a contradiction. ◀