This article has been accepted for publication in IEEE Transactions on Services Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSC.2023.3253307

A Novel Team Formation Framework based on Performance in a Cybersecurity Operations Center

Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam, Steve Hutchinson

Abstract—A Cybersecurity Operations Center (CSOC) performs various tasks to protect an organization from cyber threats. Several types of personnel collaborate to function effectively as a team to analyze the threat signals, in the form of alerts, arriving from various sources. Teams are often formed ad hoc, resulting in an imbalance in their performances and thereby increasing the risk associated with the low-performing teams. The current approach taken by behavioral scientists in forming effective teams focuses on first qualitatively assessing individuals such as analysts, who are then grouped into teams based on their credentials and expertise. Our work takes a holistic view of the CSOC by first defining team requirements and then selecting individuals to form several collaborative teams that meet these requirements for every shift of operation. We present a novel team formation framework that integrates optimization, simulation, and scoring methods to form effective teams and introduce a new collaborative score metric that measures their effectiveness. Results from simulated experiments show the formation of effective teams whose collaborative scores are maximized and balanced. Our approach is also able to identify high and low performers within the first few months of implementing the framework.

Index Terms—Cyber Team Formation, Team Collaborative Score, CSOC Performance, Combinatorial Optimization, Simulation Model, Performance Scoring Model.

1 Introduction

A Cybersecurity Operations Center (CSOC) performs several tasks that include monitoring networks, threat detection, analysis of alert logs, containment and mitigation of threats, incident response and remediation, vulnerability assessment, reporting, compliance with cybersecurity standards, and signature updates for building intelligence about cyber threats [1]. Several types of personnel are employed at a CSOC that include senior level lead, watch officers, first responders that perform triage analysis, incident handlers who analyze incidents and write reports, vulnerability engineer for configuring the vulnerability scanner and scheduling the various vulnerability scans, Information Technology (IT) systems engineer for implementing remediating actions, and subject matter experts such as researchers [2]. Threat signals (alerts) for a CSOC arrive from various sources such as Intrusion Detection Systems (IDS), Security Information and Event Management (SIEM), emails, phone calls, unauthorized access, data breach, and espionage activities by personnel. Alert analysis is a critical and complex task performed by a CSOC. Sensors monitoring different parts of the network generate varying amounts and types of alerts. Also, security personnel have varying characteristics, such as they differ in their expertise and skill levels, among others. The above illustrates that a CSOC environment is highly dynamic, and continuously changing and evolving. In a continuously evolving environment, the personnel at a CSOC collaborate with their team members and function effectively as a team to achieve the primary objective of protecting the organization from cyber threats. Hence, effective team formation is a critical need of the CSOC enterprise.

There is a limited amount of literature on team formation for cyber operations. Current research by behavioral scientists focuses on selecting team members based on their traits and working habits [2], [3], [4]. Malviya et al. [5] and Granasen and Andersson [6] cite effective team bonding as a critical trait for forming successful teams, while Valle et al. [7] and Khanna [8] cite leadership as a necessary trait to handle the time pressure of applying security skills in active cyber threat scenarios. Dawson and Thompson [9] identify team dynamics and social fitting within a cyber operation team as important traits, while Brase et al. [10] identify behavioral ratings and study their effects on patterns of responses. Khanna et al. [11] focus on traits such as leadership, attention to detail, adherence to deadlines, action over reflection, and big picture over details to form cybersecurity teams. Figure 1 shows a schematic of the current approach in team formation for cyber operations. In this approach, first, an individual profile for each of the personnel is created. The individuals answer selected questions through a survey identifying their personality traits. Next, teams are formed with a balanced mix of these traits. The teams are assigned tasks, and their performance is measured based on the quality and timeliness of the results [11], [12]. The teams are reconstituted based on the identified strengths and weaknesses obtained from their psychometric profile

- *A. Shah is with the University of South Florida, Tampa, FL 33620, USA. Email: ankitshah@usf.edu.*
- *R. Ganesan and S. Jajodia are with the Center for Secure Information Systems, George Mason University, Fairfax, VA 22030, USA. Email: {rganesan, jajodia}@gmu.edu.*
- *H. Cam is with the Army Research Laboratory, Adelphi, MD 20783, USA. Email: hasan.cam.civ@mail.mil.*
- *S. Hutchinson is with the Maryland Innovation and Security Institute, Columbia, MD, USA. Email: sehutchinson@misi.tech*

compositions. However, the direct impact of the team's performance on its members is not quantified. As a result, the performance of the individuals on the previous task(s) is not considered when forming new teams.

Analysts also have specific skill-based characteristics, which are critical for cyber operations and need to be directly taken into consideration while forming teams. Analysts have different expertise levels (junior, intermediate, and senior), tooling knowledge (such as SNORT, Suricata, Snorby, Squil, OSSEC, and Bro [13]), and credentials (such as confidential, secret, and top-secret security clearance levels). The sub-optimal grouping of analysts in a team directly impacts the performance of the CSOC. For example, in a scenario from literature [14], it was observed that the throughput of the alert investigation was significantly impacted when proper analyst credentials were not met during alert investigation tasks. In another scenario [11], it was observed that while there was an abundance of a specific skill set in a team, other essential skills needed for the successful operation of the cybersecurity tasks were absent. Hence, it is critical to identify and meet the team requirements by considering the individuals' technological skills and past performances.

Our study focuses on taking a holistic view of the CSOC by defining the team requirements and selecting individuals based on their skills and performances to form effective teams. The goal is to protect an organization from cyber threats by forming balanced teams using a quantitative approach. Each team must maximize the throughput of investigated alerts and minimize false positive and negative decisions. To quantify the team performance and its impact on the team members, we propose a new metric, *collaborative score*. This new performance metric measures the goodness of a team by considering the quality (precision and recall values) and timeliness (throughput rate) of the performed tasks. We provide the related definitions in Section 3.3. The research objective is to form optimal teams of cyber analysts as needed such that 1) all teams have their collaborative score maximized, and 2) all the collaborative scores are balanced among teams (i.e., the maximum difference between the collaborative scores of any two teams is minimized, which means that the variance among the collaborative scores of the teams is minimized). Since it is necessary always to maintain effective teams to maintain performance, it is also desired to have a decision-support framework that assists in re-forming teams when disruptions occur. Teams need to be re-formed to balance the performances in the face of disruptions that may occur due to one or a combination of factors, such as a high rate of alert generation beyond normal levels (which includes adding new sites or activities at the CSOC), personnel gaps, and deterioration in team performance.

An optimal team, also known as an effective team, is said to be formed if its collaborative team score is the highest, which means that the team has the highest possible throughput of analyzed alerts, and lowest possible false positive and negative rates, while keeping the performances among the teams balanced. The premise of the paper is that good collaborative teams lead to best overall performance of the alert analysis division of the CSOC. To our knowledge, based on our discussions with the CSOCs, the current process of team formation with only analysts and their allocation to the alert analysis is not only manual but also is executed at the local level of CSOC analysts. There does not exist an efficient decision-support approach in the literature that guides effective team formation as needed for the CSOC. The individual performance score is determined as follows using the team's collaborative score that has the following metrics. At the end of each shift, throughput metric on the number of alerts analyzed by a team is available. False positive information about an alert is available as soon as the alert is determined to be benign by a secondary check process. False negative information about an alert is a more serious situation and is often detected only when compromise is detected, which could take several shifts of operation. It is assumed that alert analysis can be traced to the team that performed the analysis (based on the time-stamp of alert generation and its categorization). Therefore, as soon as the above information on the team's performance is available, the throughput, false positive, and false negative scores for a team are uniformly divided among the team members, which will update the individual's performance scores on the above metrics. Meeting the research objectives and forming effective teams is nontrivial because of the multitude factors given above. Hence, to meet the research objectives, we present an integrated framework comprising an optimization, simulation, and team scoring model to form effective teams for each shift of operation.

The research in this paper differs from the authors' prior work to optimize the level of operational effectiveness (LOE) of a CSOC [15], dynamic scheduling of analysts [16], and grouping of sensors and allocation to analysts [14] in the following manner. The aforementioned research studies do not form teams, and it is assumed that as long as there is a minimum expertise mix of senior, intermediate, and junior analysts in a team, then such a team is supposed to have a high quality of alert analysis. In other words, this means that not only would such a team have high throughput rate of alert analysis but also have negligible false-positive and negative rates. In this paper, the above assumption is broken, and individual team members are scored based on their performance within a team. This information, as described earlier, is used to form collaborative teams that have high collaborative scores.

There are several contributions of this research. The primary contribution is the integrated framework that delivers the teams for each shift of operation such that the research objective is met. The framework benefits the CSOC that can quickly put together an efficient team based on the inputs such as team requirements, team member attributes, and desired organizational performance for each shift of operation. The collaborative team scoring scheme from which the individual member's scores are determined using the team's throughput, false positive, and false negative rates is another novel aspect of this research. It must be noted that as time progresses, more information on the team's performance becomes available, which can be used to fine tune the formation of the teams. Another unique and novel aspect of the research objective is that not only does the integrated framework seek to maximize the collaborative team score of a team that is formed but also attempts to balance the collaborative score among the teams such

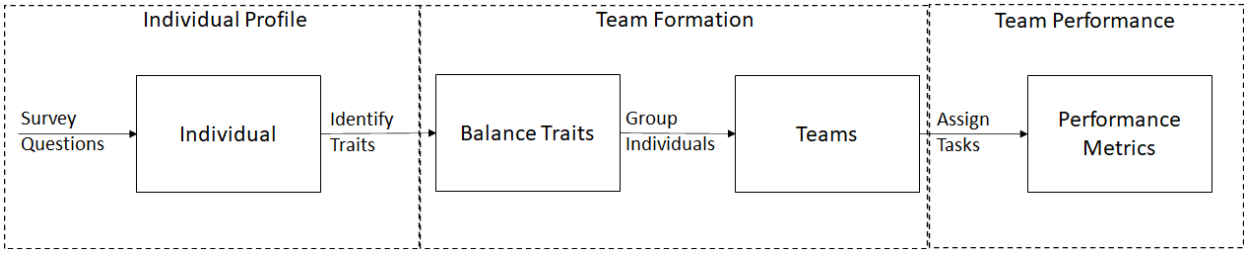

Fig. 1. Schematic of Current Approach in Team Formation

that the difference in the collaborative score between any two teams is minimum. Finally, the proposed framework is also able to quickly identify the outliers (high- and low-performers) with the performance scoring model. The paper also delivers several meta-principles that provide insights and also guide in effective team formation that can be used by practitioners and researchers alike.

The paper is organized as follows. Section 2 provides related literature pertaining to team formation. Section 3 describes the integrated framework for team formation. The optimization, simulation, and scoring models are explained in this section. Section 4 describes the experiments and provides an analysis of results. Section 5 describes the meta-principles obtained from this research study, which is followed by the conclusions and future work.

2 Related Literature

The environment of a CSOC is continuously evolving. Analysts are expected to adapt to this dynamic environment and protect the organization from cyber-threats. Analysts mitigate the threats by collaborating among each other and functioning effectively as a team. Team members share the workload comprising of various tasks and collaborate on difficult tasks [17]. Steinke et al. [4] lists forming effective teams as a critical aspect that determines the success of a CSOC. The effectiveness of a CSOC team depends on the technological and behavioral characteristics of the individual team members. In a recent study, Foley and Rooney in [18] emphasize the role of people and their experiences in the CSOC teams toward achieving optimal functioning. They use grounded theory to qualitatively assess the human experience of working in security environments. The technological characteristics of team members include proficiency with the tools needed for analyzing the tasks, member credentials, and their expertise level. Behavioral characteristics, as identified by Steinke et al. [4], include adaptation, collective problem-solving, communication, building trust, and sharing knowledge of expertise. For instance, collective problem-solving require team members with varying expertise levels so that difficult tasks can be collaborated successfully. Hence, when forming effective teams it is critical to have a mix of expertise levels and proficiency with various tools among the team members. Similarly, psychological safety or trust can be attained by building teams where the team members work well with each other.

The authors in [1] identify the following roles that are consistent in CSOC teams: (1) senior level lead, (2) incident handlers (analysts who analyze incidents and write reports), (3) first responders (that perform triage analysis [19]), and (4) subject matter experts. Killcrese et al. [20] identify the first responders and the incident handlers as the core members of the teams. The other roles are required on a need basis. A team-based approach to computer network defense has been studied by Deckard et al. [21], where the team performance is assessed by post-event surveys. Behavioral scientists have identified key characteristics [2], [3] and strategies to improve the performance of a team. The effectiveness of a team is derived from the individual behavioral characteristics of the team members. Halfhill et al. [22] define team personality composition as an aggregation of personality traits in a team that influences its effectiveness. The authors in [23] quantify and aggregate individual characteristic scores to the team level score to gauge the effectiveness of a team. In particular, they use the mean characteristic score of the team members, the lowest score of the team members, the highest score of the team members, and the variance in the scores among the team members for evaluation. Agreeableness is identified as an important personality trait in composition of effective teams [24]. Agreeable team members value affiliations and are more trustworthy compared to others [25], [26]. As a result, they are perceived to improve team cohesion.

The research so far has focused on the relationship between individual characteristics and the effectiveness of a team. In this paper, we study the characteristics of a team and form effective teams such that the performances of the teams are balanced. The literature is primarily focused on the behavioral characteristics of team members where as in this paper, we focus on technological and performance related factors. In particular, we present a quantitative approach towards forming effective teams.

3 Team Formation and Performance Evaluation Framework

Team formation and performance evaluation is a dynamic exercise because teams can change in 1) the immediate horizon (per shift) based on available personnel in a shift (accounts for absenteeism), 2) the short-term (every 14-day scheduling cycle) based on the analyst performances, and 3) the long-term as new analysts are added or existing ones leave, analyst tool knowledge is enhanced, and analyst credentials are updated. Figure 2 provides the framework for team formation and performance evaluation. It consists of three models, which are executed in a sequential manner and iterated over many scheduling cycles. The framework

aims to help the decision-makers at the CSOC by forming optimal teams, whose collaborative scores are maximized and balanced, that minimizes the risk to the organization posed by sub-optimal team(s). Upon repeating the process over several scheduling cycles, the framework aims to also help identify outlier analyst performances (high and low). Next, we explain each of these models in detail.

3.1 Optimization Model

The optimization model determines the optimal teams of analysts with maximized and balanced collaborative scores by taking into account several inputs and subject to several constraints. We first explain the inputs to the model, followed by the team requirements constraints.

Analysts have various attributes, including different expertise levels, proficiency with different tools, credentials, and performance scores acquired from their past teams' performances (explained in Section 3.3). Analysts have access to all the tools that are used to analyze alerts; however, the expertise on these tools vary. For instance, a junior analyst will have expertise on a few tools (such as SNORT), while a senior analyst will be an expert on all the tools in the entire set. Alerts coming from specific sensors must require credentials such as security clearances due to the sensitive nature of the data. A matrix of analysts and their attributes, which include expertise level, tooling, credentials, throughput score, precision score, recall score, and average time to analyze an alert are provided as inputs to the model. It is assumed that the CSOC is adequately staffed such that there are appropriate numbers and types of analysts available at the start of the shift to cover all the alerts generated under normal operating conditions. The team requirements gathered from literature and our conversations with the CSOC managers, which are provided as constraints that the team formation model must meet, are as follows. The mix of analyst skill (expertise) levels must be maintained for each team [16], [27]. Each team's credentials and tooling needs must be met [14]. Also, each team must have the required number of analysts and each analyst must be assigned to a team.

The output of the optimization model is the teams and their respective composite collaborative scores, which are maximized and balanced. The research objective of the optimization model is to form teams and meet the team requirements, maximize throughput, and minimize false positives and negatives so that in the long run optimal teams with highest collaborative scores are formed and the collaborative scores among teams are balanced. The above distinction between highest and balanced is justified as follows.

Optimization for team formation is executed before simulating the work-shift with alerts. One of the inputs to optimization is each analyst's individual throughput, false positive and negative scores. When teams are formed with analysts who report to work, the sum total of all their throughput and false positive and negative scores is a constant, regardless of how teams are formed. For example an ad hoc team formation method could result in a team with all high performers (high collaborative score) and another team with all low performers (low collaborative score), but the total throughput will still remain the same as that of two teams with equal collaborative score. This is because of a closed system, which cannot improve unless the individual analyst's performance improves. Since a team with all low performers (low collaborative score) is undesirable as the risk associated with incorrect alert categorization and low throughput will be very high with this team, the goal is to create teams of near-equal (or balanced) performances. Hence, the objective of the optimization model is to form teams such that the maximum difference among the collaborative team scores (say, d) is minimized while meeting the team formation and performance requirement constraints. The frequency of running the optimization model is every 14 days, although the conditions mentioned earlier in this section can also trigger the need for running the model to make readjustments. Next, the work-shift is simulated by first generating the alerts, followed by the process of analyzing them by the formed teams of analysts. The simulation details are explained in the next section.

TABLE 1
Definitions of Notations

Notation	Definition
Indices	
i	Analyst identity
j	Team identity
k	Skill level identity
t	Tool identity
Inputs	
I	Total number of analysts available
J	Total number of teams
K	Total number of skill levels
T	Total number of tools
S_i	Performance score of analyst i
M_j	Total number of analysts required per team j
$N_{j,k}$	Minimum number of analysts required per team j from skill level k
$O_{j,t}$	Minimum number of analysts required per team j with tool t
$Z_{i,k}$	1 if analyst i belongs to skill level k, and 0 otherwise
$U_{i,t}$	1 if analyst i is expert on tool t, and 0 otherwise
Variables	
$x_{i,j}$ (Binary)	1 if analyst i is assigned to team j, and 0 otherwise
d	Maximum difference between the collaborative scores of any two teams

3.1.1 Mathematical formulation

The model formulation consists of the objective function, constraints, and outputs, which are explained in detail below. The notations for the parameters of the exact optimization model are described in Table 1.

The constraints for the optimization model are as follows:

1) Each team must have the required number of analysts, which is given by

$$\sum_i x_{i,j} = M_j \quad \forall j. \tag{1}$$

2) Each team must meet the mix of analyst skill levels, which is given by

$$\sum_i x_{i,j} * Z_{i,k} \geq N_{j,k} \quad \forall j, k. \tag{2}$$

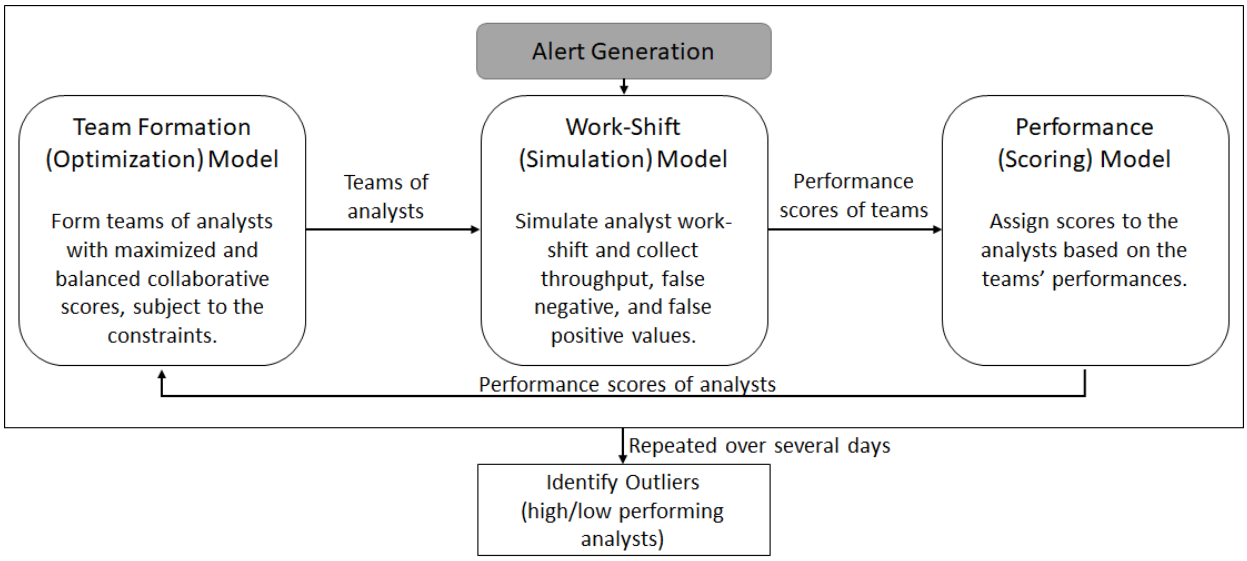

Fig. 2. Team Formation and Performance Evaluation Framework

3) Each team must meet the credentials and tooling needs, which is given by

$$\sum_i x_{i,j} * U_{i,t} \geq O_{j,t} \ \ \forall j,t. \tag{3}$$

4) Each analyst must be assigned to a team, which is given by

$$\sum_j x_{i,j} = 1 \ \ \forall i. \tag{4}$$

The objective of the optimization model is to balance the collaborative scores among teams (i.e., the maximum difference between the collaborative scores of any two teams, d, is minimized). The maximum difference in collaborative scores among any two teams is measured by

$$\sum_i x_{i,j} * S_i \leq d \ \ \forall j. \tag{5}$$

The objective function equation is as follows:

$$w = Min \ \ d. \tag{6}$$

The performance scores of the analysts ($S_i, \forall i$) are calculated by the scoring model (explained later in Section 3.3) and provided as an input to the optimization model. The outputs from the optimization model are (1) teams of analysts, i.e., $x_{i,j} \ \ \forall i,j$, and (2) collaborative scores of all the teams, i.e., $\sum_i x_{i,j} * S_i \ \ \forall j$.

Algorithm 1: Optimization Algorithm for Team Formation.

Input: Total number of analysts available I, total number of teams J, total number of skill levels K, performance score of each analyst S_i, total number of analysts required per team M_j, minimum number of analysts required per team from each skill level $N_{j,k}$, minimum number of analysts required per team j with tool t $O_{j,t}$, and skill levels and tool expertise of all the analysts ($U_{i,t}$ and $Z_{i,k}$).
Output: Analysts assignments to teams, $x_{i,j} \ \ \forall i,j$; collaborative scores of the teams, $\sum_i x_{i,j} * S_i \ \ \forall j$.
/*Initiate a solution search using an integer programming solver */
repeat
  for *a set of* $x_{i,j} = 1$, */*Potential solution obtained in a search*/ Check for feasibility:* **do**
    $\sum_i x_{i,j} = M_j \ \ \forall j$ /* equation 1 */
    $\sum_i x_{i,j} * Z_{i,k} \geq N_{j,k} \ \ \forall j,k$ /* equation 2 */
    $\sum_i x_{i,j} * U_{i,t} \geq O_{j,t} \ \ \forall j,t$ /* equation 3 */
    $\sum_j x_{i,j} = 1 \ \ \forall i$ /* equation 4 */
    $\sum_i x_{i,j} * S_i \leq d \ \ \forall j$ /* equation 5 */
  end
  if *Feasible* **then**
    $w = Min \ \ d$ /*Min score difference*/
  end
until *Stopping criteria /*Optimal value for d is found*/*;
return $x_{i,j} \ \forall i,j$; $\sum_i x_{i,j} * S_i \ \ \forall j$.

3.1.2 Algorithm for the Optimization Model

Algorithm 1 provides the implementable steps for the optimization model described above.

3.1.3 Computational Complexity

The decision problem presented in this paper of assigning each member to a team is similar to a political districting problem, where districts are formed by partitioning a territory. The political districting problem is known to be NP-complete [28]. The research problem is as hard as that problem. The complexity of the algorithm is of the order $O(2^{I*J})$. However, due to a smaller number of teams to be formed in a CSOC and constraints in forming the teams, the number of potential solutions are decreased. The instances shown in the experiments have a computational run time of less than five minutes. The algorithm was developed in

Python and solved using the Gurobi solver [29]. It is to be noted that the teams are formed biweekly and there is ample time between each decision to allow for more complex data (in terms of number of teams and organization-specific constraints) at a CSOC.

3.1.4 Advantages of Optimization Model

1) It allows for the constraints to be met via the mathematical programming model, which could be difficult to achieve manually via ad hoc (experience-based) team formation.
2) It gives data-driven insights/recommendations to a decision-maker and thereby helps in reducing human error.
3) It helps in balancing the collaborative score, which ensures that there are no teams with very low collaborative score (high risk).

3.2 Work-Shift (Simulation) Model

In real-world conditions, the output of the optimization model, i.e., the analyst teams will be recommended to the CSOC manager and schedules with the new teams will be generated. To mimic the real-world CSOC operations, we develop a simulated environment. A work-shift is simulated using a discrete event simulation model as shown in Figure 2. Each simulation run corresponds to a 14-day time period, which further consists of two shifts of 12 hours each per day, and alerts are generated from various sensors in a network using a Markovian distribution (Poisson arrival process) [16]. The teams of analysts formed as a result of the optimization model are used as inputs in the simulation model, which are scheduled throughout the 14-day time period [16]. Analysts have attributes such as the throughput, false positive and false negative rates. First, sensors are grouped together to form clusters [14]. This process helps with balancing the alert workload and analyst credential requirements. For each shift of operation, alerts are enqueued based on their generation time. The alert analysis process is then simulated. The analysts analyze alerts based on the CSOC's picking strategy. In this study, we consider the selection of alerts that are oldest in the queue. If the alerts are prioritized based on their criticality level, then they can select alerts that are not only the oldest but also have the highest criticality level. Finally, the total number of alerts analyzed, false positives, and false negatives are reported for each of the teams.

3.2.1 Algorithm for the Work-Shift (Simulation) Model

Algorithm 2 describes the steps in the simulation model for the work-shift.

Algorithm 2: Work-Shift Simulator Algorithm.

Input: Average number of alerts generated per sensor, analyst teams from optimization model, number of analysts and their attributes: tooling/credentials, throughput rate, false positive rate, and false negative rate.
Output: Number of alerts analyzed, number of false positives and false negatives.
repeat
  Step 1: Simulate a work-shift;
  Create groups of sensors to form clusters;
  Create a queue to collect alerts generated from sensors;
  Create teams of analysts (obtained from the optimization model);
  $timeindex$ = 0;
  repeat
    Step 2: Generate alerts per sensor using a probability distribution with the given average alert generation rate;
    Step 3: Enqueue alerts generated from the sensors based on their arrival times;
    Step 4: Alerts are picked up for alert analysis process by idle analysts using the CSOC's alert picking strategy;
    Alerts are analyzed by the analysts by picking a random value within their range of throughput rates;
    Alerts are marked as false positives and false negatives based on the respective analyst's false positive and negative rates by picking a random value within their ranges;
    $timeindex$ ++ $(1\ hr\ time\ steps)$;
  until $timeindex$ = *12 hrs*2 shifts*14 days*;
  Report the analyzed, false positive and false negative alerts;
until *the number of simulation replications have been completed*;
return *Number of alerts analyzed, number of false positives and false negatives.*

3.3 Performance (Scoring) Model

The performance of the CSOC is measured in many ways [12]. In this paper, we consider the following metrics to measure the performance of the alert analysis division.

1) Throughput of a team: It measures the number of alerts investigated by a team. It is also a direct function of 1) time elapsed between detected time of potential attack/compromise (alert generation time) and report writing (time at which the alert has completed investigation), 2) the type of sensor (alert type), and 3) tooling availability and credentials of the team investigating the alert.
2) False positives of a team: It measures the number of alerts that were incorrectly identified as attacks. False positives result in wastage of team resource by focusing their time on innocuous alerts. This in turn reduces the throughput. Precision is defined as follows:

$$Precision = \frac{TruePositives}{(TruePositives + FalsePositives)} \tag{7}$$

3) False negatives of a team: It measures the number of alerts that were incorrectly identified as benign. False negatives result in real harm to the organiza-

tion. Recall is defined as follows:

$$Recall = \frac{TruePositives}{(TruePositives + FalseNegatives)} \tag{8}$$

4) Collaborative score of a team: It measures the goodness of a team. This proposed metric is a composite score of three metrics that captures the combined throughput of all the individual members on a team, their combined precision, and their combined recall scores. The collaborative score will follow a value model wherein the value of normal or a predetermined throughput (best value) is 1 on a 0-1 scale, and the value of precision and recall is already expressed on a 0-1 scale where 1 is the best value. The collaborative score C_j for team j is then a weighted value of throughput, precision, and recall, and is calculated for every shift in the day.

$$C_j = W_T V(T_j) + W_P V(P_j) + W_R V(R_j), \tag{9}$$

where $V()$ is the value function and W_T, W_P, and W_R are the weights of throughput T, precision P, and recall R, respectively. The weights provide the flexibility for the CSOC to form teams by prioritizing either of the three performance metrics, throughput, precision, or recall, based on the problem type and importance of consequences.
The variance among C_j for J total number of teams formed is given by

$$Var(C_j) = \frac{\sum_{j=1}^{J}(C_j - \bar{C})^2}{J-1}, \tag{10}$$

where $\bar{C}$ is the average collaborative score of all teams. $Var(C_j)$ is minimized, as explained later in the optimization model (in Section 3.1). The overall performance of the organization can be defined as the sum of the collaborative scores of the J number of individual teams, which can be calculated as

$$Total\ Collaborative\ Score = \sum_{j=1}^{J} C_j. \tag{11}$$

Based on the output of the work-shifts observed over each 14-day period, i.e., the biweekly performances of the teams, the scoring model assigns individual scores to the analysts. First, the throughput, precision, and recall values are obtained for each team and their respective collaborative scores are calculated using Equation 9. Next, these values are equally distributed among the members of each team and their respective performance scores are updated by taking an average over the n number of biweekly periods. The equation for calculating the performance score, S_i, of an analyst i from team j is given by

$$S_i = \frac{S_i * (n-1) + C_j}{n} \quad \forall i, j. \tag{12}$$

These scores ($S_i, \forall i$) are then used in the optimization model (Equation 5) to form teams.

3.3.1 Performance Evaluation

The steps in this framework, as shown in Figure 2, are repeated over many iterations. The CSOC environment is simulated to observe the performances of the teams and the individuals. The team performances are evaluated biweekly, and the analyst scores are updated accordingly. The framework then reports outliers, i.e., high (top) or low (weak) performers, if any, over the long run.

For the overall performance metric, the collaborative team scores are measured and the maximum difference between the collaborative scores of any two teams is reported. In an ideal situation, real-world data must be used to measure aggregated throughput, and false positives and negatives both with and without efficient team formation over the same long period of time. There is a lack of such data due to confidentiality reasons and the research would prove the hypothesis that by balancing the collaborative scores of the teams formed via optimization, the overall risk associated with low performing teams (with lower throughput and higher rates of false negative and positive decisions) is minimized. Next, we present the experiments and the analysis of results.

4 Experiments and Analysis of Results

In this section, we first describe the setup for conducting experiments, followed by the analysis of results.

4.1 Experimental Setup

We have been actively studying the problem of alert management and have had many discussions with managers from various CSOCs (government, public, and private), which include in-house and third-party service providers. Our design decisions were based on these conversations and observations. We consider three cases in our experiments, which were obtained from our discussions with them. These include one nominal case and two outlier cases, which are explained as follows.

4.1.1 Nominal Case

We consider a CSOC with 40 analysts and four teams. There are 12 senior analysts, 14 intermediate analysts, and 14 junior analysts. The analyst expertise mix required in a team is a minimum of 20% seniors, 30% intermediates, and 30% juniors [27]. The number of members required in each team is 10. All the junior analysts have access to basic tools (such as SNORT), intermediate analysts have access to more advanced tools, and seniors have access to the entire tool set required at a CSOC [14]. An analyst expertise level is highly correlated with the throughput and false positive and negative rates in cybersecurity tasks. For example, a senior level analyst is expected to have the highest throughput and the lowest false positive and negative rates. Table 2 shows the analyst attribute ranges for the throughput, false positive, and false negative rates used for the experiments, which were obtained from our conversations with the CSOCs and the literature survey [16], [27], [30].

Based on these conversations, we considered equal weights ($W_T = W_P = W_R = 1$) for the performance metrics observed at a CSOC in the experiments shown

This article has been accepted for publication in IEEE Transactions on Services Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSC.2023.3253307

TABLE 2
Input: Analyst Attributes

Attributes (Range)	Senior	Intermediate	Junior
Analysis rate/hour	300-420 alerts	200-280 alerts	160-200 alerts
False positive rate	4-6%	6-8%	8-10%
False negative rate	0.04-0.06%	0.06-0.08%	0.08-0.10%

in this section. These weights are used to calculate the collaborative score, C_j, of each team (in Equation 9). The value for $V(T_j)$ is calculated on a 0-1 scale as follows. The highest throughput rate possible in a 14-day period is determined for each team, which will be obtained if all the analysts in the team analyzed the number of alerts (or more) expected from them. This number takes a value of 1. A value of 0 is assigned if no alerts are analyzed by the team. The actual value of $V(T_j)$ at the end of the 14-day period is obtained by exponentially normalizing between these two numbers. The value of $V(P_j)$ and $V(R_j)$ are obtained based on equations 7 and 8, respectively, which are also described in the introduction section.

We also set up a baseline case to compare against our methodology. In this baseline case, we use an ad hoc approach to form teams without the optimization model, which is explained as follows. We measure each team's attribute values at the end of the biweekly period. If this value does not meet a predetermined threshold, then we select the lowest performing analyst(s) from the respective team(s) and swap them with the average performer(s) of the same skill level from the top performing team(s). Note that we do not swap the lowest performing member with the best performer from another team, as that will create even more imbalanced teams. We selected the mean values of the ranges shown in Table 2 (i.e., throughput of 240 alerts, false positive rate of 7%, and false negative rate of 0.07%) as the respective thresholds.

Next, we developed outlier cases for the experiments to evaluate our methodology. These cases were developed based on our conversations with the CSOC managers. Through these conversations, we found multiple cases in which it becomes difficult for a CSOC manager to attain balanced teams, including cases with (a) a minimal number of analysts and (b) many analysts with outlier characteristics (under and over performing). We considered various cases in our experiments and describe in detail two of these cases to demonstrate the effectiveness of our approach. One case considers only one outlier analyst, while the other considers multiple outlier analysts from individual skill levels.

4.1.2 Outlier Case 1

The performances of two analysts are changed in this case. 1) The attributes of analyst ID 3 (a senior analyst) are upgraded by 20% (i.e., analyst's throughput, true positive and true negative rates are upgraded by 20% from the numbers listed in Table 2). 2) The attributes of analyst ID 38 (a junior analyst) are downgraded by 30%. All the other inputs are used from the nominal case, as shown in Table 2.

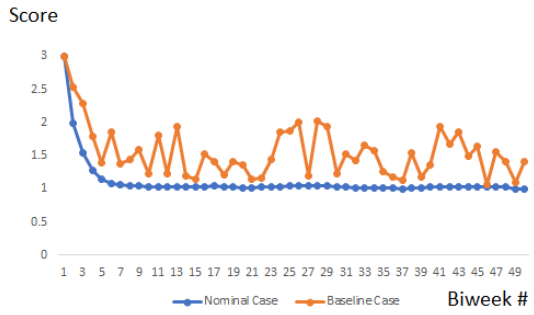

Fig. 3. Convergence in Nominal Case

4.1.3 Outlier Case 2

We setup another outlier case, in which there are multiple outliers considered in a CSOC. The attributes of analyst IDs 3, 6 (senior analysts) and 16 (an intermediate analyst) are upgraded by 30%. Whereas, the attributes of analyst IDs 18 (an intermediate analyst) and 34 (a junior analyst) are downgraded by 30%. All the other inputs are used from the nominal case, as shown in Table 2.

In addition to the above case studies, we also fluctuated the alert arrival rates by randomly generating adverse events (Shah et al. [31]), to observe the impact on team formation. Each experiment is run over 26 iterations (equivalent to 52 weeks) to observe the characteristics of the teams and the individual members. Next, the results are analyzed and presented.

4.2 Analysis of Results

Below we analyze the results obtained from the experiments conducted for the three cases.

4.2.1 Nominal Case

We ran this experiment by starting with the least optimal team configurations in the first iteration. We selected members for each team to maximize the difference among the collaborative team scores in biweek #1. We wanted to determine the maximum time it would take for teams to balance their collaborative scores using our methodology. In Figure 3, we plot the maximum difference among the collaborative team scores (d) over a longer (two years) period to observe any fluctuations. It can be seen that this value converges around the eighth biweekly period. We ran various experiments with different random seeds and noticed similar results. We observed that it takes up to four months for a CSOC to establish teams, where the maximum difference in collaborative scores converges, when starting from sub-optimal configurations. These results are due to the sequential (biweekly) application of the optimization algorithm in our framework and the correlation between the analysts' attributes (throughput, false positive, and false negative values) and their expertise levels (senior, intermediate, and junior), which exists in cybersecurity tasks.

It is to be noted that without optimization, the team formation is ad hoc and there will be fluctuations in the difference in collaborative scores among the teams and the value of this difference will be larger. Figure 3 shows the

This article has been accepted for publication in IEEE Transactions on Services Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSC.2023.3253307

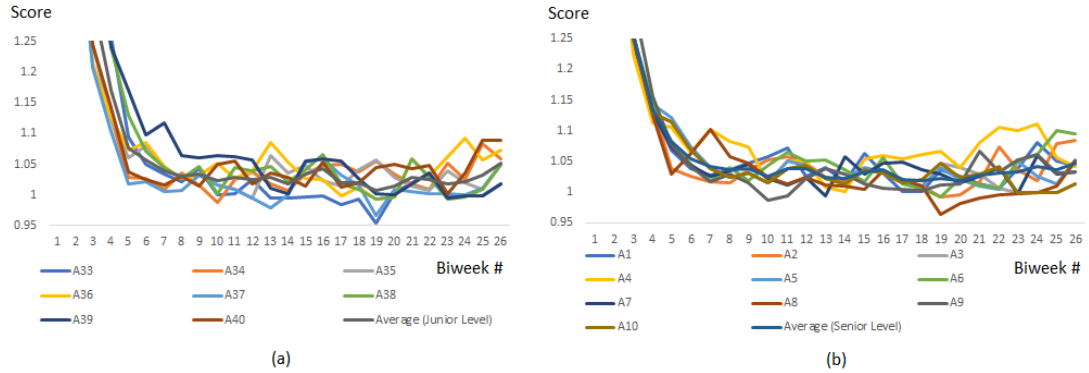

Fig. 4. Individual Performances in Nominal Case: (a) Junior Analysts and (b) Senior Analysts

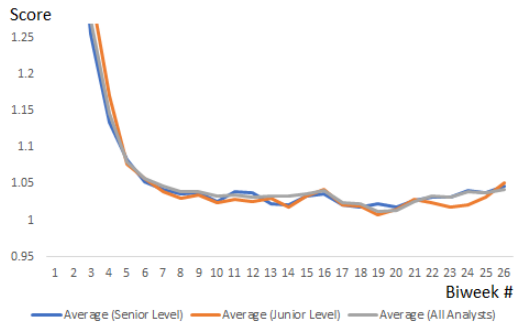

Fig. 5. Average Analyst Performances in Nominal Case

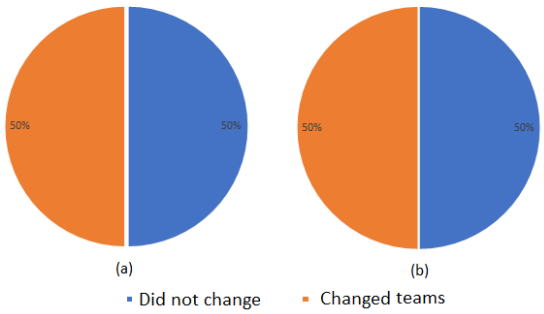

Fig. 6. % of Times an Analyst Changed Teams in Nominal Case: (a) Top and (b) Weakest Performers

fluctuation using the baseline case for this experiment. This indicates that without optimization there will be teams with lower collaborative scores than the average observed among the teams, creating a high-risk scenario for the CSOC.

Figure 4 (a) shows the performance of the junior level analysts and Figure 4 (b) shows the performance of the senior level analysts. We have modified (magnified) the scale on the y-axis to show the temporal patterns in the analyst performance scores. In a nominal case where the analysts are performing according to their expected attribute values, there is no significant difference observed among their individual scores over the 52 weeks. Similar observations were made for the intermediate level analysts. Furthermore, Figure 5 shows that the average performance score of analysts is similar across all the expertise levels. This is because the teams' performances are nearly balanced from the sixth biweekly period (see Figure 3), and the individual scores of the analysts are calculated from their respective teams' scores. It is also observed that in a nominal case, the % of times an analyst changed teams was the same (50%) among all the performers. Figure 6 shows the % change for the top and the weakest performers. In the baseline case, the top performers did not change teams, and only the weakest performers changed teams.

4.2.2 *Outlier Case 1*

This is the case in which two outliers (analysts), in terms of their performances, were considered in a CSOC. It is to be noted that this information was not directly known to the optimization model. Figure 7 shows that the maximum

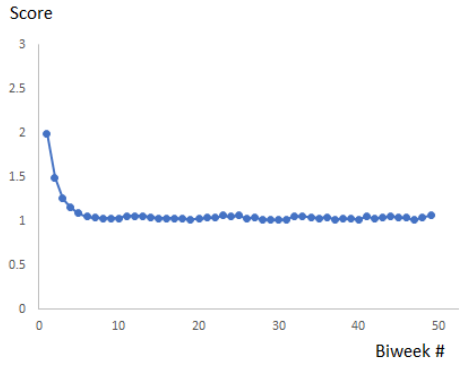

Fig. 7. Convergence in Outlier Case 1

difference among the team scores converges, approximately, around the four-month period, which was also observed in the nominal case. Through various experiments, we observed that it takes around four months for a CSOC to establish teams with the difference among their collaborative scores optimally minimized (converged), irrespective of outliers being present among the team members.

Figure 8 (a) shows the performance scores of a sample set of senior analysts. It is easy to spot the top performer using this visualization. It can also be noted from Figure 8 (b) that a top performer can be identified very quickly, within the first two months. Similarly, Figure 9 (a) shows the performance scores of a sample set of junior analysts. It could be seen that it may take about three to four months

This article has been accepted for publication in IEEE Transactions on Services Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TSC.2023.3253307

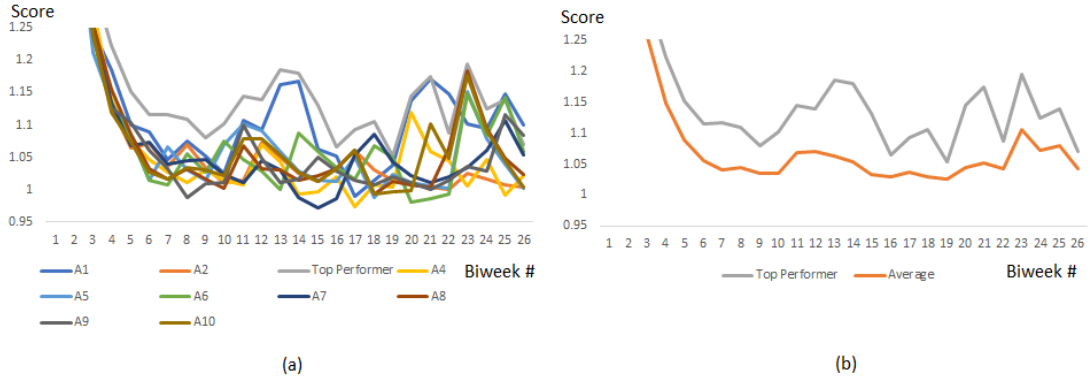

Fig. 8. Individual Performances in Outlier Case 1: (a) Senior Analysts and (b) Top vs Average Performers

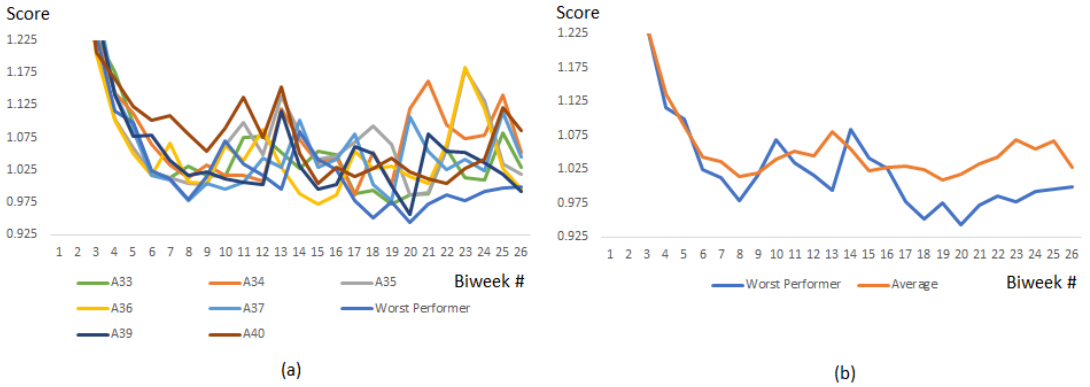

Fig. 9. Individual Performances in Outlier Case 1: (a) Junior Analysts and (b) Weak vs Average Performers

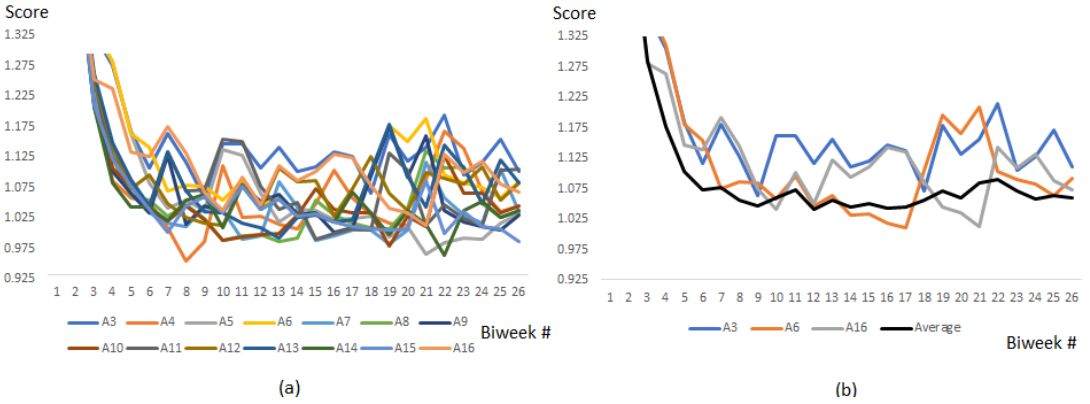

Fig. 10. Performances in Outlier Case 2: (a) Senior and Intermediate Analysts and (b) Top vs Average Performers

to clearly identify an outlier in terms of the weakest performance (see Figure 9 (b)). Figure 12 shows the performance of the outliers compared to the average performance of all the team members. It is to be noted that such a visualization of the analyst scores can help the CSOC manager effectively identify the outliers in the teams. Figure 13 shows the % of times an analyst changed teams. It can be seen that the top performer changed teams most frequently compared to all the other members. It is also noted that the outliers are often assigned to different teams when compared with an average team member. This is the reason why the scores of weak performers fluctuate in the first few months before significantly deviating from the average score of all the analysts (see biweek #16 onward in Figure 9 (b)).

4.2.3 *Outlier Case 2*

In this case, more outliers, compared to the previous case, were added to the CSOC work-shift simulator. Figure 14 shows that the maximum difference among the team scores converges around the similar four-month period, as observed in the other two cases. We observed similar results for convergence in other experiments performed with varying number of outliers. Figure 10 (a) shows the performance scores of a sample set of senior and intermediate analysts.

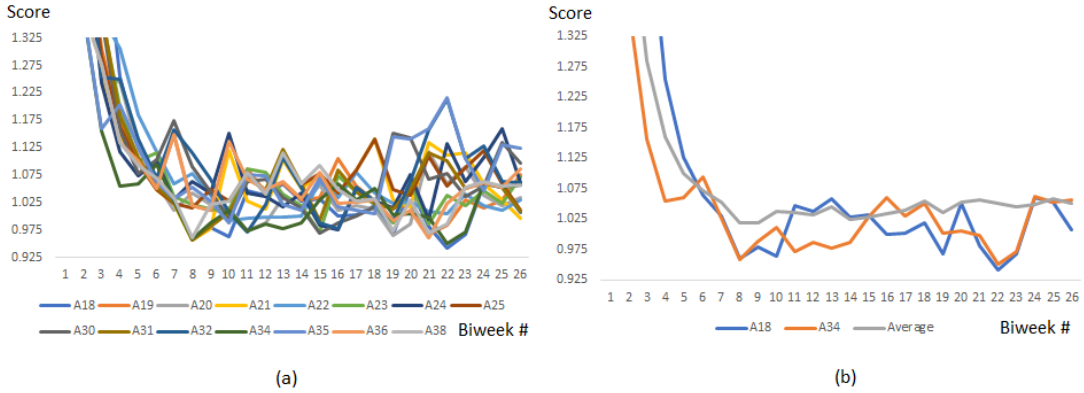

Fig. 11. Performances in Outlier Case 2: (a) Junior and Intermediate Analysts and (b) Weak vs Average Performers

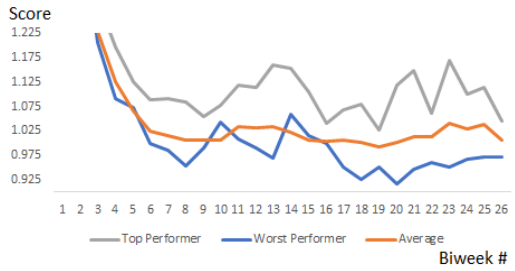

Fig. 12. Top vs Weakest Analyst Performances in Outlier Case 1

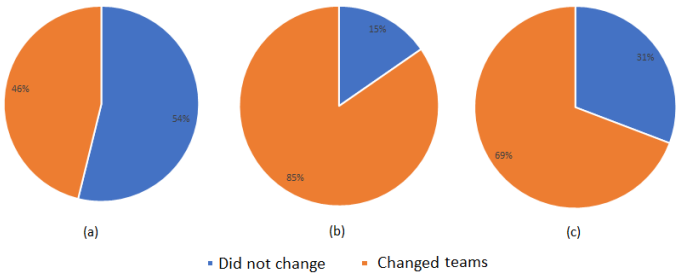

Fig. 13. % of Times an Analyst Changed Teams in Outlier Case 1: (a) Average, (b) Top, and (c) Weakest Performers

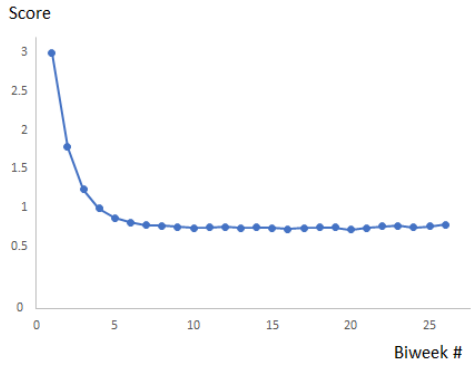

Fig. 14. Convergence in Outlier Case 2

It is easy to spot the top performers using the visualization shown in Figure 10 (b), which uses the average performance score of the team members to identify the top performers. Figure 11 (a) shows the performance scores of a sample set of junior and intermediate analysts. Similar to the earlier figure for top performers (Figure 10 (b)), it is easy to spot the weak performers using the visualization shown in Figure 11 (b). In this case with more outliers, it takes under four months to identify the weak performers. We noted that % of times various analysts (average, top, and weakest performers) changed teams were identical to those observed in the results for the outlier case 1 (Figure 13). The top performer changed teams most frequently compared to all the other members and the outliers were assigned to a larger number of (different) teams when compared to an average team member.

We had conducted several other experiments, with different weighing schemes for the performance metric weights in Equation 9. We considered a higher weight for each of them compared to the other two in different sets of experiments. We obtained similar results to those described in this section, where effective teams were formed that met the CSOC requirements and performance goals, and converged the collaborative score difference among them. Next, we present our conclusions.

5 Conclusions, Summary of Meta-Principles, and Future Work

In this paper, we investigated a quantitative approach using a novel amalgamation of optimization, simulation, and performance scoring models to form effective teams, which are able to maximize throughput and minimize the false positive and negative rates in the alert analysis process at a CSOC. Our team formation and performance evaluation framework can optimally balance the teams' performances, measured using a new metric that calculates the collaborative effort of the team members to achieve the objectives of the CSOC. Any CSOC can adopt our mathematical formulation for real-world implementation and set up the organization-specific input parameter values to form effective teams. We developed the simulator to mimic the real-world CSOC operations based on published literature and our communication with the CSOCs. We demonstrated

that our framework could quickly form effective teams and identify high and low-performing team members.

Below we present a summary of meta-principles obtained from the experiments conducted in this research:

1) Our quantitative approach that integrates optimization, simulation, and scoring models, along with the proposed metric that measures the effectiveness of the teams, helped form effective teams that were able to maximize and balance their performances.
2) Our proposed team formation and performance evaluation framework is able to form teams such that the difference among their collaborative scores is optimally minimized and reaches convergence in no more than four months, even when starting from sub-optimal configurations. We investigated our approach under varying conditions, including the presence of outliers (causing deterioration in team performances) and fluctuations in alert generation rate (indicative of the addition of new sites or activities at the CSOC).
3) The outliers were quickly identified with our framework in the simulated experiments. The top performers were identified within the first two months, and the weakest performers were identified within the first four months.
4) It was interesting to observe that the top-performing analysts were selected to change teams more frequently compared to the average-performing analysts to balance the collaborative scores and thereby minimize the risk of low-performing teams.
5) In a team environment, identifying low performing members is often challenging. We observed that, in general, team members with below-average performances are assigned to a greater number of teams over a longer time compared to the others.
6) A temporal visualization of individual analyst scores compared against the average score of all the analysts can help a CSOC manager effectively identify the outliers. In the case of low-performing team members, such visualization or dashboard can help identify the opportunities for intervention in the form of specific training programs to improve their performances.

As part of the future work, many improvements can be made over time, influencing the analyst performance scores. These improvements include using new signatures for exploits, employing better network analysis tools, and finding correlations among alerts from various sensors that may reduce the false positives and negatives. Our focus was on forming effective teams in this research study, and we did not consider the impact of these factors on the performance of the CSOC. However, this will be an interesting future direction for the cybersecurity research community. Another future research direction may include determining optimal interventions through training programs to help security analysts maintain the level of operational effectiveness of the CSOC.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Jennifer Cowley who works for the Chief Digital and AI Office at the US Defense Information Systems Agency (DISA) for many helpful discussions. This work was partially supported by ARO grant W911NF-13-1-0421, by ONR grants N00014-20-1-2407 and N00014-18-1-2670, and by NSF grant CNS-1822094.

REFERENCES

[1] G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek, "State of the practice of computer security incident response teams (csirts)," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU/SEI-2003-TR-001, 2003.

[2] S. J. Zaccaro, "Trait-based perspectives of leadership," *American Psychologist*, vol. 62, no. 1, p. 6, 2007.

[3] S. T. Bell, "Deep-level composition variables as predictors of team performance: A meta-analysis." *Journal of applied psychology*, vol. 92, no. 3, p. 595, 2007.

[4] J. Steinke, B. Bolunmez, L. Fletcher, V. Wang, A. J. Tomassetti, K. M. Repchick, S. J. Zaccaro, R. S. Dalal, and L. E. Tetrick, "Improving cybersecurity incident response team effectiveness using teams-based research," *IEEE Security & Privacy*, vol. 13, no. 4, pp. 20–29, 2015.

[5] A. Malviya, G. A. Fink, L. Sego, and B. Endicott-Popovsky, "Situational awareness as a measure of performance in cyber security collaborative work," in *2011 Eighth International Conference on Information Technology: New Generations*. IEEE, 2011, pp. 937–942.

[6] M. Granåsen and D. Andersson, "Measuring team effectiveness in cyber-defense exercises: a cross-disciplinary case study," *Cognition, Technology & Work*, vol. 18, no. 1, pp. 121–143, 2016.

[7] A. M. Valle, P. Khanna, and S. Prabhu, "Comprehensive learning incorporating ako–a tertiary education approach at wintec," in *Proceedings of ITx 2018 New Zealand's conference of IT*, 2018.

[8] P. Khanna, "Cognitive education framework for cyber security: A collaborative community approach aligning to tenets of ako," 2019 ITP Research symposium, Napier, New Zealand, Tech. Rep., 2019. [Online]. Available: http://researcharchive.wintec.ac.nz/6832/

[9] J. Dawson and R. Thomson, "The future cybersecurity workforce: going beyond technical skills for successful cyber performance," *Frontiers in psychology*, vol. 9, p. 744, 2018.

[10] G. L. Brase, E. Y. Vasserman, and W. Hsu, "Do different mental models influence cybersecurity behavior? evaluations via statistical reasoning performance," *Frontiers in psychology*, vol. 8, p. 1929, 2017.

[11] P. Khanna, D. Abuaiadah, and C. Baker, "Forming team for cybersecurity and cyber-forensics operations using individual profiling," in *Proceedings of the 12th Annual CITRENZ Conference*, 2021, pp. 49–65.

[12] C. Zimmerman, *The strategies of a world-class cybersecurity operations center*. McLean, VA: The MITRE Corporation, 2014.

[13] H. Cam, M. Ljungberg, A. Oniha, and A. Schulz, "Dynamic analytics-driven assessment of vulnerabilities and exploitation," in *Big Data Analytics in Cybersecurity*. Auerbach Publications, 2017, pp. 53–80.

[14] A. Shah, R. Ganesan, S. Jajodia, and H. Cam, "Optimal assignment of sensors to analysts in a cybersecurity operations center," *IEEE Systems Journal*, vol. 13, no. 1, pp. 1060–1071, 2018.

[15] ——, "Dynamic optimization of the level of operational effectiveness of a csoc under adverse conditions," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 9, no. 5, pp. 1–20, 2018.

[16] R. Ganesan, S. Jajodia, A. Shah, and H. Cam, "Dynamic scheduling of cybersecurity analysts for minimizing risk using reinforcement learning," *ACM Transactions on Intelligent Systems and Technology*, vol. 8, no. 1, pp. 1–21, Jul. 2016. [Online]. Available: http://doi.acm.org/10.1145/2882969

[17] J. R. Goodall, W. G. Lutters, and A. Komlodi, "I know my network: collaboration and expertise in intrusion detection," in *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, 2004, pp. 342–345.

[18] S. N. Foley and V. Rooney, "A grounded theory approach to security policy elicitation," *Information & Computer Security*, 2018.

[19] A. D'Amico and K. Whitley, "The real work of computer network defense analysts," in *VizSEC 2007: Proceedings of the Workshop on Visualization for Computer Security*, 2008.
[20] G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek, "State of the practice of computer security incident response teams (csirts)," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU/SEI-2003-TR-001, 2003, table 9, page 66.
[21] G. M. Deckard and L. J. Camp, "Measuring efficacy of a classroom training week for a cybersecurity training exercise," in *Technologies for Homeland Security (HST), 2016 IEEE Symposium on*. IEEE, 2016, pp. 1–6.
[22] T. Halfhill, E. Sundstrom, J. Lahner, W. Calderone, and T. M. Nielsen, "Group personality composition and group effectiveness: An integrative review of empirical research," *Small group research*, vol. 36, no. 1, pp. 83–105, 2005.
[23] M. S. Prewett, A. A. Walvoord, F. R. Stilson, M. E. Rossi, and M. T. Brannick, "The team personality–team performance relationship revisited: The impact of criterion choice, pattern of workflow, and method of aggregation," *Human Performance*, vol. 22, no. 4, pp. 273–296, 2009.
[24] J. E. Hoch and J. H. Dulebohn, "Team personality composition, emergent leadership and shared leadership in virtual teams: A theoretical framework," *Human Resource Management Review*, vol. 27, no. 4, pp. 678–693, 2017.
[25] P. T. Costa and R. R. McCrae, "Normal personality assessment in clinical practice: The NEO personality inventory." *Psychological Assessment*, vol. 4, no. 1, p. 5, 1992.
[26] P. T. Costa Jr and R. R. McCrae, "Four ways five factors are basic," *Personality and Individual Differences*, vol. 13, no. 6, pp. 653–665, 1992.
[27] R. Ganesan, S. Jajodia, and H. Cam, "Optimal scheduling of cybersecurity analyst for minimizing risk," *ACM Transactions on Intelligent Systems and Technology*, vol. 8, no. 4, Feb. 2017.
[28] C. Puppe and A. Tasnádi, "Optimal redistricting under geographical constraints: Why pack and crack does not work," *Economics Letters*, vol. 105, no. 1, pp. 93–96, 2009.
[29] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2022. [Online]. Available: https://www.gurobi.com
[30] A. Shah, R. Ganesan, S. Jajodia, and H. Cam, "An outsourcing model for alert analysis in a cybersecurity operations center," *ACM Transactions on the Web (TWEB)*, vol. 14, no. 1, pp. 1–22, 2020.
[31] ——, "A methodology to measure and monitor level of operational effectiveness of a CSOC," *International Journal of Information Security, Springer*, vol. 17, no. 2, pp. 121–134, 2018.

Sushil Jajodia is University Professor, BDM International Professor, and the founding director of Center for Secure Information Systems in the College of Engineering and Computing at the George Mason University, Fairfax, Virginia. He is also the director of the NSF IUCRC Center for Cybersecurity Analytics and Automation (CCAA).

Dr. Jajodia has made research contributions to diverse aspects of security and privacy, including access control, multilevel secure databases, vulnerability analysis, moving target defense, cloud security, and steganography, as well as replicated and temporal databases and algebraic topology. He has authored or coauthored seven books, edited 53 books and conference proceedings, and published more than 500 technical papers in the refereed journals and conference proceedings. He is also a holder of 28 patents, 17 of which have been licensed by successful startups. He is a fellow of ACM, IEEE, and IFIP; and recipient of numerous awards including the IEEE Computer Society W. Wallace McDowell Award. According to the Google Scholar, he has over 50,000 citations and his h-index is 112. He has supervised 27 doctoral dissertations; thirteen of these graduates hold academic positions while rest are in successful industrial positions.

Hasan Cam is the founder and CEO of SiberYZ. He currently works on the projects involved with providing cybersecurity resilience over smart factories, grids, and homes. Previously, he was a Principal Machine Learning Scientist at Best Buy, and a Computer Scientist at the U.S. Army Research Laboratory. He is a Senior Member of IEEE.

Ankit Shah is an Assistant Professor of Industrial and Management Systems Engineering at the University of South Florida. His research interests lie at the intersection of computer science, operations research, data science, and information technology, with a strong focus on artificial intelligence for cybersecurity. His current research work is in the area of deep reinforcement learning and adversarial machine learning for threat reduction and secure systems.

Rajesh Ganesan is an Associate Professor of systems engineering and operations research with George Mason University, Fairfax, VA, USA. His research interests include stochastic optimization (approximate dynamic programming), Bigdata analytics, multiscale statistical data analysis using wavelets, and engineering education, with applications in cybersecurity, healthcare, defense, air transportation, and nanomanufacturing.

Steve Hutchinson is with the Maryland Innovation and Security Institute, Columbia, MD, USA.