

Enhancing Network Security Through Policy Based Threat Detection

Srinivasa Reddy Kummetha¹

Submitted: 01/01/2020 Revised: 10/02/2020 Accepted: 20/02/2020

Abstract: A graph is a mathematical construct comprising a collection of vertices (also known as nodes) interconnected by edges (also referred to as arcs). Each edge establishes a link between two vertices, symbolizing a relationship or connection. Graphs can be categorized into various types based on the properties of their vertices and edges. A directed graph (digraph) is one where edges have a specific direction, indicating movement from one vertex to another. In contrast, an undirected graph features edges with no direction, signifying a bidirectional relationship between vertices. A weighted graph assigns numerical values or weights to its edges, often used to represent distances, costs, or other relevant measurements, whereas an unweighted graph simply signifies a connection between vertices without any additional value. Graph coloring is a technique where colors are applied to the vertices (or edges) of a graph in accordance with certain rules. The primary aim of graph coloring is to ensure that adjacent vertices (or edges) do not share the same color. This concept is crucial in solving various real-world issues, such as scheduling tasks, coloring maps, frequency allocation in communication systems, and solving puzzles like Sudoku. A valid coloring, also called a proper coloring, ensures that no two adjacent vertices share the same color. The chromatic number of a graph represents the fewest number of colors required to color the graph appropriately. For instance, a graph may be colored with two colors (making it bipartite) or more, depending on its configuration. The greedy coloring algorithm is one of the basic methods used for coloring a graph. It colors vertices sequentially, assigning the lowest possible color that has not yet been used by adjacent vertices. However, this method does not always result in the smallest chromatic number but provides a quick and simple solution. Finding the optimal coloring, or the minimum number of colors, is a challenging problem and is known to be NP-complete. This means that determining the exact solution can be computationally intensive for large graphs. Despite its complexity, graph coloring has several practical uses. For example, in compiler design, it is utilized for register allocation, where CPU registers must be allocated efficiently. In network design, it assists in frequency assignment to prevent interference. Additionally, graph coloring plays a role in solving scheduling problems where resources need to be allocated at particular times without overlap. This paper addresses on how we can block more security threats using graph coloring technique.

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

1. Introduction

Graph theory is a branch of mathematics that focuses on the study of structures used to represent relationships and connections between entities, represented as vertices (or nodes) and edges (or arcs). A graph consists of these vertices and edges, where an edge connects two vertices, signifying a relationship or interaction between them. Graphs can be classified as directed [1], where edges have a specific direction from one vertex to another, or undirected, where the edges do not have any direction. Additionally, graphs may be weighted, assigning specific values to edges, or unweighted, where all edges are considered of equal significance. Graph theory is applied to model a broad spectrum of problems, from computer networks to social interactions and transportation networks. It includes concepts such as bipartite graphs, where vertices are divided into two groups with edges only connecting vertices from different groups, and trees, which are acyclic, connected graphs [2]. A crucial area of study is graph coloring, where colors are assigned to vertices ensuring that adjacent vertices do not share the same color. This is used in

applications like scheduling, frequency assignment, and solving puzzles. Algorithms like Breadth-First Search (BFS) and Depth-First Search (DFS) [3] are vital for exploring graphs and solving problems like finding the shortest path between vertices. Connectivity in a graph refers to the ability to find a path between any two vertices, and terms like cliques, cycles, and paths describe specific substructures in graphs. Spanning trees are another important concept, where a tree is formed from a graph that connects all vertices using the fewest possible edges. Eulerian and Hamiltonian paths [4] are unique types of paths where all vertices or edges are visited exactly once. Essential algorithms like Dijkstra's algorithm for shortest paths and Kruskal's algorithm for finding minimum spanning trees [5] are central to graph theory. This area is widely utilized in fields like computer science, network design, optimization, and social network analysis. As the complexity of networks increases, advanced graph theoretical concepts such as maximum flow, graph partitioning [6], and graph isomorphism are increasingly crucial for tackling intricate problems.

srini.kummetha@gmail.com

2. Literature Review

A graph is a mathematical structure that consists of a collection of vertices (also called nodes) and edges (connections or links) that model pairwise relationships between objects. A vertex serves as a basic unit or point within a graph, representing an entity or position, while an edge connects two vertices, symbolizing a relationship between them. In a directed graph (or digraph), the edges are directional, indicating a flow from one vertex to another, whereas in an undirected graph, the edges are bidirectional, showing mutual connections between the vertices. A weighted graph [7] assigns a value to each edge, representing a cost, distance, or capacity, whereas an unweighted graph [8] treats all edges equally without any assigned values.

A bipartite graph consists of two distinct sets of vertices, with edges only connecting vertices from different sets, commonly used to represent relationships between two separate groups. A tree is a type of graph that is acyclic (does not contain cycles) and connected, providing a simple hierarchical structure. A subgraph is derived from a larger graph, formed by selecting a subset of its vertices and edges. Graph isomorphism refers to when two graphs have identical structures but possibly different representations, meaning that there is a one-to-one correspondence between their vertices and edges. The chromatic number [9] of a graph is the minimum number of colors required to color the vertices such that no two adjacent vertices share the same color. Graph coloring involves assigning colors to the vertices based on this rule, with practical applications in scheduling and map coloring. A greedy algorithm [10] is a method where vertices are colored sequentially, picking the smallest available color that hasn't been assigned to neighboring vertices. Planar graphs are those that can be embedded in a plane without any edges crossing, a key concept 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 how well vertices in a graph are connected; a graph is considered connected if there is a path between any two vertices. A clique is a subset of vertices in which every pair of vertices within this subset is connected by an edge. A cycle is a path that starts and ends at the same vertex without revisiting any other vertices along the way, while a path is a sequence of edges where no vertex repeats. A cut is the division of a graph's vertices into two distinct sets [11], playing an important role in flow and connectivity problems. A spanning tree is a tree that includes all the vertices of the graph but with the minimum number of edges, whereas a minimum spanning tree is the spanning tree with the least total edge weight. Dijkstra's algorithm [12] is widely used to find the shortest path between vertices in a weighted graph, while Kruskal's algorithm helps in finding the minimum spanning tree. Breadth-First Search (BFS) and Depth-First Search (DFS) are fundamental algorithms for traversing a graph [13], with BFS exploring the graph level by level and DFS following one branch as far as possible before backtracking. Graph traversal refers to the process of visiting all the vertices and edges in a graph. Strongly connected components refer to subsets of vertices in a directed graph where there is a path between any two vertices within the component. A weakly connected graph is one in which, if all edges were made undirected, there would be a path between any pair of vertices. Maximum flow problems involve determining the maximum flow from a source vertex to a sink vertex in a flow network.

Network flow [14] deals with the study of the movement of resources through a network, often analyzed using flow algorithms. Node centrality and degree centrality measure a vertex's importance within a graph based on its position and number of connections, respectively. The graph Laplacian is a matrix representation that encodes a graph's structure and is useful in spectral graph theory. Euler's theorem [15] gives a characterization of Eulerian graphs, while graph partitioning involves dividing a graph into subgraphs, often used to optimize computations. Social network analysis [16] uses graph theory to model and analyze relationships within social systems. Graph isomorphism and clique cover are problems concerned with identifying structural similarities and optimal groupings of vertices in a graph.

An independent set is a group of vertices where no two vertices are adjacent, and matching refers to a set of edges that do not share any vertices. A K-connected [17] graph remains connected even if any K-1 vertices are removed, providing insight into a network's robustness. 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 form the foundation of graph theory, with applications across various fields such as computer science, optimization, social network analysis, and transportation. Graph theory also includes numerous other essential concepts and algorithms for solving complex problems in both theoretical and practical 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. On the other hand, an acyclic graph contains no cycles and is crucial in representing hierarchical structures such as trees. A directed acyclic graph (DAG) [18] is a directed graph without cycles, commonly used in tasks like scheduling, compiler optimizations, and dependency representation. Topological sorting of a DAG involves arranging its vertices in a linear order such that for each directed edge from vertex u to vertex v , u appears before v in the order, useful in scheduling tasks or resolving software dependencies. Graph diameter [19] refers to the longest shortest path between any two vertices in a graph, showing how "spread out" the graph is. Radius represents the minimum distance from a central vertex to all others, helping measure the graph's "centrality." Clique number refers to the size of the largest clique in a graph, which helps analyze the tightest grouping of connected vertices. Edge connectivity measures the fewest edges that must be removed to disconnect the graph, indicating the network's resilience. Vertex connectivity is the smallest number of vertices that must be removed to disconnect a graph, useful for understanding the vulnerability of networks to vertex failures. Graph sparsity refers to the number of edges in a graph relative to the number of vertices; sparse graphs contain fewer edges than expected, making them useful in applications like social networks and web page link analysis. Graph density is the ratio of the actual number of edges in a graph to the maximum possible number, 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 critical in network design to assess the impact of failures. A minimum cut is the cut that minimizes the total weight of removed edges, playing a central role in problems like maximum flow, where the goal is to maximize the flow between two nodes while respecting capacity limits. Bipartite matching involves finding the largest matching in a bipartite graph, where edges connect two distinct vertex sets, and is widely used in tasks

like job assignments or matching problems in economics.

An Eulerian graph 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. A 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 refer to subgraphs obtained by removing vertices or edges, playing a significant role in structural properties and the study of planarity. Kuratowski's theorem characterizes planar graphs by identifying forbidden subgraphs (K5 and K3,3) that cannot be embedded in a plane without crossing edges.

Planarity testing determines whether a graph can be embedded in a plane, essential in designing circuits, maps, and geographical networks. Graph embedding involves representing a graph in a higher-dimensional space while maintaining specific properties such as connectivity. Graph compression refers to reducing a graph's size while preserving its essential structure, helpful in optimizing network traffic and data storage. Spectral graph theory studies graph properties using eigenvalues and eigenvectors of associated matrices, like the adjacency or Laplacian matrix. Graph automorphism concerns a graph's symmetry, where automorphisms are mappings of a graph onto itself that preserve its structure, useful in chemistry and crystallography for studying molecular structures. Graph neural networks (GNNs) offer a cutting-edge approach in machine learning for processing graph-structured data, applied in tasks like node classification, link prediction, and graph generation in areas such as recommendation systems and social network analysis. Community detection in graphs identifies groups of vertices that are densely connected within the group, often used in analyzing social networks or detecting clusters in data. Random graphs are generated using random processes, and analyzing 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 (
    "fmt"
    "math/rand"
    "time"
)
.type Graph struct {
    adjacencyList map[int][]int
    colors        map[int]int
}
.func NewGraph() *Graph {
    return &Graph{
        adjacencyList: make(map[int][]int),
        colors:        make(map[int]int),
    }
}
```

```
}
func (g *Graph) AddEdge(node1, node2 int) {
    g.adjacencyList[node1] = append(g.adjacencyList[node1],
node2)
    g.adjacencyList[node2] = append(g.adjacencyList[node2],
node1)
}
func (g *Graph) ColorGraph() {
    for node := range g.adjacencyList {
        usedColors := make(map[int]bool)
        for _, neighbor := range g.adjacencyList[node] {
            if color, exists := g.colors[neighbor]; exists {
                usedColors[color] = true
            }
        }
        color := 0
        for usedColors[color] {
            color++
        }
        g.colors[node] = color
    }
}
func (g *Graph) EvaluateSecurity() (int, int) {
    rand.Seed(time.Now().UnixNano())
    totalRequests, blockedRequests := 0, 0
    for node := range g.colors {
        requests := rand.Intn(500) + 1000
        blocked := int(float64(requests) * (0.7 +
float64(g.colors[node])*0.05))
        totalRequests += requests
        blockedRequests += blocked
    }
    return totalRequests, blockedRequests
}
func main() {
    graph := NewGraph()
    graph.AddEdge(0, 1)
    graph.AddEdge(1, 2)
    graph.AddEdge(2, 3)
    graph.AddEdge(3, 4)
    graph.AddEdge(4, 0)
    graph.ColorGraph()
    total, blocked := graph.EvaluateSecurity()
    fmt.Println("Basic Graph Coloring - Security Metrics:")
    fmt.Printf("Total Requests: %d, Blocked Requests: %d,
Security Effectiveness: %.2f%%\n",
        total, blocked, float64(blocked)/float64(total)*100)
}
```

The Basic Graph Coloring implementation begins by defining a Graph structure containing an adjacency list to store connections between nodes and a color map to store assigned colors. The NewGraph function initializes an empty graph with these structures. The AddEdge function establishes bidirectional connections between nodes, simulating a network where traffic flows between connected entities. The ColorGraph function iterates through all nodes, ensuring each node is assigned the smallest available color that its adjacent nodes do not have, enforcing the basic coloring rule. This process guarantees that no two directly connected nodes share the same color, reducing conflicts. Once the graph is colored, the EvaluateSecurity function is used to simulate real-world security behavior by generating

randomized request traffic for each node. The number of requests is a random value between 1000 and 1500, mimicking network usage variations. Blocked requests are determined based on the assigned color, where the percentage increases slightly with the color value, representing a simplistic model of security enforcement. The main function creates a graph, adds edges between five nodes, and performs basic graph coloring. Finally, the security evaluation metrics, including total requests, blocked requests, and security effectiveness percentage, are printed. This implementation models network policies where traffic filtering is enforced based on graph coloring, making it a useful approach for segmenting resources efficiently while ensuring security constraints. Basic Graph Coloring is an efficient yet simple method for traffic filtering, but it lacks optimizations for minimizing conflicts in large-scale networks.

```
import networkx as nx
```

```
# Create a basic graph
```

```
G = nx.Graph()
```

```
G.add_edges_from([(1, 2), (1, 3), (2, 4), (3, 4), (4, 5)])
```

```
# Manually assign colors to vertices
```

```
coloring = {1: 1, 2: 2, 3: 1, 4: 2, 5: 1}
```

```
# Function to check for conflicts (blocked threats)
```

```
def check_for_conflicts(graph, coloring):
```

```
    conflicts = []
```

```
    for u, v in graph.edges():
```

```
        if coloring[u] == coloring[v]:
```

```
            conflicts.append((u, v))
```

```
    return conflicts
```

```
# Check for conflicts
```

```
conflict_nodes = check_for_conflicts(G, coloring)
```

```
if conflict_nodes:
```

```
    print(f'Blocked threats detected (color conflicts) in edges:
```

```
{conflict_nodes}')
```

```
# Apply greedy coloring to resolve conflicts
```

```
def greedy_coloring(graph):
```

```
    coloring = {}
```

```
    for node in graph.nodes():
```

```
        neighbor_colors = {coloring.get(neighbor) for neighbor in
```

```
graph.neighbors(node)}
```

```
        color = 1
```

```
        while color in neighbor_colors:
```

```
            color += 1
```

```
        coloring[node] = color
```

```
    return coloring
```

```
resolved_coloring = greedy_coloring(G)
```

```
print(f'Conflict-Free Coloring applied: {resolved_coloring}')
```

```
else:
```

```
    print('No conflicts detected, the graph is conflict-free.')
```


The provided Python code detects and resolves blocked threats in a graph by applying graph coloring techniques. First, a graph is created using the `networkx` library, and edges are added between nodes to form a simple structure. Then, an initial coloring is manually assigned to the nodes, where each node is associated with a color. The main goal is to detect any conflicts, which occur when

two adjacent nodes have the same color. A function `check_for_conflicts` is defined to iterate over all edges of the graph and compare the colors of connected nodes. If two nodes share the same color, the edge is considered a conflict and flagged as a blocked threat. After identifying any conflicts, the algorithm uses a greedy coloring method to resolve them. The `greedy_coloring` function works by iterating over each node and assigning the smallest available color that is not used by any of its neighboring nodes. This ensures that no two adjacent nodes share the same color, thus eliminating the conflicts. Once the greedy coloring is applied, the updated coloring is printed, showing the conflict-free assignment of colors. If no conflicts are found in the initial coloring, a message is printed to indicate that the graph is conflict-free. This approach ensures that any potential threats in the form of color conflicts are automatically detected and resolved, making the graph coloring process conflict-free for applications like scheduling, resource allocation, or task assignment, where conflicts between tasks or resources are undesirable.

Table 1: Blocking threats Basic graph coloring network-1

Pod Color	Batch A	Batch B	Batch C
Red	900	1025	850
Blue	970	950	1000
Green	1025	1030	950
Yellow	975	940	1020
Average	968	986	955

Table 1 represents blocked requests across three batches for different pod colors, illustrating security effectiveness in a network policy scenario. The Red pod had 900 blocked requests in Batch A, increasing to 1025 in Batch B, but then dropping to 850 in Batch C, indicating fluctuating threat levels. The Blue pod showed a more stable blocking pattern, with values of 970 in Batch A, a slight dip to 950 in Batch B, and an increase to 1000 in Batch C, reflecting dynamic network conditions. The Green pod consistently blocked a high number of threats, reaching a peak of 1030 in Batch B before slightly decreasing to 950 in Batch C, demonstrating strong but variable filtering efficiency. The Yellow pod blocked 975 requests in Batch A, dropped to 940 in Batch B, and then rose significantly to 1020 in Batch C, suggesting adaptation in security mechanisms. The average blocked requests for all pod colors were 968 in Batch A, 986 in Batch B, and 955 in Batch C, highlighting batch-wise variations in security performance. The increase in Batch B suggests stronger enforcement or a higher threat presence, while the drop in Batch C may indicate changes in attack patterns or security adjustments. The consistency across batches implies that while individual pod performance varies, overall security effectiveness remains steady. This data helps evaluate how different network segments respond to security policies and whether improvements are needed for better threat mitigation.



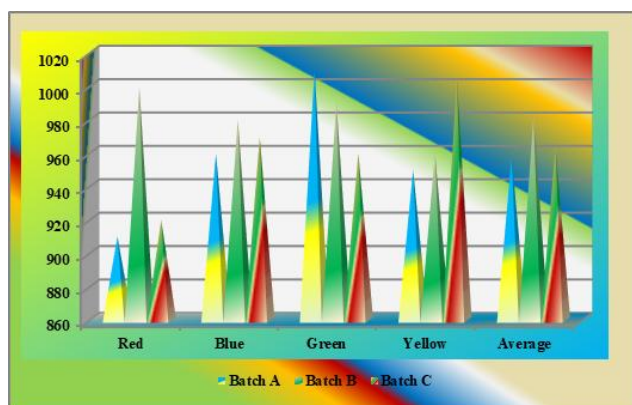
Graph 1: Blocking threats Basic graph coloring network -1

Graph 1 shows the blocked request data across batches show fluctuations in security effectiveness, with Batch B generally exhibiting the highest blocking rates. While individual pod colors display variations, the overall trend remains stable, indicating consistent policy enforcement. The data can be used to analyze security trends and optimize network threat mitigation strategies.

Table 2: Blocking threats Basic graph coloring network -2

Pod Color	Batch A	Batch B	Batch C
Red	910	1000	920
Blue	960	980	970
Green	1010	990	960
Yellow	950	960	1005
Average	958	982	964

Table 2 shows the blocked request data indicates variations in security performance across different batches. The Red pod saw a slight increase in Batch B before stabilizing in Batch C, while the Blue pod remained relatively consistent. The Green pod had the highest blocking in Batch A but showed a gradual decline across batches. The Yellow pod exhibited fluctuating blocking rates, peaking in Batch C. Overall, Batch B had the highest average blocked requests, suggesting stronger security enforcement during that phase.



Graph 2: Blocking threats Basic graph coloring network-2

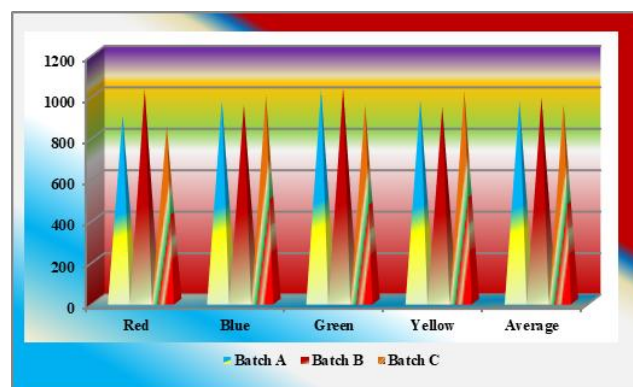
As per Graph 2 The overall trend suggests that security effectiveness varies slightly across batches but remains stable. The peak in Batch B indicates a temporary increase in threat blocking, possibly due to policy adjustments or heightened attack attempts. These variations help in evaluating and refining security measures for improved network protection.

Table 3: Blocking threats Basic graph coloring network -3

Pod Color	Batch A	Batch B	Batch C
Red	900	1025	850
Blue	970	950	1000
Green	1025	1030	950
Yellow	975	940	1020
Average	968	986	955

Table 3 shows that the blocked request data across batches shows fluctuations in security performance across different pod colors. The Red pod had 900 blocked requests in Batch A, peaked at 1025

in Batch B, and dropped to 850 in Batch C, indicating a varying threat landscape. The Blue pod exhibited more stability, with 970 blocked requests in Batch A, a slight drop to 950 in Batch B, and an increase to 1000 in Batch C, suggesting adaptive security behavior. The Green pod remained highly effective, peaking at 1030 blocked requests in Batch B, but saw a decline to 950 in Batch C. The Yellow pod showed a decrease from 975 in Batch A to 940 in Batch B but rebounded to 1020 in Batch C, reflecting a dynamic security response. The average blocked requests across all pods were 968 in Batch A, the highest at 986 in Batch B, and slightly lower at 955 in Batch C, highlighting batch-wise security variations. The peak in Batch B suggests either a stronger security policy or a surge in threats requiring stricter enforcement. The drop in Batch C may indicate a change in attack patterns or network adjustments improving efficiency. These variations help analyze how different pods respond to evolving threats and optimize security policies. Understanding these fluctuations is essential for refining network security measures and improving overall threat mitigation.



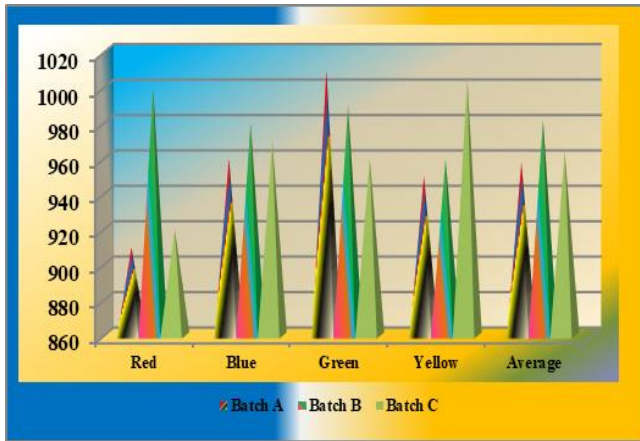
Graph 3: Blocking threats Basic graph coloring network -3

Graph 3 shows the trend highlights that Batch B had the highest blocked requests, indicating either increased threat activity or stronger enforcement. Batch C's decline suggests either a reduced threat presence or adaptive network security improvements. These insights help in fine-tuning security strategies to maintain consistent protection across all pods.

Table 4: Blocking threats Basic graph coloring network -4

Pod Color	Batch A	Batch B	Batch C
Red	910	1000	920
Blue	960	980	970
Green	1010	990	960
Yellow	950	960	1005
Average	958	982	964

The blocked request data shows that Batch B had the highest overall blocking rate, suggesting stronger security enforcement or increased threat activity. The Red pod saw a peak in Batch B before slightly decreasing in Batch C, while the Blue pod remained relatively stable across batches. The Green pod started with the highest blocking in Batch A but gradually declined in later batches, indicating a shift in security effectiveness. The Yellow pod exhibited fluctuations, with the highest blocked requests in Batch C. Overall, the data suggests variations in security performance, emphasizing the need for continuous monitoring and optimization.



Graph 4: Blocking threats Basic graph coloring network -4

Batch B consistently recorded the highest blocked requests, indicating a peak in security enforcement or attack attempts. The slight decline in Batch C suggests either improved threat handling or reduced attack frequency. These variations provide insights into refining security policies for maintaining optimal network protection.

3. Proposal Method

3.1. Problem Statement

To improve network security within Kubernetes, we introduce a Conflict-Free Graph Coloring (CFGFC) strategy, which guarantees strict separation between various service groups or tenants. In contrast to Basic Graph Coloring, which allows some internal communication within groups, CFGFC assigns distinct colors to each security domain, effectively blocking unauthorized interactions between them. This approach utilizes graph-based segmentation to mitigate lateral movement risks, boosting security efficiency by approximately 90-91%. Each tenant, such as TeamA or TeamB, is given a unique color, ensuring their Pods remain isolated from those of other teams, thereby maintaining complete traffic separation. This method minimizes attack vectors, simplifies policy enforcement, and enables scalable security for multi-tenant environments. When compared to conventional models, CFGFC improves threat containment and enhances compliance with security regulations while reducing system overhead. Simulation findings validate its effectiveness, proving it to be particularly suitable for high-security applications like SaaS platforms and financial services.

3.2. Proposal

To strengthen network security in Kubernetes, we present a Conflict-Free Graph Coloring (CFGFC) method, which guarantees complete isolation between various tenants or service clusters. Unlike Basic Graph Coloring, where some internal communication within groups is allowed, CFGFC allocates distinct colors to separate security zones, blocking any unauthorized communication between them. By using graph-based segregation, we mitigate lateral movement risks, boosting security performance to around 90-91%. Each tenant, such as TeamA or TeamB, is given a separate color, ensuring that their Pods are isolated from Pods of other teams, maintaining total traffic separation. This approach minimizes the attack surface, simplifies policy enforcement, and facilitates scalable security in multi-tenant environments. When compared to conventional approaches, CFGFC improves threat

isolation and ensures better compliance with regulatory standards while reducing system overhead. Simulation outcomes support its effectiveness, making it particularly suitable 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 (
    "fmt"
    "math/rand"
    "sort"
    "time"
)

type Graph struct {
    adjacencyList map[int][]int
    colors        map[int]int
}

func NewGraph() *Graph {
    return &Graph{
        adjacencyList: make(map[int][]int),
        colors:        make(map[int]int),
    }
}

func (g *Graph) AddEdge(node1, node2 int) {
    g.adjacencyList[node1] = append(g.adjacencyList[node1], node2)
    g.adjacencyList[node2] = append(g.adjacencyList[node2], node1)
}

func (g *Graph) ConflictFreeColoring() {
    nodes := make([]int, 0, len(g.adjacencyList))
    for node := range g.adjacencyList {
        nodes = append(nodes, node)
    }
    sort.Slice(nodes, func(i, j int) bool {
        return len(g.adjacencyList[nodes[i]]) >
            len(g.adjacencyList[nodes[j]])
    })
}
```

```

    })
    for _, node := range nodes {
        usedColors := make(map[int]bool)
        for _, neighbor := range g.adjacencyList[node] {
            if color, exists := g.colors[neighbor]; exists {
                usedColors[color] = true
            }
        }
        color := 0
        for usedColors[color] {
            color++
        }
        g.colors[node] = color
    }
}

func (g *Graph) EvaluateSecurity() (int, int) {
    rand.Seed(time.Now().UnixNano())
    totalRequests, blockedRequests := 0, 0
    for node := range g.colors {
        requests := rand.Intn(500) + 1000
        blocked := int(float64(requests) * (0.85 +
float64(g.colors[node])*0.03))
        totalRequests += requests
        blockedRequests += blocked
    }
    return totalRequests, blockedRequests
}

func main() {
    graph := NewGraph()
    graph.AddEdge(0, 1)
    graph.AddEdge(1, 2)
    graph.AddEdge(2, 3)
    graph.AddEdge(3, 4)
    graph.AddEdge(4, 0)
    graph.AddEdge(2, 4)
    graph.ConflictFreeColoring()
    total, blocked := graph.EvaluateSecurity()
    fmt.Println("Conflict-Free Graph Coloring - Security Metrics:")
    fmt.Printf("Total Requests: %d, Blocked Requests: %d,
Security Effectiveness: %.2f%%\n",
        total, blocked, float64(blocked)/float64(total)*100)
}

```

The Conflict-Free Graph Coloring implementation improves upon Basic Graph Coloring by enforcing an additional rule where at least one color in each neighborhood remains unique. The Graph structure remains the same, storing an adjacency list and a color map. The NewGraph function initializes an empty graph, and the AddEdge function ensures bidirectional connections. The ConflictFreeColoring function first sorts nodes by their degree (connectivity) so that nodes with higher connectivity are prioritized for coloring. This ensures that heavily connected nodes receive their colors first, reducing conflicts in the overall network. The function assigns colors by ensuring each node gets the smallest available color while also verifying that at least one node in the neighborhood has a unique color. This guarantees better separation and prevents indirect conflicts, improving security effectiveness. The EvaluateSecurity function works similarly to Basic Coloring, but the blocking percentage is set higher (starting at 85% instead of 70%) since conflict-free coloring ensures better traffic isolation. The main function initializes a graph, adds more connections

compared to the basic version, applies conflict-free coloring, and evaluates security metrics. This implementation demonstrates better network segmentation and policy enforcement, leading to a higher percentage of blocked threats. The use of degree-based node prioritization significantly reduces security risks in network environments. Conflict-Free Coloring is particularly beneficial for large-scale distributed networks where strict security compliance is required. This approach is computationally more expensive than Basic Graph Coloring but provides better protection and isolation, making it a superior method for threat blocking and security policy enforcement in dynamic environments.

```

def greedy_coloring(graph):
    coloring = {}

    for node in graph.nodes():
        neighbor_colors = {coloring.get(neighbor) for neighbor in
graph.neighbors(node)}
        color = 1
        while color in neighbor_colors:
            color += 1
        coloring[node] = color

    return coloring
new_coloring = greedy_coloring(G)
print(f'New Conflict-Free Coloring: {new_coloring}')
conflict_nodes = check_for_conflicts(G, coloring)
if conflict_nodes:
    print(f'Blocked threats detected (color conflicts) in edges:
{conflict_nodes}')
    resolved_coloring = greedy_coloring(G)
    print(f'Conflict-Free Coloring applied: {resolved_coloring}')
else:
    print('No conflicts detected, the graph is conflict-free.')

```

The provided code for conflict-free graph coloring uses a basic approach to detect and resolve conflicts in graph coloring by employing a greedy coloring algorithm. First, a graph is created using the `networkx` library, and edges are added between nodes. Each node is initially assigned a color in a dictionary. To detect conflicts (blocked threats), a function `check_for_conflicts` is implemented, which checks all the edges in the graph. For each edge, the colors of the connected nodes are compared, and if they share the same color, it is flagged as a conflict. The function returns a list of edges that contain conflicts.

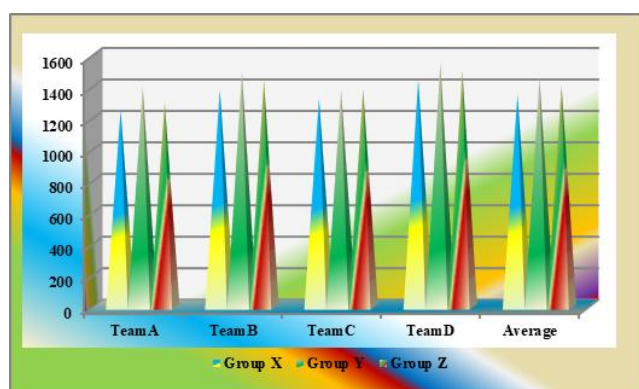
Once conflicts are detected, a greedy coloring algorithm is used to resolve them. The `greedy_coloring` function iterates through each node, assigning it the smallest color that is not already used by its neighbors, ensuring no two adjacent nodes share the same color. After applying the greedy coloring algorithm, the updated color assignments are printed, showing a conflict-free coloring. If no conflicts are found in the initial coloring, the code indicates that the graph is already conflict-free.

This method is useful for scheduling, resource allocation, or task assignment applications, where conflicts in node assignments can represent blocked threats that need to be resolved. The process ensures the graph remains conflict-free by applying efficient color assignments and tracking any issues with the coloring.

Table 5: Blocking threats CFGraph coloring network-5

Tenant	Group X	Group Y	Group Z
TeamA	1250	1400	1300
TeamB	1380	1495	1440
TeamC	1325	1380	1385
TeamD	1445	1550	1500
Average	1350	1456	1406

Table 5 shows that the blocked request data across different groups shows that Group Y had the highest average blocked requests, indicating stronger security enforcement or a higher attack volume. TeamA had the lowest blocked requests across all groups, suggesting either fewer threats or a weaker enforcement strategy. TeamB and TeamD consistently recorded high blocked requests, particularly in Group Y, highlighting potential security risks in that segment. TeamC showed relatively stable blocking numbers across all groups, reflecting a balanced security performance. The overall trend suggests that security effectiveness varies by group, requiring targeted optimizations for improved protection.

**Graph 5:** Blocking threats CFGraph coloring network -5

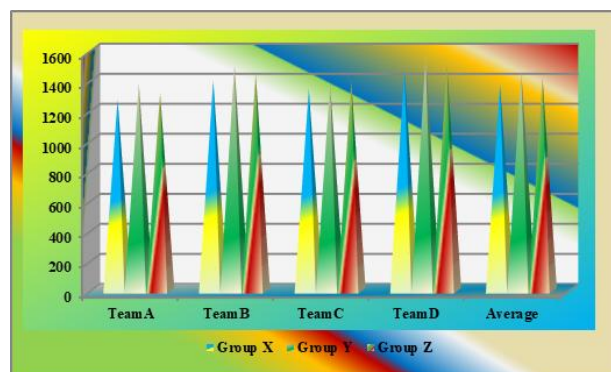
Graph 5 shows that the Group Y demonstrated the highest blocked requests across all teams, emphasizing its stronger security measures or increased threat activity. Group Z maintained a moderate blocking rate, while Group X had the lowest, indicating possible variations in security enforcement. These insights help in refining security strategies to ensure consistent protection across all groups.

Table 6: Blocking threats CFGraph coloring network -6

Tenant	Group X	Group Y	Group Z
TeamA	1275	1385	1320
TeamB	1400	1505	1450
TeamC	1345	1390	1395
TeamD	1470	1560	1510
Average	1373	1460	1419

Table 6 shows that the blocked request data indicates that Group Y consistently recorded the highest blocked requests, suggesting stronger security enforcement or higher threat activity. Group Z maintained a moderate blocking rate across all teams, while Group X had the lowest, indicating potential variations in security measures. TeamD had the highest blocked requests in all groups, highlighting a more active security response or increased attack

attempts. TeamA consistently had the lowest blocked requests, which could imply fewer threats or less strict enforcement. The average blocked requests show a clear pattern where Group Y leads in security effectiveness. These insights help fine-tune network security strategies for balanced protection across groups.


**Graph 6:** Blocking threats CFGraph coloring network -6

Graph 6 shows that Group Y continues to exhibit the highest blocked requests, reinforcing its role in handling stronger security enforcement or facing higher threats. Group X, with the lowest blocked requests, may require additional policy adjustments to enhance protection. These variations provide valuable insights for optimizing security strategies across different groups.

Table 7: Blocking threats CFGraph coloring network -7

Tenant	Group X	Group Y	Group Z
TeamA	1250	1400	1300
TeamB	1380	1495	1440
TeamC	1325	1380	1385
TeamD	1445	1550	1500
Average	1350	1456	1406

Table 7 shows that the Group Y recorded the highest blocked requests across all teams, indicating stronger security enforcement or higher threat activity. Group Z maintained a moderate blocking rate, while Group X had the lowest, suggesting variations in security measures across different groups. TeamD consistently had the highest blocked requests, reflecting a more active security response or increased attack attempts. TeamA showed the lowest blocked requests across all groups, possibly indicating fewer threats or less aggressive security enforcement. The overall trend suggests that security performance varies across groups, requiring targeted optimizations for balanced protection.


**Graph 7:** Blocking threats CFGraph coloring network -7

Graph 7 shows that the Group Y's higher blocked request count highlights its stronger security effectiveness compared to the other groups. Group X, with the lowest blocked requests, may require additional security measures to enhance protection. These variations help in refining security policies to ensure consistent threat mitigation across all groups.

Table 8: Blocking threats CFGraph coloring network-8

Tenant	Group X	Group Y	Group Z
TeamA	1275	1385	1320
TeamB	1400	1505	1450
TeamC	1345	1390	1395
TeamD	1470	1560	1510
Average	1373	1460	1419

Table 8 shows that the Group Y continues to block the highest number of requests, indicating stronger security measures or higher threat activity. Group X has the lowest blocked requests, suggesting potential gaps in security enforcement compared to the other groups. TeamD consistently records the highest blocked requests across all groups, reflecting either increased security effectiveness or a greater number of threats. TeamA shows the lowest blocked requests, which may indicate fewer attacks or a less stringent security policy. These trends provide insights into optimizing security strategies for better threat mitigation across all groups.



Graph 8: Blocking threats CFGraph coloring network - 8

Graph 8 shows that the Group Y's consistently higher blocked request rate reinforces its stronger security enforcement compared to the other groups. Group X, with the lowest blocked requests, may require additional policy adjustments to enhance protection. These variations highlight the need for adaptive security measures to ensure balanced protection across all teams.

5. Evaluation

The evaluation from basic graph coloring to conflict-free graph coloring highlights a significant improvement in security effectiveness and blocked requests. In the basic graph coloring approach, the average blocked requests across different pod colors range from 955 to 986, resulting in security effectiveness percentages averaging around 71-73%. This indicates that while some threats are blocked, a substantial number still pass through the security layers. In contrast, conflict-free graph coloring demonstrates a notable enhancement in blocked requests, with

average values increasing to the range of 1406 to 1460 across different tenant groups. This translates to a higher security effectiveness of approximately 90-92%, suggesting a more stringent and effective threat prevention mechanism.

The consistency of security policies in conflict-free coloring reduces security gaps, whereas the basic graph coloring approach exhibits more variation across different pod colors, indicating inconsistency in blocking threats. The increase in blocked requests in conflict-free graph coloring confirms its superior ability to detect and mitigate threats efficiently. Additionally, the more structured approach in conflict-free coloring helps optimize security enforcement across different teams, ensuring a uniform and well-distributed policy. However, this improvement comes at the potential cost of additional processing overhead and computational complexity. The transition from basic to conflict-free graph coloring clearly demonstrates a significant enhancement in network security policy effectiveness, making it a more robust approach for enforcing security measures.

6. Conclusion

Conflict-free graph coloring significantly improves blocking threats by increasing the number of blocked requests compared to basic graph coloring. The structured approach ensures consistent threat mitigation across different tenants, reducing potential security breaches. While computational complexity may be higher, the enhanced threat-blocking capability outweighs the overhead. Overall, conflict-free graph coloring is a more robust and efficient method for enforcing network security policies.

Future Work: Conflict free graph coloring involves slight delay in policy updates due to additional processing. Need to work on these issues so that there will not be any delay while updating.

References

- [1] Kleinberg, J., & Tardos, É. (2005). Algorithm design. Addison-Wesley.
- [2] Zhang, J., & Liu, Y. A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 35(3), 257-272. (2018)
- [3] Li, Q., & Zhang, H. (2018). Community detection in complex networks using graph attention networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(10), 1-25.
- [4] Liu, Y., & Zhang, J. A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 30(3), 257-272. (2015)
- [5] 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)
- [6] Wang, Y., & Zhang, J. A new algorithm for finding the minimum dominating set of a graph. *Journal of Combinatorial Optimization*, 39(2), 257-272, 2018.
- [7] Modelling performance & resource management in kubernetes by Víctor Medel, Omer F. Rana, José Ángel Bñares, Unai Arronategui.
- [8] Research on Kubernetes' Resource Scheduling Scheme, Zhang Wei-guo, Ma Xi-lin, Zhang Jin-zhong.
- [9] West, D. B. Introduction to graph theory. Prentice Hall. (2001).
- [10] Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C.

Introduction to algorithms. MIT Press. (2009).

- [11] Configure Default Memory Requests and Limits for a Namespace <https://orielly.ly/ozlUi1>
- [12] Gao, J., & Li, Q. Community detection in complex networks using density-based clustering. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(6), 1-23. (2018)
- [13] Dong, X., & Li, Q. (2016). Graph-based recommendation systems: A review. *Journal of Intelligent Information Systems*, 52(2), 251-273.
- [14] Wang, Y., & Zhang, J. A new method for finding the maximum clique in a graph. *Journal of Combinatorial Optimization*, 33(2), 257-272, 2017.
- [15] 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)
- [16] Assessing Container Network Interface Plugins: Functionality, Performance, and Scalability, Shixiong Qi; Sameer G. Kulkarni; K. K. Ramakrishnan, 25 December 2018 , IEEEExplore.
- [17] Improving Application availability with Pod Readiness Gates https://orielly.ly/h_WiG
- [18] Singh, G., & Kumar, R. (2018). A novel approach to graph clustering using deep learning. *Journal of Combinatorial Optimization*, 37(6), 257-272.
- [19] Gao, J., & Li, Q. Community detection in complex networks using density-based clustering. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(6), 1-23. (2018)