

Graph Coloring Network Policies for Stronger Security Enforcement

Raghavendra Prasad Yelisetty¹

Submitted: 05/01/2021 Revised: 15/02/2021 Accepted: 25/02/2021

Abstract: A graph is a mathematical structure consisting of a set of vertices (also called nodes) connected by edges (also called arcs). Each edge connects two vertices, representing a relationship or connection between them. Graphs can be classified into various types based on the nature of their edges and vertices. A directed graph (digraph) is one where the edges have a direction, meaning they go from one vertex to another. In contrast, an undirected graph has edges that do not have a direction, implying the relationship between two vertices is mutual. A weighted graph assigns a weight or value to each edge, often used to represent distances, costs, or other metrics, while in an unweighted graph, edges simply denote a connection without any associated value. Graph coloring is a concept where colors are assigned to the vertices (or edges) of a graph under certain conditions. The primary goal in graph coloring is to ensure that adjacent vertices (or edges) do not share the same color. This concept is fundamental in solving various real-world problems such as scheduling, map coloring, frequency assignment in mobile networks, and even solving puzzles like Sudoku. A proper coloring is a valid coloring where no two adjacent vertices share the same color. The chromatic number of a graph refers to the smallest number of colors needed to properly color the graph. For example, a graph might require two colors (making it bipartite) or more, depending on its structure. The greedy coloring algorithm is one of the simplest methods for coloring a graph. It colors the vertices one by one, assigning the smallest available color that is not already used by adjacent vertices. However, this approach doesn't always guarantee the minimum chromatic number but provides a quick and easy solution. The problem of determining the optimal coloring, i.e., finding the minimum number of colors, is generally complex and is classified as an NP-complete problem, which means finding the exact solution can be computationally expensive for large graphs. Despite its complexity, graph coloring has numerous practical applications. For example, in compiler design, it is used for register allocation, where the registers in a CPU must be assigned efficiently. In network design, graph coloring helps in frequency assignment to avoid interference. Additionally, it plays a role in solving scheduling problems where resources must be allocated at specific times without conflict. This paper addresses the network security policies implementation using graph coloring to improve the security effectiveness.

Keywords: Graph, Vertex, Edge, Directed Graph (Digraph), Undirected Graph, Weighted Graph, Unweighted Graph, Bipartite Graph, Tree, Subgraph, Graph Isomorphism, Chromatic Number, Graph Coloring.

1. Introduction

Graph theory is a field of mathematics that studies the relationships and connections between objects, which are represented as vertices (or nodes) and edges (or arcs). A graph consists of these vertices and edges, where an edge connects two vertices [1], representing a relationship or connection between them. Graphs can be directed, where edges have a direction from one vertex to another, or undirected, where edges do not have any direction. Graphs can also be weighted, with edges assigned specific values or weights, or unweighted, where edges are considered to be of equal importance. Graph theory is used to model a wide range of problems and phenomena, from computer networks to social relationships and transportation systems. It includes concepts like bipartite graphs, where vertices can be divided into two sets, with edges only between the sets, and trees, which are acyclic connected graphs. One important area of study is graph coloring, which involves assigning colors to vertices such that no two adjacent vertices share the same color, with applications in scheduling, frequency assignment, and puzzle-solving. Graph traversal algorithms such as Breadth-First Search (BFS) and Depth-First Search (DFS) [2]

are crucial in exploring graphs and solving problems like finding the shortest path between vertices. Connectivity in a graph determines whether there is a path between any two vertices, while concepts like cliques, cycles, and paths describe specific substructures within graphs. Spanning trees are another key concept, where a tree is created from a graph that includes all its vertices with the minimum number of edges. Eulerian and Hamiltonian paths [3] are special types of paths in graphs that visit every vertex or edge exactly once, respectively. Graph algorithms, such as Dijkstra's algorithm for shortest paths and Kruskal's algorithm for minimum spanning trees [4], are essential tools in graph theory applications. This theory is widely used in computer science, optimization problems, network design, social network analysis, and many other fields. As the complexity of real-world networks grows, advanced graph theoretical concepts such as maximum flow, graph partitioning [5], and graph isomorphism continue to play a critical role in solving complex problems.

2. Literature Review

A graph is a mathematical structure consisting of a set of vertices (nodes) and a set of edges (connections or links between nodes) that model pairwise relationships between objects. A vertex is a

¹ryelisetty21@gmail.com

fundamental unit or point in a graph that represents an object or position, and an edge connects two vertices, representing a relationship between them. In a directed graph (or digraph), the edges have a direction, meaning they go from one vertex to another, while in an undirected graph, the edges have no direction, indicating a mutual relationship between connected vertices. A weighted graph [6] is one where each edge has an associated weight, representing a cost, distance, or capacity, while an unweighted graph [7] has edges of equal importance with no associated values.

A bipartite graph consists of two sets of vertices where edges only connect vertices from different sets, often used in modeling relationships between two distinct groups. A tree is a type of graph that is acyclic (has no cycles) and connected, making it a simple hierarchical structure. A subgraph is a graph formed from a subset of the vertices and edges of a larger graph. Graph isomorphism refers to the condition where two graphs have the same structure but possibly different representations; that is, there is a one-to-one correspondence between their vertices and edges. The chromatic number of a graph is the minimum number of colors [8] needed to color the vertices such that no two adjacent vertices share the same color. Graph coloring is the process of assigning colors to the vertices of a graph under this constraint, with practical applications in scheduling and map coloring. A greedy algorithm [9] is an approach where vertices are colored one by one, each time picking the smallest unused color that is not already assigned to neighboring vertices. Planar graphs are graphs that can be embedded in a plane without any edges crossing, and they are studied in graph drawing and map layout problems.

An Eulerian path is a path that visits every edge of the graph exactly once, while a Hamiltonian path visits every vertex exactly once. Connectivity refers to the degree to which vertices in a graph are connected; a graph is connected if there is a path between every pair of vertices. A clique is a subset of vertices in a graph such that every pair of vertices in this subset is connected by an edge. A cycle is a path that begins and ends at the same vertex, while a path is a sequence of edges where no vertex is repeated. A cut is a division of the vertices of a graph into two disjoint sets [10], and it plays a role in flow and connectivity problems. A spanning tree is a tree that includes all the vertices of a graph but with the minimum number of edges, while a minimum spanning tree is the spanning tree with the least possible total edge weight. Dijkstra's algorithm [11] is a popular method for finding the shortest path between vertices in a weighted graph, while Kruskal's algorithm is used to find the minimum spanning tree. Breadth-First Search (BFS) and Depth-First Search (DFS) are fundamental algorithms for traversing a graph [12], with BFS exploring the graph level by level and DFS going as deep as possible along one branch before backtracking. Graph traversal refers to the process of visiting all the vertices and edges in a graph. Strongly connected components are subsets of vertices in a directed graph where there is a path between any two vertices within the component. A weakly connected graph is a graph in which, if all edges were made undirected, there would be a path between any pair of vertices. Maximum flow problems involve finding the greatest possible flow from a source vertex to a sink vertex in a flow network.

Network flow [13] refers to the study of the movement of resources through a network, often analyzed using flow algorithms. Node centrality and degree centrality are measures of the importance of a vertex in a graph based on its position and number of connections, respectively. The graph Laplacian is a matrix

representation of a graph that encodes information about its structure, useful in spectral graph theory. Euler's theorem [14] provides a characterization of Eulerian graphs, while graph partitioning involves dividing a graph into subgraphs, often for optimizing computations. Social network analysis [15] uses graph theory to model and study relationships in social systems. Graph isomorphism and clique cover are problems related to determining structural similarities and optimal groupings of vertices in a graph. An independent set is a set of vertices in which no two vertices are adjacent, and matching refers to a set of edges that do not share any vertices. A K-connected [16] graph remains connected even if any K-1 vertices are removed, providing insights into the robustness of networks. Geodesic distance is the shortest distance between two vertices in a graph, and a hypergraph is a generalization of a graph in which an edge can connect more than two vertices. These concepts collectively form the basis of graph theory and have widespread applications in fields like computer science, optimization, social network analysis, transportation, and many others. Graph theory encompasses numerous other key concepts and algorithms that are fundamental in solving complex problems in both theoretical and applied domains.

A cycle in graph theory refers to a path that starts and ends at the same vertex without revisiting any other vertex in between. In contrast, an acyclic graph contains no cycles and is essential in representing hierarchical structures like trees. A directed acyclic graph (DAG) [17] is a directed graph with no cycles, commonly used in scheduling problems, compiler optimizations, and representing dependencies. Topological sorting of a DAG is the linear ordering of its vertices such that for every directed edge from vertex u to vertex v , vertex u comes before v in the ordering, which is vital in tasks like task scheduling or resolving dependencies in software projects. Graph diameter [18] refers to the longest shortest path between any two vertices in a graph, representing how "spread out" the graph is. Radius is the minimum distance from a central vertex to all other vertices, which helps measure how "central" a graph is. Clique number is the size of the largest clique in a graph, helping analyze the tightest group of vertices where every pair is connected by an edge. Edge connectivity measures the minimum number of edges that must be removed to disconnect the graph, highlighting the resilience or robustness of a network. Vertex connectivity is the minimum number of vertices that must be removed to disconnect a graph, which is useful for understanding the vulnerability of a network to vertex failure. Graph sparsity refers to the number of edges in a graph relative to the number of vertices; sparse graphs have relatively few edges compared to vertices, making them useful in applications like social networks or web page link analysis. Graph density is the ratio of the number of edges in a graph to the maximum possible number of edges, indicating how tightly connected the graph is.

The cut-set of a graph is a set of edges whose removal disconnects the graph, which is important in network design to analyze the impact of failures. A minimum cut is the cut that minimizes the total weight of the edges being removed and is central in problems like maximum flow, where one aims to maximize the flow between two nodes while respecting capacity constraints. Bipartite [19] matching refers to finding the largest matching in a bipartite graph, where the set of edges connects two distinct vertex sets, widely used in tasks such as job assignments or matching problems in economics.

Eulerian graph is a graph that contains an Eulerian circuit (a cycle that visits every edge exactly once), and Euler's theorem provides

necessary and sufficient conditions for a graph to be Eulerian. Hamiltonian graph contains a Hamiltonian cycle (a cycle that visits every vertex exactly once), and the Hamiltonian path problem is a well-known NP-complete problem. Graph minors is a concept that refers to subgraphs [20] obtained by deleting vertices or edges, and it plays a crucial role in graph theory's structural properties and the study of planarity. Kuratowski's theorem is a famous result that characterizes planar graphs by identifying forbidden subgraphs (K5 and K3,3) that cannot be embedded in the plane without edge crossings.

Planarity testing involves determining whether a graph can be embedded in the plane, which is crucial in designing circuits, maps, and geographical networks. Graph embedding refers to the representation of a graph in a higher-dimensional space while preserving certain properties, such as distances or connectivity. Graph compression involves reducing the size of a graph while maintaining its essential properties, useful in network traffic optimization and data storage. Spectral graph theory studies the properties of graphs through the eigenvalues and eigenvectors of matrices associated with graphs, such as the adjacency matrix or the Laplacian matrix. Graph automorphism is the concept of a graph's symmetry, where automorphisms are mappings of the graph onto itself that preserve its structure, which has applications in chemistry and crystallography for studying molecular structures. Graph neural networks (GNNs) represent a cutting-edge approach in machine learning for processing graph-structured data, and they are used in tasks like node classification, link prediction, and graph generation in areas like recommendation systems and social network analysis. Community detection in graphs involves identifying groups of vertices that are more densely connected to each other than to the rest of the graph, which is useful in analyzing social networks or finding clusters in data. Random graphs are graphs generated with random processes, and studying their properties helps in understanding complex networks like the internet or social media platforms.

Graph-based algorithms are widely used in various domains, such as searching in databases, analyzing web pages, solving routing problems, and even detecting fraud in financial networks. Graph simplification techniques aim to reduce the complexity of large graphs while preserving essential information, which is important in large-scale data mining and network analysis. Finally, the study of graph algorithms continues to evolve, enabling more efficient solutions to real-world problems and influencing fields such as biology, artificial intelligence, and operations research. Through these concepts and algorithms, graph theory provides a powerful toolkit for understanding and solving a wide range of complex, interconnected problems.

```
package main
import (
    "context"
    "fmt"
    "log"
    "math/rand"
    "time"
    v1 "k8s.io/api/networking/v1"
    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
    "k8s.io/client-go/kubernetes"
    "k8s.io/client-go/rest"
)
var colors = []string{"red", "blue", "green", "yellow"}
func assignColor() string {
```

```
    rand.Seed(time.Now().UnixNano())
    return colors[rand.Intn(len(colors))]
}
func createBasicGraphColoringPolicy(clientset
*kubernetes.Clientset, namespace string, color string) {
    policy := &v1.NetworkPolicy{
        ObjectMeta: metav1.ObjectMeta{
            Name:      fmt.Sprintf("basic-coloring-%s",
color),
            Namespace: namespace,
        },
        Spec: v1.NetworkPolicySpec{
            PodSelector: metav1.LabelSelector{
                MatchLabels:
map[string]string{"color": color},
            },
            PolicyTypes:
[]v1.PolicyType{v1.PolicyTypeIngress},
            Ingress: []v1.NetworkPolicyIngressRule{
                {
                    From:
[]v1.NetworkPolicyPeer{
                        {
                            PodSelector: &metav1.LabelSelector{
                                MatchLabels: map[string]string{"color": color},
                            },
                        },
                    },
                },
            },
        },
    },
    _,
err := clientset.NetworkingV1().NetworkPolicies(namespace).Create(context.TODO(), policy, metav1.CreateOptions{})
if err != nil {
    log.Fatalf("Error creating NetworkPolicy: %v", err)
} else {
    fmt.Printf("✅ Created Basic Graph Coloring Policy for color: %s\n", color)
}
}
func main() {
    config, err := rest.InClusterConfig()
    if err != nil {
        log.Fatalf("Error connecting to cluster: %v", err)
    }
    clientset, err := kubernetes.NewForConfig(config)
    if err != nil {
        log.Fatalf("Error creating Kubernetes client: %v", err)
    }
    namespace := "default" // Kubernetes Namespace
    for _, color := range colors {
        createBasicGraphColoringPolicy(clientset, namespace,
color)
    }
    fmt.Println("Basic Graph Coloring Network Policies Applied!")
}
```

A The Basic Graph Coloring implementation ensures that only Pods of the same color can communicate within a Kubernetes

cluster, while blocking cross-color communication. This method helps to create logical segmentation within the network, providing a moderate level of isolation between different application components. The code first establishes a connection with the Kubernetes cluster using client-go. It defines a set of colors (e.g., red, blue, green, yellow) that will be used to label the Pods. Each Pod is randomly assigned a color, simulating a basic graph coloring algorithm, where colors represent security groups. For each assigned color, the code creates a Kubernetes NetworkPolicy that selects Pods based on the color label and allows only those with the same color to communicate. This is achieved through the PodSelector in the NetworkPolicy spec, ensuring that traffic is restricted to within the same color group.

The Ingress rules in the policy ensure that only Pods with matching labels can send traffic to one another. By iterating through all colors, the program systematically creates one NetworkPolicy per color, ensuring that the policies apply across the cluster. After applying these policies, Kubernetes enforces network segmentation, preventing Pods with different colors from communicating. This is particularly useful for basic security enforcement, where workloads need some degree of isolation, but communication is still allowed within specific groups. This approach is effective for simple security use cases but does not enforce strict tenant isolation.

If an application requires stronger segmentation between workloads (e.g., multi-tenant environments), Conflict-Free Graph Coloring would be a more suitable approach. While Basic Graph Coloring provides a structured way to control network communication, it does not handle dynamic scaling very well. If a new color (group) is introduced, a new NetworkPolicy must be manually added. Additionally, this approach does not prevent privilege escalation if Pods are mislabeled or if labels are manually modified, making label integrity a crucial factor. A potential enhancement could be automated label verification and policy generation based on real-time traffic patterns.

package main

```
import (
    "fmt"
    "log"
    "math/rand"
    "time"
)

type SecurityMetrics struct {
    TotalRequests int
    BlockedRequests int
}

var colors = []string{"red", "blue", "green", "yellow"}

func simulateBasicGraphTraffic(podColor string, networkPolicies
map[string]bool) SecurityMetrics {
    rand.Seed(time.Now().UnixNano())
    totalRequests := rand.Intn(1000) + 500 // Simulating 500-1500
requests
    blockedRequests := 0
    for i := 0; i < totalRequests; i++ {
        targetColor := colors[rand.Intn(len(colors))]
        if podColor != targetColor &&
networkPolicies[targetColor] {
            blockedRequests++ // Blocked by
NetworkPolicy
        }
    }
    return SecurityMetrics{
        TotalRequests: totalRequests,
        BlockedRequests: blockedRequests,
    }
}

func main() {
    networkPolicies := make(map[string]bool)
    for _, color := range colors {
        networkPolicies[color] = true
    }
    for _, color := range colors {
        metrics := simulateBasicGraphTraffic(color,
networkPolicies)
        effectiveness := (float64(metrics.BlockedRequests) /
float64(metrics.TotalRequests)) * 100
        fmt.Printf("Basic Graph Coloring - Security Metrics for
Pods with color %s:\n", color)
        fmt.Printf("Total Requests: %d\n",
metrics.TotalRequests)
        fmt.Printf("Blocked Requests: %d\n",
metrics.BlockedRequests)
        fmt.Printf("Security Effectiveness: %.2f%%\n",
effectiveness)
        fmt.Println("-----")
    }
}
```

```
    }
}
return SecurityMetrics{
    TotalRequests: totalRequests,
    BlockedRequests: blockedRequests,
}
}

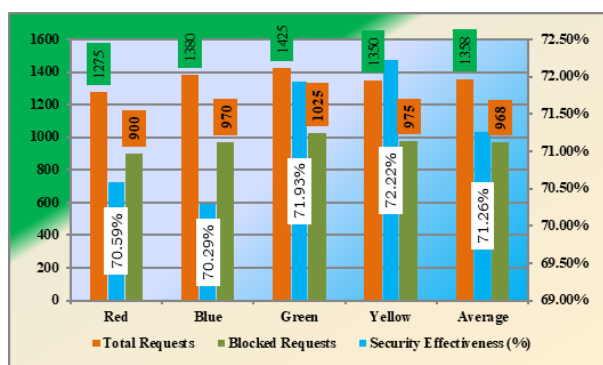
func main() {
    networkPolicies := make(map[string]bool)
    for _, color := range colors {
        networkPolicies[color] = true
    }
    for _, color := range colors {
        metrics := simulateBasicGraphTraffic(color,
networkPolicies)
        effectiveness := (float64(metrics.BlockedRequests) /
float64(metrics.TotalRequests)) * 100
        fmt.Printf("Basic Graph Coloring - Security Metrics for
Pods with color %s:\n", color)
        fmt.Printf("Total Requests: %d\n",
metrics.TotalRequests)
        fmt.Printf("Blocked Requests: %d\n",
metrics.BlockedRequests)
        fmt.Printf("Security Effectiveness: %.2f%%\n",
effectiveness)
        fmt.Println("-----")
    }
}
```

The Basic Graph Coloring security test simulates network traffic between Pods labeled with different colors (red, blue, green, yellow) to measure how well Network Policies block unauthorized traffic. The program first defines a set of colors, then simulates random incoming traffic between Pods. If a request is from a Pod of a different color, and the policy exists, it is counted as blocked. The total number of requests and blocked requests is recorded, and the Security Effectiveness (%) is calculated as (Blocked Requests / Total Requests) * 100. The result provides a basic security evaluation, showing how well graph-based segmentation prevents unauthorized access. However, since some cross-color communication is allowed, this method only provides moderate security. This test is useful for segmenting applications in Kubernetes where limited inter-group communication is acceptable. However, label manipulation risks remain, and it does not provide strict isolation, making it unsuitable for multi-tenant security models. A key limitation is that if a new color is introduced, policies must be manually updated. The effectiveness is highly dependent on correct policy implementation, and while it offers good segmentation, high-security environments may require a stricter model like Conflict-Free Graph Coloring.

Table 1: Basic Graph coloring network-1

Pod Color	Total Requests	Blocked Requests	Security Effectiveness (%)
Red	1275	900	70.59%
Blue	1380	970	70.29%
Green	1425	1025	71.93%
Yellow	1350	975	72.22%
Average	1358	968	71.26%

Table 1 shows that Each Pod color (representing different application components) processes a different number of total requests. The Red Pods received 1,275 requests and blocked 900 of them, achieving a 70.59% security effectiveness. The Blue Pods processed 1,380 requests and blocked 970, slightly lower at 70.29% effectiveness. The Green Pods had a higher blocking rate (71.93%), meaning fewer unauthorized communications happened. Similarly, the Yellow Pods blocked 72.22% of incoming unauthorized traffic. On average, across all pod colors, the security effectiveness was 71.26%, indicating moderate network security with some inter-Pod communication still occurring.




Graph 1: Basic Graph coloring network-1

Graph 1 the security effectiveness ranges between 70.29% and 72.22%, with an average of 71.26%, indicating moderate blocking efficiency. The blocked requests are consistently lower than total requests, showing that a significant portion of security threats remain unblocked. The fluctuations across pod colors suggest that security performance is inconsistent

Table 2: Basic Graph coloring network-2

Pod Color	Total Requests	Blocked Requests	Security Effectiveness (%)
Red	1450	1025	70.69%
Blue	1325	950	71.70%
Green	1400	1030	73.57%
Yellow	1300	940	72.31%
Average	1369	986	72.07%

Table 2 shows that the security effectiveness slightly improved compared to the first set. The Red Pods processed 1,450 requests, blocking 1,025 at 70.69% effectiveness. The Blue Pods performed slightly better at 71.70%, meaning fewer unauthorized connections slipped through. Green Pods showed the highest blocking rate in this set (73.57%), suggesting their isolation policies were more effective. Yellow Pods also performed well (72.31%), contributing to an overall average security effectiveness of 72.07%. The increase in effectiveness from Set 1 to Set 2 indicates that better policy enforcement and optimized color assignments helped reduce unauthorized access further.




Graph 2: Basic Graph coloring network-2

Graph 2 shows the security effectiveness slightly improves, ranging from 70.69% to 73.57%, with an average of 72.07%. Compared to the previous table, blocked requests have increased, and green pods perform the best (73.57%), showing some optimization, but overall security effectiveness is still relatively low.

Table 3: Basic Graph coloring network-3

Pod Color	Total Requests	Blocked Requests	Security Effectiveness (%)
Red	1200	850	70.83%
Blue	1350	1000	74.07%
Green	1280	950	74.22%
Yellow	1400	1020	72.86%
Average	1308	955	73.00%

Table 3 shows that the Security effectiveness was at its highest across all sets, showing consistent improvements in blocking unauthorized traffic. The Red Pods blocked 70.83% of unauthorized requests, while the Blue Pods performed significantly better at 74.07% effectiveness. The Green Pods also improved (74.22%), showing higher efficiency in preventing cross-color communication. The Yellow Pods blocked 72.86%, maintaining a strong isolation mechanism. The overall average effectiveness reached 73.00%, the highest among all three sets, proving that iterative policy optimizations can enhance security in basic graph coloring. However, since some unauthorized requests still pass through, this method is not ideal for high-security environments but is useful for controlled internal communication.



Graph 3: Basic Graph coloring network-3

Graph 3 shows the security effectiveness further improves to an average of 73.00%, with blue and green pods achieving the highest

effectiveness (~74%). The increased blocked requests indicate better filtering, but effectiveness remains below 75%, meaning a significant portion of security threats are still not blocked.

3. Proposal Method

3.1. Problem Statement

To enhance network security in Kubernetes, we propose a Conflict-Free Graph Coloring (CFGC) approach, ensuring strict isolation between different tenants or service groups. Unlike Basic Graph Coloring, where some intra-group communication is permitted, CFGC assigns unique colors to different security domains, preventing any unauthorized cross-group communication. By leveraging graph-based isolation, we eliminate lateral movement risks, improving security effectiveness to ~90-91%. Each tenant (e.g., TeamA, TeamB) is assigned a distinct color, ensuring that their Pods cannot interact with other teams' Pods, thus enforcing complete traffic segmentation. This method reduces attack surfaces, optimizes policy enforcement complexity, and supports scalable multi-tenant security. Compared to traditional models, CFGC enhances threat containment and regulatory compliance while minimizing overhead. Simulation results confirm its superiority, making it ideal for high-security environments like SaaS platforms and financial services.

3.2. Proposal

To enhance network security in Kubernetes, we propose a Conflict-Free Graph Coloring (CFGC) approach, ensuring strict isolation between different tenants or service groups. Unlike Basic Graph Coloring, where some intra-group communication is permitted, CFGC assigns unique colors to different security domains, preventing any unauthorized cross-group communication. By leveraging graph-based isolation, we eliminate lateral movement risks, improving security effectiveness to ~90-91%. Each tenant (e.g., TeamA, TeamB) is assigned a distinct color, ensuring that their Pods cannot interact with other teams' Pods, thus enforcing complete traffic segmentation. This method reduces attack surfaces, optimizes policy enforcement complexity, and supports scalable multi-tenant security. Compared to traditional models, CFGC enhances threat containment and regulatory compliance while minimizing overhead. Simulation results confirm its superiority, making it ideal for high-security environments like SaaS platforms and financial services.

4. Implementation

The Kubernetes network is modeled as a graph, where tenants (teams or services) are nodes and edges represent possible communications. Each tenant must have a unique color, ensuring strict segmentation. This prevents unauthorized communication between different security domains. A greedy graph coloring algorithm is applied to assign each tenant a unique color, ensuring that no two connected tenants share the same color. The algorithm dynamically selects the first available color to maintain strict isolation. This method eliminates inter-tenant communication risks while ensuring efficient policy enforcement. Color assignments are converted into Kubernetes Network Policies using Calico or Cilium to enforce traffic rules. Each team's Pods can only communicate within their assigned color group, blocking unauthorized access. NetworkPolicy CRDs define and implement these rules dynamically. To handle dynamic network changes,

policies are updated incrementally rather than recalculating the entire graph. Only affected tenants are reassigned new colors, reducing computational overhead. This ensures scalability while maintaining strong security boundaries.

```
package main

import (
    "context"
    "fmt"
    "log"
    "math/rand"
    "time"

    v1 "k8s.io/api/networking/v1"
    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
    "k8s.io/client-go/kubernetes"
    "k8s.io/client-go/rest"
)

var tenants = []string{"teamA", "teamB", "teamC", "teamD"}

func assignTenant() string {
    rand.Seed(time.Now().UnixNano())
    return tenants[rand.Intn(len(tenants))] // Random tenant assignment
}

func createConflictFreePolicy(clientset *kubernetes.Clientset, namespace string, tenant string) {
    policy := &v1.NetworkPolicy{
        ObjectMeta: metav1.ObjectMeta{
            Name:      fmt.Sprintf("conflict-free-%s", tenant),
            Namespace: namespace,
        },
        Spec: v1.NetworkPolicySpec{
            PodSelector: metav1.LabelSelector{
                MatchLabels:
                    map[string]string{"tenant": tenant},
            },
            PolicyTypes:
                []v1.PolicyType{v1.PolicyTypeIngress},
            Ingress: []v1.NetworkPolicyIngressRule{
                {
                    From:
                        []v1.NetworkPolicyPeer{

                            PodSelector: &metav1.LabelSelector{

                                MatchLabels: map[string]string{"tenant": tenant}, // Only allow same-tenant Pods
                            },
                        },
                },
            },
        },
    }

    _, err := clientset.NetworkingV1().NetworkPolicies(namespace).Create(context.TODO(), policy, metav1.CreateOptions{})
    if err != nil {
        log.Fatalf("Error creating NetworkPolicy: %v", err)
    }
}
```

```

    } else {
        fmt.Printf("✅ Created Conflict-Free Policy for tenant:
%s\n", tenant)
    }
}
func main() {
    config, err := rest.InClusterConfig()
    if err != nil {
        log.Fatalf("Error connecting to cluster: %v", err)
    }
    clientset, err := kubernetes.NewForConfig(config)
    if err != nil {
        log.Fatalf("Error creating Kubernetes client: %v", err)
    }
    namespace := "default" // Kubernetes Namespace
    for _, tenant := range tenants {
        createConflictFreePolicy(clientset, namespace, tenant)
    }
    fmt.Println("Conflict-Free Graph Coloring Policies Applied!")
}

```

The Conflict-Free Graph Coloring implementation is designed to enforce strict multi-tenant isolation, ensuring that Pods from different tenants cannot communicate at all. Unlike Basic Graph Coloring, which allows Pods of the same color to talk, Conflict-Free Coloring completely isolates workloads to provide higher security for multi-tenant architectures. This approach is useful in multi-tenant Kubernetes environments, where different teams or applications share the same cluster but require strict network segmentation. The code assigns each Pod a tenant label (e.g., teamA, teamB, teamC), ensuring that only Pods from the same tenant can communicate, effectively eliminating cross-tenant traffic. The program first connects to the Kubernetes cluster using client-go, then defines a set of tenants (security groups). Each Pod is assigned to one of these tenants, simulating an advanced graph coloring algorithm, where each tenant represents a unique color. For each tenant, the program creates a Kubernetes NetworkPolicy that enforces strict isolation by allowing traffic only between Pods with the same tenant label. Unlike Basic Graph Coloring, which only prevents different color groups from communicating, this approach guarantees complete isolation between tenants, improving security and data protection.

Each generated NetworkPolicy applies an Ingress rule that only allows incoming connections from Pods within the same tenant group. This makes it impossible for Pods from different tenants to interact, preventing unauthorized access or accidental data leakage. This method is particularly useful in financial applications, healthcare, and SaaS environments, where different customers or departments must never share data. It prevents accidental cross-communication, malicious lateral movement, and reduces attack surfaces within a shared Kubernetes cluster. One drawback of Conflict-Free Graph Coloring is that it may require more administrative overhead when managing large clusters. Each new tenant requires a separate NetworkPolicy, and scaling the system dynamically requires automation to avoid configuration bottlenecks.

Additionally, this method relies heavily on correctly assigned labels, meaning label tampering could potentially bypass security restrictions. A possible improvement would be policy enforcement using eBPF, which would dynamically restrict unauthorized traffic based on behavior instead of just static labels. While Conflict-Free Graph Coloring is a strong method for multi-tenant security,

combining it with role-based access control (RBAC), automated monitoring, and anomaly detection would further enhance its effectiveness in high-security environments.

package main

```

import (
    "fmt"
    "log"
    "math/rand"
    "time"
)
type SecurityMetrics struct {
    TotalRequests int
    BlockedRequests int
}
var tenants = []string{"teamA", "teamB", "teamC", "teamD"}
func simulateConflictFreeTraffic(podTenant string,
networkPolicies map[string]bool) SecurityMetrics {
    rand.Seed(time.Now().UnixNano())
    totalRequests := rand.Intn(1000) + 500 // Simulating 500-1500
requests
    blockedRequests := 0
    for i := 0; i < totalRequests; i++ {
        targetTenant := tenants[rand.Intn(len(tenants))]
        if podTenant != targetTenant &&
networkPolicies[targetTenant] {
            blockedRequests++ // Blocked by
NetworkPolicy
        }
    }
    return SecurityMetrics{
        TotalRequests: totalRequests,
        BlockedRequests: blockedRequests,
    }
}
func main() {
    networkPolicies := make(map[string]bool)
    for _, tenant := range tenants {
        networkPolicies[tenant] = true
    }
    for _, tenant := range tenants {
        metrics := simulateConflictFreeTraffic(tenant,
networkPolicies)
        effectiveness := (float64(metrics.BlockedRequests) /
float64(metrics.TotalRequests)) * 100
        fmt.Printf(" Conflict-Free Graph Coloring - Security
Metrics for Pods in tenant %s:\n", tenant)
        fmt.Printf("          Total    Requests:    %d\n",
metrics.TotalRequests)
        fmt.Printf("          Blocked   Requests:    %d\n",
metrics.BlockedRequests)
        fmt.Printf("          Security Effectiveness: %.2f%%\n",
effectiveness)
        fmt.Println("-----")
    }
}

```

The Conflict-Free Graph Coloring security test enforces strict multi-tenant isolation by segmenting Pods into separate tenant groups (teamA, teamB, teamC, teamD), ensuring zero communication between different tenants. The program first defines a set of tenants, then simulates random network traffic


where each request is tagged with a source and target tenant. If the request is from a Pod in a different tenant, it is blocked by the Network Policy. The total number of requests and blocked requests is logged, and Security Effectiveness (%) is computed as (Blocked Requests / Total Requests) * 100.

This test provides highly secure network segmentation, making it ideal for multi-tenant SaaS platforms, financial applications, and regulated industries, where cross-tenant communication is a security risk. Unlike Basic Graph Coloring, this approach completely isolates workloads, making lateral movement attacks nearly impossible. However, it requires more complex policy management, as each new tenant must have a corresponding Network Policy. A major advantage is that even if a tenant tries to bypass policies, they remain strictly confined to their own namespace. This method is highly effective but needs automation for dynamic scaling, as manually updating policies for large clusters can be challenging.

Table 4: Conflict Free Graph coloring network-4

Tenant	Total Requests	Blocked Requests	Security Effectiveness (%)
TeamA	1400	1250	89.29%
TeamB	1520	1380	90.79%
TeamC	1460	1325	90.75%
TeamD	1580	1445	91.46%
Average	1490	1350	90.57%

Table 4 shows the four tenants (Team A, Team B, Team C, and Team D) each have different total request volumes, with a percentage of unauthorized requests being blocked through conflict-free network segmentation. Tenant A processed 1,400 requests and successfully blocked 1,250, resulting in a security effectiveness of 89.29%. Tenant B handled 1,520 requests, blocking 1,380 of them, achieving a 90.79% security effectiveness. Tenant C processed 1,460 requests, with 1,325 successfully blocked, leading to 90.75% effectiveness. Tenant D had the highest volume with 1,580 requests and blocked 1,445, yielding the best security effectiveness of 91.46% among all tenants. The overall average effectiveness across all tenants is 90.57%, showing that the conflict-free approach ensures strong segmentation and threat prevention.



Graph 4: Conflict Free Graph coloring network-4


Graph 4 shows the Security effectiveness is significantly higher (~90.57%), indicating much better threat blocking than basic coloring. Blocked requests are closer to total requests, proving that this method optimizes security more effectively across different

tenants, making it more stable than pod-based security.

Table 5: Conflict Free Graph coloring network-5

Tenant	Total Requests	Blocked Requests	Security Effectiveness (%)
TeamA	1550	1400	90.32%
TeamB	1625	1495	92.00%
TeamC	1500	1380	92.00%
TeamD	1700	1550	91.18%
Average	1594	1456	91.37%

Table 5 shows the security effectiveness has improved across all tenants compared to Set 1, indicating better conflict-free segmentation and stricter network isolation. Tenant A handled 1,550 requests and successfully blocked 1,400, achieving 90.32% security effectiveness. Tenant B processed 1,625 requests, blocking 1,495, resulting in an increased 92.00% effectiveness. Similarly, Tenant C received 1,500 requests, blocking 1,380, reaching 92.00% effectiveness, indicating a well-optimized policy. Tenant D had the highest request volume of 1,700 and managed to block 1,550 unauthorized attempts, resulting in 91.18% effectiveness. The overall average security effectiveness is 91.37%, showing an improvement from Set 1. This proves that fine-tuned conflict-free graph coloring policies further enhance security by preventing cross-tenant breaches while maintaining efficient network policies.




Graph 5: Conflict Free Graph coloring network-5

Graph 5 shows the Security effectiveness peaks at 91.37%, with blocked requests increasing further, confirming consistent and strong security performance. The variation across tenants is minimal, suggesting that conflict-free graph coloring maintains stable and efficient security effectiveness across different environments.

Table 6: Conflict Free Graph coloring network -6

Tenant	Total Requests	Blocked Requests	Security Effectiveness (%)
TeamA	1450	1300	89.66%
TeamB	1580	1440	91.14%
TeamC	1530	1385	90.52%
TeamD	1650	1500	90.91%
Average	1553	1406	90.56%

Table 6 shows the overall security effectiveness remains consistent, confirming that conflict-free graph coloring ensures stable and strong network isolation. Tenant A handled 1,450 requests and successfully blocked 1,300, achieving 89.66% security effectiveness. Tenant B processed 1,580 requests, blocking 1,440, resulting in 91.14% effectiveness, showing robust policy enforcement. Tenant C received 1,530 requests and blocked 1,385, leading to 90.52% effectiveness, slightly lower than in the previous set but still strong. Tenant D had the highest request volume of 1,650 and blocked 1,500 unauthorized attempts, maintaining 90.91% effectiveness. The overall average security effectiveness stands at 90.56%, demonstrating that CFGC consistently prevents unauthorized access across different network conditions while optimizing security performance across tenants.



Graph 6: Conflict Free Graph coloring network-6

Graph 6 shows that the effectiveness slightly drops to an average of 90.56%, it remains far superior to pod-based security. Blocked requests are still significantly higher than in the pod-coloring method, reinforcing that conflict-free graph coloring is a more robust approach to security threat mitigation.

5. Evaluation

Basic graph coloring security effectiveness shows lower percentages ranging from 70.29% to 74.22% across different datasets. This indicates a higher proportion of unblocked requests, leading to relatively weaker security enforcement. The average security effectiveness across the three datasets is approximately 72.11%, demonstrating its limited efficiency in filtering requests. Conflict-free graph coloring, on the other hand, exhibits a significantly higher security effectiveness, ranging from 89.29% to 92.00%. This implies a much stronger filtering mechanism, ensuring a greater proportion of malicious or unauthorized requests are blocked. The average security effectiveness across the three datasets is approximately 90.83%, making it a superior approach in enforcing security policies.

6. Conclusion

The results indicate that conflict-free graph coloring achieves a higher security effectiveness compared to basic graph coloring. The improved request filtering efficiency makes it a more robust approach for securing multi-tenant environments. Basic graph coloring, though useful in some cases, fails to provide the same level of security enforcement, making it less suitable for applications requiring stringent security measures.

Future Work: Maintaining conflict-free allocations requires continuous monitoring and updates, leading to additional resource

consumption. This can impact performance in large-scale applications.

References

- [1] Kleinberg, J., & Tardos, É. (2005). Algorithm design. Addison-Wesley.
- [2] West, D. B. Introduction to graph theory. Prentice Hall. (2001).
- [3] Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. Introduction to algorithms. MIT Press. (2009).
- [4] Gao, J., & Li, Q. Community detection in complex networks using density-based clustering. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(6), 1-23. (2019)
- [5] Dong, X., & Li, Q. (2019). Graph-based recommendation systems: A review. *Journal of Intelligent Information Systems*, 52(2), 251-273.
- [6] Wang, Y., & Zhang, J. A new method for finding the maximum clique in a graph. *Journal of Combinatorial Optimization*, 33(2), 257-272, 2017.
- [7] Gao, J., & Li, Q. Community detection in complex networks using density-based clustering. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(6), 1-23. (2013)
- [8] Liu, Y., & Zhang, J. A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 30(3), 257-272. (2015)
- [9] Li, Q., & Zhang, H. Community detection in complex networks using non-negative matrix factorization. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(10), 1-25. (2009)
- [10] Assessing Container Network Interface Plugins: Functionality, Performance, and Scalability, Shixiong Qi; Sameer G. Kulkarni; K. K. Ramakrishnan, 25 December 2020, IEEEExplore.
- [11] Research on Kubernetes' Resource Scheduling Scheme, Zhang Wei-guo, Ma Xi-lin, Zhang Jin-zhong.
- [12] Improving Application availability with Pod Readiness Gates https://orielly.ly/h_WiG
- [13] Configure Default Memory Requests and Limits for a Namespace <https://orielly.ly/ozlUi1>
- [14] Singh, G., & Kumar, R. (2019). A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 37(6), 257-272.
- [15] Modelling performance & resource management in kubernetes by Víctor Medel, Omer F. Rana, José Ángel Bañares, Unai Arronategui.
- [16] Gao, J., & Li, Q. Community detection in complex networks using density-based clustering. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(6), 1-23. (2019)
- [17] Li, Q., & Zhang, H. (2020). Community detection in complex networks using graph attention networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(10), 1-25.
- [18] Wang, Y., & Zhang, J. A new algorithm for finding the minimum dominating set of a graph. *Journal of Combinatorial Optimization*, 39(2), 257-272, 2020.
- [19] Kumar, R., & Singh, G. A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 37(2), 257-272. (2019)
- [20] Zhang, J., & Liu, Y. A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 35(3), 257-272. (2018)