MODELING AND SIMULATION OF COOPERATION AND LEARNING IN CYBER SECURITY DEFENSE TEAMS

Pasquale Legato and Rina Mary Mazza
Department of Informatics, Modeling, Electronics and System Engineering
University of Calabria
Via P. Bucci 42C - 87036, Rende (CS), Italy
e-mail: {legato, rmazza}@dimes.unical.it

KEYWORDS
Simulation optimization, cyber security, team formation and cooperation.

ABSTRACT

Cyber security analysts may be organized in teams to share skills and support each other upon the occurrence of cyber attacks. Team work is expected to enforce the mitigation capability against unpredictable attacks addressed against a set of cyber assets requiring protection. A conceptual model for evaluating the expected performances of cooperating analysts by reproducing their learning process within a team is proposed. Analytical approaches to solve the underlying state-space model under stochastic evolution and discrete-event simulation are both discussed. The basic assumption is that a set of regeneration points corresponds to skill achievement through learning. A Simulation-based Optimization (SO) tool ranging from the inner level modeling of the cooperation-based learning process to the outer assignment of analysts to assets is then presented. Team formation may be supported by the SO tool for obtaining the team composition, in terms of individuals and skills, that maximizes system performance measures. Numerical results are reported for illustrative purposes.

INTRODUCTION

In today's rapidly changing threat landscape, cyber attacks are becoming much more common and damaging. To keep pace with this change, one must start by recognizing that cyber defence is as much about people as it is about technology. As a matter of fact, governments, organizations and industry worldwide are seeking ways to both manage and improve the expertise of their human resources in order to prevent, mitigate and recover from cyber attacks (NATO 2016).

Generally speaking, with respect to human resource management, a company may decide to deploy its cyber defense security analysts according to different *modus operandi*. Cyber defense analysts may be called to *i*) work alone according to traditional individualistic approaches or *ii*) in consultation with other analysts who are committed to a common mission and are willing to share the knowledge that is necessary to fulfill that mission (Kvan and Candy 2000). In the former case, it is a matter of being in charge of one's own achievements, concentration and schedule when deciding what to do and when to do it. In the latter, it is about teaming two or more individuals, with complementary background and skills, who organize their efforts in a mutual supportive way, share experience and complete common tasks.

Testing and evaluating the suitability of either of these two alternatives in a cyber attack scenario, under stochastic attack arrivals and mitigation services, requires a systematic approach in order to *i*) evaluate overall attack tolerance with regard to system performance degradation and *ii*) assess the effectiveness of using cooperation-driven learning among teammates as a countermeasure against cyber attacks.

To this purpose, a state-space model is exploited to mimic the learning process of an analyst when working in consultation with other analysts. This learning model, applied to the cooperating members of the same team, allows to account for the acquisition of new skills or the growth of expertise on pre-existing skills. So, cooperation-driven learning becomes a countermeasure against attacks by increasing positive attack mitigation in whatever be scenario. Both the analytical and simulation solution of the model are discussed as possible evaluation tools within a more general and powerful framework (Legato and Mazza 2016) aimed at optimizing the benefits from cooperation-based learning.

The paper is organized as follows. A description of the cooperation and learning process is provided in section 2. The conceptual model of the attack arrival and mitigation process and how it can account for cooperation-based learning among analysts is introduced in section 3. Analytical and simulation methodologies for model solution are discussed in section 4, while section 5 focuses on the choice to embed an SO procedure in the overall solution framework. The SO tool is presented in section 6 and conclusions are drawn in section 7.

COOPERATION AND LEARNING

When working in a team, cooperating with colleagues is at the basis of the learning process experienced by any given analyst during his/her daily task of defending an assigned cyber asset against unpredictable attacks. The capability of an analyst to gain and, in turn, transfer

Proceedings 31st European Conference on Modelling and
Simulation ©ECMS Zita Zoltay Paprika, Péter Horák, Kata Váradi,
Péter Tamás Zwierczyk, Ágnes Vidovics-Dancs, János Péter Rádics (Editors)
ISBN: 978-0-9932440-4-9/ ISBN: 978-0-9932440-5-6 (CD)

knowledge depends on which skills he/she bears, along with the specific level held for each skill. In the attempt to quantify measurable knowledge, here we consider a four-level scale: no level of skill, basic level, intermediate level and expert level. These levels are in a continuum meaning that anyone standing at any of these levels can pursue progressive skill acquisition and, thus, learn through the continuum until he/she has reached the expert level. Put in other terms, unskilled analysts can learn from basic level analysts, intermediate level analysts and expert level analysts; basic level analysts can learn from intermediate level analysts and expert level analysts; intermediate level analysts can learn from expert level analysts. Obviously, expert level analysts can only act as hand-on knowledge workers, in addition to their individual role as cyber attack mitigation units.

Let us now consider a possible evaluation program according to which credits are awarded to analysts for every attack their skills allow them to mitigate. A credit is a measure of security performance whose amount depends on the type of attack – the more dangerous the attack, the higher the amount of credits rewarded. Whether working alone or in cooperation with others, an analyst's behavior is meant to collect as many credits as possible over time. If the analyst works alone and is skilled to manage an incoming threat, mitigation occurs according to a service time that depends on the type of attack. As a result of attack mitigation, the analyst is rewarded with the entire credit and free again to face new threats. If no such skill is held by the analyst, the lack of ability to mitigate the malicious attack may produce a negative impact on the entire system and likely cause a loss of overall performance, unless he/she works in consultation with other analysts. Practically speaking, the analyst may consult with skilled team members, if any, and thus start acquiring the necessary knowledge to manage the attack. If such a teammate exists and is not already engaged in other activities, attack mitigation may begin; otherwise, the analyst must wait. As a result of the ongoing interaction process, the "enquiring" analyst starts to accumulate knowledge (e.g. one scalar unit for each mitigation completed under cooperation or for each time unit spent in a cooperation state) because of the learning process he/she is undergoing. The team members that took part in the knowledge-sharing process then share the related credits. If none of the team members hold the appropriate skill to manage the attack, similar to the work alone *modus operandi*, this lack produces a negative impact on the entire system and causes a loss of overall performance.

A new skill achieved through the learning process may be recognized to the analyst after a (positive) periodic verification, provided that a fixed threshold on the number of attacks mitigated in cooperation has been reached. This skill recognition is the result of an examination whose timeframes and procedures are approximately scheduled by the senior management. Our focus is to investigate the extent to which there is a practical possibility of building a stochastic model of an analyst's evolution through his/her work by means of cooperation-based learning activities and then solving it by analytical approaches and simulation techniques.

CONCEPTUAL MODEL

From a conceptual point of view, the attack arrival and mitigation process experienced by an asset which is protected by a suitably skilled analyst may be represented by the client-server paradigm. The attack is a client who arrives randomly to a given asset and should be mitigated (serviced) by the analyst dedicated to that asset. A queued attack represents a non-detected status, but, as the attack "waits" to be detected, usually the damage delivered to the asset becomes bigger. Once detected, the attack either receives a mitigation service from the dedicated analyst or the analyst is forced to ask for cooperation from a colleague dedicated to another asset. Whatever be the case, the time required by the analyst(s) to mitigate the attack is random and may grow larger due to both a greater time of detection and a delay in starting the mitigation caused by a lack of cooperation.

The client-server model at hand is illustrated in Figure 1 for a couple of assets with two cooperating analysts. Clearly, it results in a non-standard queueing system not only because of the correlation among service duration and waiting time, but also because, when called to cooperate with a teammate, a server-analyst appears to be on "vacation" to his/her asset and to other colleagues as well.

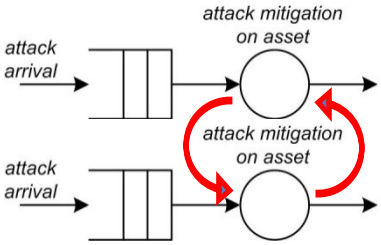

Figure 1: The Client-Server Model for the Attack Arrival and Mitigation Process under Cooperation

After using the above client-server system to represent working and cooperation among analysts, now we need to model the underlying learning process. A possible learning model should quantify, by means of a scalar quantity, the amount of knowledge incorporated by an analyst as a result of his/her attack mitigation experience under different scenarios and management policies.

In this respect, state-space models have already proven to be successful (Distefano et al. 2012) in capturing dynamic effects in reliability and availability studies. So, we choose to model the analyst's evolution over time by means of a sequence of states, each representing the different conditions in which an analyst may be found. These conditions are: idle, busy, waiting for colleague, teaching and learning, as illustrated in Figure 2.

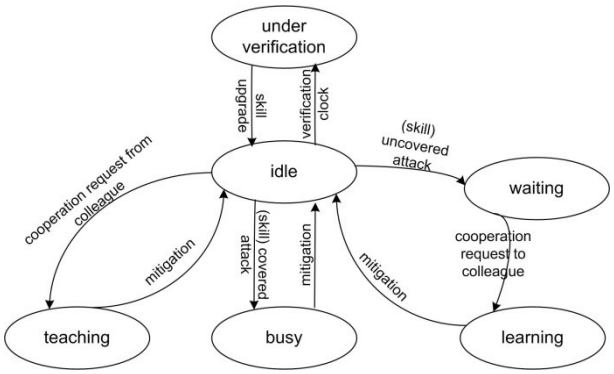

Figure 2: State-based Representation of Analyst Evolution

In principle, the learning state could be completely characterized by a mathematical real-valued function capturing the instantaneous measure of the learning growth. Similarly to the instantaneous hazard function in reliability modeling (Trivedi 2002), the rate of learning may be dependent or not from the age in a proper state. As for the skill upgrade policy, the real domain suggests that skill upgrade is the result of a successful verification of the knowledge gain achieved by an analyst between two successive verification instants. Verification activities occur rather cyclically along the analyst's working life and is aimed at evaluating the (cumulative) time spent by the analyst in the cooperation-based learning state (through repeated visits).

METHODOLOGIES FOR SOLUTION

Markov and generalized Markov models (Kulkarni 2009) are at the basis of the analytical tractability of a state-space based stochastic model. Nowadays, efficient tools are available to also manage the case where a large set of states need to be considered (Trivedi and Sahner 2009). However, besides the size of the model at hand, the mathematical tractability of the state-space model in Figure 2 relies upon the existence of points within the process where the memoryless property occurs. In our case, these points may occur at the time epochs of any given verification, followed by the upgrade of the analyst's skill level. Moreover, the further assumption that any future state of the analyst between two successive regeneration points can only depend from his/her state at the latest regeneration point, leads us to the concept of Markov renewal sequences. Therefore, Markov regenerative processes (Logothetis et al. 1995) appear to be the most powerful analytical tool for the stochastic modeling of the cooperation-based learning process of our interest. In particular, MRGPs allow generally distributed clock times for skill verification and upgrade. To this respect, one should formulate the MRGP through the definition of the related three matrix valued functions concerning transition probabilities, global kernel and local kernel (op. cit.). In particular, the local kernel matrix describes the behavior of the MRGP between two consecutive regeneration time points. It is required to compute the steady-state solution of the MRGP, provided that the discrete-time Markov chain embedded at the regeneration points is finite, aperiodic and irreducible and, therefore, returns the unique steady-state solution for the state visit ratios.

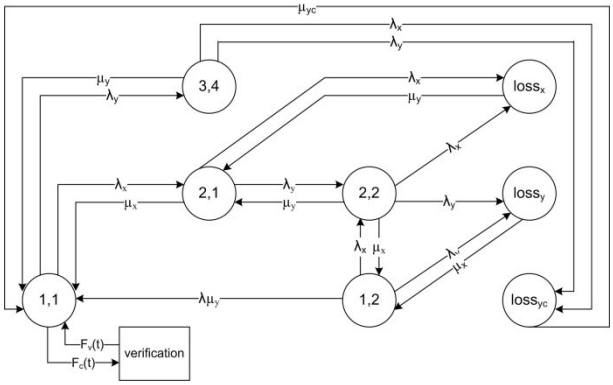

Figure 3: State-Diagram for a Simple Model

For illustrative purposes, let us consider the simple model in Figure 3 in which two cyber assets are both subject to two different attacks - attack x, attack y with rates λ_x and λ_y, respectively. Asset 1 is protected by analyst 1 who bears only the skill required to mitigate x-type attacks with mitigation rate μ_x. Similarly, asset 2 is protected by analyst 2 who bears only the skill required to mitigate y-type attacks with mitigation rate μ_y. Cooperation between the two analysts may occur when one requires the skill of the other to cover the attach on his/her asset. In this example, we assume that mitigation, with or without cooperation, should start immediately; otherwise, a "loss" (i.e. degradation) in system performance occurs. Thus, the state of an analyst is one of the following: 1 – idle, 2 – busy, 3 – learning, and 4 – teaching. Observe that the "verification" state, refers to skill verification for both the analysts. Its occurrence is regulated by the distribution function of the verification clock time, $F_c(t)$, while its time to positive completion is regulated by the distribution function $F_v(t)$ and leads to a new regeneration cycle under an upgraded skill level.

The state-diagram amenable to an MRGP-based analysis is described as follows. State (1,1) means that both analysts are idle and verification occurs only in this state; state (2,1) means that analyst 1 is busy because of the occurrence of an x-type attack on his/her own asset; (2,2) means that both analysts are busy mitigating an x-type and y-type attack on their respective assets; state (1,2) means that analyst 2 is busy because of the occurrence of an y-type attack on his/her own asset; (loss$_*$) represents the impossibility to provide mitigation for neither attacks, whether separately (loss$_x$, loss$_y$) or in cooperation (loss$_{yc}$, loss$_{xc}$); (3,4) means that analyst 1 is learning by cooperating with analyst 2 in mitigating an y-type attack on his/her own asset;); (4,3) means that analyst 2 is learning by cooperating with analyst 1 in mitigating an x-type attack on his/her own asset. For

simplicity, in Figure 3 we represent only half of the entire model, i.e. the part referred to attack and mitigation activities occurring on asset 1. The corresponding activities on asset 2 are defined by analogy.

In principle, the state-diagram embedded in the regeneration cycle could be a continuous-time Markov chain provided that we assume exponential distributions for both attack occurrences and mitigation services, under independent sojourn times per visit in any given state. One could relax the exponential assumption to get a semi-Markov embedded process. However, the independence assumption on the sojourn time per visit would somehow limit the representation of the underlying memory effect of the learning activity accomplished by correlated successive returns in the same learning state.

This stated, it is our belief that providing a simulation framework for both performance evaluation and optimization is worth becoming the major goal of our current research contribution. Simulation allows us to obtain greater flexibility in setting a more realistic queuing-based description of the management policies regarding team formation and analyst cooperation in the cyber security real domain (Poste Italiane, the Italian national postal service) that has stimulated our work. Here, dealing with at least 10 different assets and ten analysts is quite common. Moreover, the optimal assignment of skilled analysts to assets may be pursued by embedding a suitable meta-heuristic based search process for better feasible assignments and, thus, learning by cooperation within a simulation based optimization tool.

Within the simulation framework, the quantitative analysis of the analyst's cooperation-based learning process between two (and also over many) successive regeneration epochs is carried out by regenerative discrete-event simulation (Shedler 1992). The effectiveness of regenerative simulation relies upon the practical possibility of replicating a suitable number of sample trajectories containing regeneration epochs and, therefore, estimate the expected performance measures of interest by both intra- and inter-cycle sample means. In particular, we may estimate the long-run probability of finding an analyst mitigating an attack in cooperation with a colleague. This probability is given by the ratio between the expected time spent by the analyst in a learning state divided by the expected duration of the overall time needed by that analyst to change his/her skill level:

$$P[\text{analyst in learning state}] = \frac{E[\text{time spent in learning state}]}{E[\text{time between change of skill level}]}$$

So, the simulation tool may also be used to validate analytical tools based on MGRPs against real policies, data and statistical distributions. This will be the subject of a companion paper.

SIMULATION-BASED OPTIMIZATION

In the illustrative example presented in the previous section, the analysts have already been assigned to a specific cyber asset and their teaming into a group of two is the only option available. When the number of analysts, assets, skills and skill levels grows larger, analyst-asset assignment and team formation is far from being so straightforward. Due to all the possible combinations, one may have to first *generate* teams by means of a search process and then *evaluate* them by means of an evaluation process. This situation may benefit from introducing a simulation-based optimization procedure (SO) (Fu and Nelson 2003) in the framework under development. In an SO approach a structured iterative approach calls an optimization algorithm to decide how to change the values for the set of input parameters (e.g. analyst-asset assignment and team formation) and then uses the responses generated by simulation runs to guide the selection of the next set. The logic of this approach is shown in Figure 4.

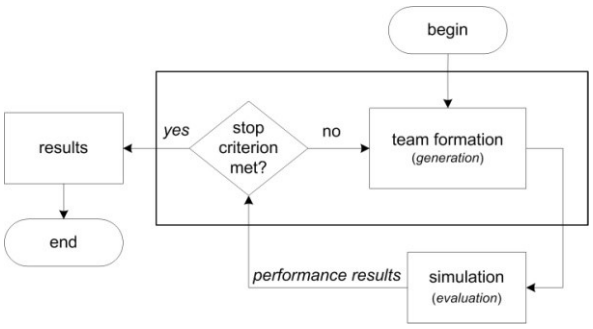

Figure 4: The Logic of Simulation-based Optimization

On the *generation* side, the first step in assembling cyber defence teams consists in assigning the analysts to the cyber assets, at least one for each asset. In the second step, the actual team formation is carried out by deciding the maximum number of teammates that can form a team, along with the specification of which colleagues should be part of the team. This step is guided by the idea of grouping analysts with complementary skills in order to cover a wider range of attacks and, thus, favor the learning process among as many teammates as possible. On the other hand, too many unskilled members on a team may prevent the skilled teammates from protecting their own assets because too busy in answering support requests from teammates.

During team generation, should an exhaustive coverage of all the possible system combinations be not reasonable, nor affordable from a computational point of view, then metaheuristic-based approaches would have to be addressed. Here we use a *simulated annealing* procedure (Kirkpatrick et al. 1983) which was first introduced by developing the similarities between combinatorial optimization problems and statistical mechanics. In the field of metal sciences, the annealing

process is used to eliminate the reticular defects from crystals by heating and then gradually cooling the metal. In our case, a reticular defect could be seen as grouping analysts in teams that are not able to "properly" protect cyber assets and, thus, guarantee a given quality of service level when the above assets undergo an attack. Technically speaking, the annealing process is aimed to generate feasible teams of analysts, explore them in a more or less restricted amount and, finally, stop at a satisfactory solution. To avoid getting caught in local minima, during the exploration process a transition to a worse feasible solution can occur with probability

$p=exp(\Delta/T)$

where Δ is the difference between the values of the objective function (measure of learning) of the current solution (state) θ and the candidate solution θ_t and T is the process temperature. A prefixed value of T determines the stop of the entire process and it usually decreases according to a so-called *cooling schema*. Unfortunately, in the literature there is no algorithm that can determine "correct" values for the initial temperature and cooling schema, but, as suggested by empirical knowledge simple cooling schemas seem to work well (Ingber 1993).
In the following, some pseudo-code is given for the original SA algorithm for a minimization problem.

```
Algorithm: Simulated Annealing
1: θ ← initial solution
2: for time = 1 to time-budget do
3:     T ← cooling-schema[time]
4:     if T=0 then
5:         Present current solution as the estimate of
           the optimal solution and stop
6:     Generate a random neighbor θt of the current
       solution θ by performing a move.
7:     Δ = f(θ) − f(θt)
8:     if Δ > 0 then
9:         θ ← θt
10:    else
11:        θ ← θt (with probability p=exp(Δ/T))
12: end for
```

When customizing the SA algorithm to our problem, some choices need to be made.
To begin with, choosing the proper cooling schema has great impact on reaching a global minimum. In particular, it affects the number and which analysts are assigned to a team (solutions) that will be evaluated by running the SA algorithm. To this end, the so-called simple mathematical cooling schema $T_{i+1}=\alpha \cdot T_i$ has been tested, and the best results are returned for an initial temperature $T_0=100$ and a decreasing rate $\alpha \cong 0.9$.
The "move" definition for neighborhood generation is very context-sensitive. For our problem, a move must be defined with respect to the feasibility (or lack thereof) of a team by taking into account the analysts' skills. Some examples of moves are:

- move analyst *l* from team *i* to team *j* ($i <> j$);
- swap analyst *l* and analyst *k* ($l <> k$), originally assigned to team *i* and team *j* ($i <> j$), respectively.

As far as the stopping criteria are concerned, designers can choose among the following possibilities:

- stop when the algorithm has reached a fixed number of iterations *n* or an upper bound on the available time-budget;
- stop when the current solution has not been updated in the last *m* iterations;
- stop when the cooling schema has reached a *lower bound* on the temperature.

Although we do not use this algorithm to perform an exhaustive search of the sample space, nor are we provided with any sort of control running on which part of the feasible set is being explored, the solutions returned as final output are likely to belong to the set of optimal global solutions (Banks et al. 2000) that under the cooperation-based learning policy allows to deliver:

- knowledge gain;
- percentage of attacks mitigated;
- resource (analyst) utilization;
- number of credits gained;
- number of cyber defense security analysts per team;
- cyber defense security team composition in terms of skill types and levels held by every single analyst assigned to every single team.

On the *evaluation* side, the simulation model for the attack arrival & mitigation process has been conceived according to an attack-centric point of view: attacks are entities flowing through a cyber network that may damage cyber assets and call for mitigation by (a group of) skilled analysts who seize and/or release resources while doing so.
An attack is defined by a record:

```
type attackrecord
   Eventtype
   Arrivaltime
   Analyst
   teammate(s)
   Asset
   Attacktype
   Operationtime
   Queue
end type
```

The primary attack events (in italics) of the simulation model are listed in Table 1, along with their effect on the system state in terms of actions and resources seized and/or released. Each event marks the beginning or the end of a given model activity and must be counted only once. An event always triggers the beginning (end) of a specific activity, but, for the sake of shorter notation, any "begin" ("end") prefix (suffix) is omitted from the event name.

Table 1: Events of the Discrete-Event Simulation Model

Event	Actions	Resources	
		Seize	Release
attack_arrival	schedules *attack_outcome*	analyst	-
	schedules *attack_outcome*	teammate	
	queues request on analyst	-	-
	queues request on teammate	-	-
	updates statistics	-	-
attack_outcome	schedules *attack_arrival*	-	analyst
	schedules *attack_arrival*		teammate
	unqueues request on analyst	analyst	-
	unqueues request on teammate	teammate	-
	updates statistics	-	-

THE SO TOOL

The SO tool is illustrated in Figure 5. It has been designed and implemented in compliance with all the conventional steps used to guide a thorough and sound simulation study (Banks et al. 2000).

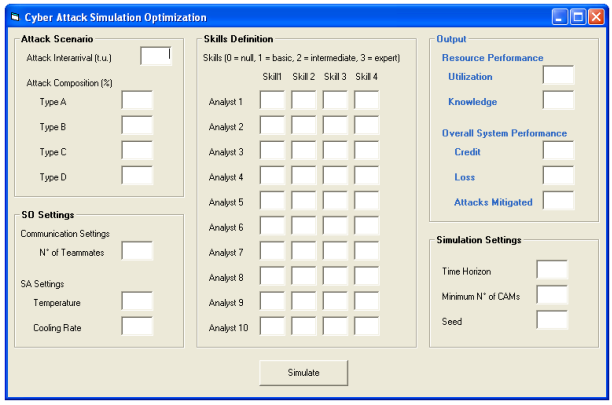

Figure 5: Snapshot of the SO Tool

All experiments have been run on a personal computer equipped with a 2.26Hz Inter Core™2 duo processor and 3 GB of RAM.

The GUI panel in Figure 5 has been conceived to meet the suggestions given by the cyber security senior management of Poste Italiane. It allows to easily specify the input data and SO parameters by means of proper sections. In the particular case at hand, we consider an attack scenario in which 10 analysts are used to defend 10 cyber assets against 4 types of attacks. Both attack interarrivals (in time units) and composition (in percentage) with respect to different types of attacks need to be specified. The skills of the 10 analysts are then defined by specifying for every analyst which skills he/she features and the level of competence for each skill (0=no skill, 1=basic, 2=intermediate, 3=expert). After inserting the maximum number of teammates in a group (here ranging between 1 and 10), the input stage is then completed by providing the SA and simulation settings. These are, respectively, the initial temperature along with the cooling rate of the SA procedure, the overall time horizon, the minimum number of cooperative attack mitigations (CAMs) completed per verification cycle and the simulation seed.

We now define the scenario for the preliminary set of experiments in order to evaluate the management of analysts and the resulting system performance. The expected measure of learning, skill upgrade, best team(s) composition (in number and skills) are returned for the former, while system credit and system loss are returned for the latter.

On average, attacks occur every 100 time units according to an exponential renewal process (λ, the average interarrival rate, is thus equal to 1/100). Arrivals are characterized by a combination of 4 different types of attacks (i.e. 70% type A, 15% type B, 10% type C and 5% type D). The 10 analysts are able to provide attack mitigation according to their own skills which are reported in Table 2.

Table 2: Skill Types and Levels of the Cyber Defense Security Staff

Analyst/Skill	A	B	C	D
analyst 1	0	0	1	0
analyst 2	3	0	1	0
analyst 3	0	2	0	3
analyst 4	0	0	3	0
analyst 5	3	0	0	0
analyst 6	3	3	0	0
analyst 7	0	1	0	3
analyst 8	1	0	0	0
analyst 9	3	3	3	0
analyst 10	3	3	0	0

In the given scenario, analysts respond to attacks by working alone (n° of teammates=1) or in cooperation with other analysts (n° of teammates>1). The rate (μ) of the attack mitigation activity depends on the type of attack, the skill level held by the analyst and if mitigation occurs alone or in cooperation with other analysts. In the later case, mitigation times are inflated by 30%.

The initial temperature and the cooling rate of the SA scheme are set equal to 100 and at least 0.948, respectively, so that at least 100 different team-formation and assignment configurations are considered for the given scenario. The time horizon is fixed at 14400 time units and both point estimates and 95% confidence intervals can be obtained for the measures of learning, system credit and system loss. Here, for clarity of illustration, in Figures 6 through 8 we prefer plotting the central value within the interval estimates to show some preliminary numerical results.

Let us start by considering the (average) measure of learning (i.e. one scalar unit for each mitigation completed under cooperation). Figure 6 shows that this measure grows approximately linearly with the number of analysts per team. In other terms, thanks to cooperation, the cumulative number of attacks mitigated by all the analysts in the fixed time horizon goes from 0 (no cooperation) to 180 (complete cooperation among the 10 analysts).

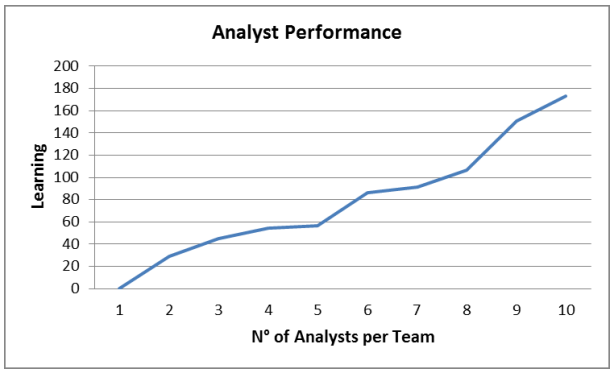

Figure 6: Trend of Analyst Learning

From the analyst's individual point of view, the learning benefit is resumed in Table 3 which, for each analyst, reports his/her skill upgrade achieved through cooperation along the time horizon.

Table 3: Skill Upgrade per Analyst

Analyst	Initial Skills	Final Skills
1	C	A, B, C
2	A, C	A, B, C, D
3	B, D	A, B, C, D
4	C	A, B, C
5	A	A, B, C
6	A, B	A, B
7	B, D	A, B, C, D
8	A	A, B, C
9	A, B, C	A, B, C, D
10	A, B	A, B

As one may see from Table 4, this knowledge growth is accomplished in conjunction with a specific asset-analyst assignment and subsequent team formation in which analysts with complementary skills have been teamed together.

Table 4: Details of the Asset-Analyst Assignment and Team Skills when Max No. of Teammates=7

Analyst	Asset	Teammates	Team Skills
1	6	2, 4, 5, 6, 8 & 10	A, B, C
2	9	1, 3, 4, 5, 7 & 10	A, B, C, D
3	10	2, 4, 5, 6, 7 & 9	A, B, C, D
4	7	1, 2, 3, 5, 6 & 10	A, B, C, D
5	8	1, 2, 3, 4, 6 & 9	A, B, C, D
6	4	1, 3, 4, 5, 7 & 8	A, B, C, D
7	2	2, 3, 6, 8, 9 & 10	A, B, C, D
8	5	1, 6, 7, 9 & 10	A, B, C, D
9	1	3, 5, 7, 8 & 10	A, B, D
10	3	1, 2, 4, 7, 8 & 9	A, B, C, D

As for the remaining performance measures, let us first consider the system credit recalling that the more dangerous the attack, the higher the amount of credits rewarded. In this set of experiments, A-type attacks are the less dangerous (1 credit rewarded per mitigation), while D-type attacks are the most dangerous (4 credits rewarded per mitigation). As shown in Figure 7, the behavior of the system credit follows a bathtub curve as the number of analysts per team grows. For small number of teammates, the benefit of cooperation is surmounted by the waiting times experienced by the requiring analysts when asking (a limited number) of skilled teammates for support. For middle-size teams, system credit remains rather stable. This is likely due to the greater number of skills and, thus, attacks covered by the teammates which affects the waiting times in a positive way. For a large number of analysts per team, both cooperation and waiting times benefit from the availability of all the skills against incoming attacks.

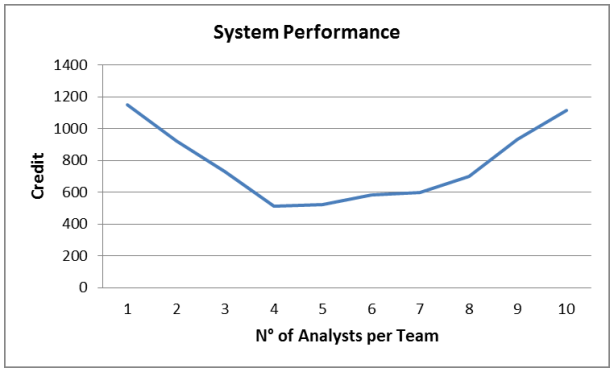

Figure 7: Trend of System Credit

As for system loss, this is a measure of the number of attacks that cannot be mitigated by the analysts due to a lack of the skills required to fulfil this purpose. By analogy with system credit, it is equal to 1 per non mitigated A-type attack and reaches 4 per non mitigated D-type attacks.

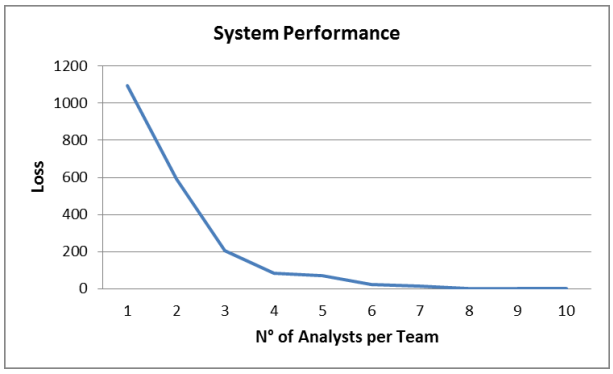

Figure 8: Trend of System Loss

Figure 8 shows that in the scenario under examination system loss is totally overcome when the number of teammates is equal to or greater than 7. As previously stated, this corresponds to the possibility of always finding a free skilled teammate upon request by an analyst.

CONCLUSIONS

A client-server system with multiple cooperating servers bearing different skills and learning capabilities has been discussed and illustrated. The centrality of an MRGP as an evaluation tool for the analysis has been highlighted. Then, a more general simulation tool has been proposed with the aim of pursuing the optimality of dynamic team formation to favor cooperation and learning among security analysts each dedicated to their own asset. The simulation-based approach allows for an effective evaluation and optimization of the whole organizational process, starting by the assignments of analysts to assets and reproducing the occurrence of attacks followed by cooperation and learning. The tool may also incorporate a relaxed MRGP aimed at reproducing the learning process of an analyst when working in consultation with other analysts. The learning model, applied to the cooperating members of the same team, allows to account for the acquisition of new skills or the growth of expertise on pre-existing skills. We have shown how the tool may be used to *i*) evaluate overall attack tolerance, in terms of credit and loss measures, with respect to system performance and *ii*) assess the effectiveness of using cooperation-driven learning as a countermeasure against cyber attacks, in terms of new skills achieved by analysts.

ACKNOWLEDGEMENTS

This work was partly supported by the Cyber Security Technological District funded within the Italian National Operational Programme for Research and Competitiveness 2007-2013 under grant number PON03PE_00032_2.

REFERENCES

Banks, J., J.S. Carson, B.L. Nelson and D.M. Nicol. 2000. *Discrete-Event System Simulation*. 3rd Edition. Prentice-Hall, Inc., Upper Saddle River, New Jersey.

Distefano, S., F. Longo, F., and K.S. Trivedi. 2012. "Investigating Dynamic Reliability and Availability through State-Space Models". In *Computers & Mathematics with Applications*, 64, No.12, 3701-3716.

Fu, M. and B. Nelson. 2003. Guest Editorial. *ACM Transactions on Modeling and Computer Simulation* 13, No.2, 105-107.

Ingber, L. 1993. "Simulated Annealing: Practice versus Theory". *Mathematical Computer Modelling* 18, No.11, 29-57.

Kirkpatrick, S., C.D. Gelatt and M.P. Vecchi. 1983. "Optimization by Simulated Annealing". *Science*, New Series, 220, No.4598, 671-680.

Legato, P. and R.M. Mazza. 2016. "A Simulation Optimisation-based Approach for Team Building in Cyber Security". *International Journal of Simulation and Process Modelling* 11, No.6, 430-442.

Logothetis, D., K.S. Trivedi and A. Puliafito. 1995. "Markov Regenerative Models". In: *Proceedings of the IEEE International Computer Performance and Dependability Symposium* (Erlangen, DE, April 24-26), 134-142.

Kulkarni, V. G. 2009. *Modeling and Analysis of Stochastic Systems*. 2nd edition. Chapman & Hall, London.

Kvan, T. and L. Candy. 2000. "Designing Collaborative Environments for Strategic Knowledge in Design". *Knowledge-Based Systems* 13, No.6 (Nov), 429-438.

NATO. 2016. Cyber Defence. http://www.nato.int/cps/en/natohq/topics_78170.htm [Last updated: 17 Jan. 2017, accessed on 14 Jul. 2016].

Shedler, G. S. 1992. *Regenerative Stochastic Simulation*. Academic Press - Elsevier, Oxford.

Trivedi, K. S. 2002. *Probability and Statistics, with Reliability, Queuing and Computer Science Applications*. 2nd edition. John Wiley & Sons, Inc., NY, NY.

Trivedi, K. S. and R.A. Sahner. 2009. "SHARPE at the Age of Twenty Two". *ACM SIGMETRICS Performance Evaluation Review* 36, No.4, 52-57.

AUTHOR BIOGRAPHIES

PASQUALE LEGATO is an Associate Professor of Operations Research in the Department of Informatics, Modeling, Electronics and System Engineering (DIMES) at the University of Calabria, Rende (CS, Italy). He has been a member of the Executive Board of the University of Calabria as well as university delegate for the supervision of associations and spin-offs from the University of Calabria. He has been involved in several EEC funded research projects aimed at the technological transfer of SO procedures and frameworks in logistics. He is a member of the INFORMS Simulation Society. His research activities focus on predictive stochastic models for cyber security, queuing network models, stochastic simulation and the integration of simulation techniques with combinatorial optimization algorithms. His e-mail address is: legato@dimes.unical.it and his web-page can be found at http://wwwinfo.dimes.unical.it/legato.

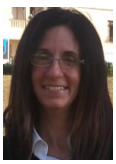

RINA MARY MAZZA is the Research Manager of the Department of Informatics, Modeling, Electronics and System Engineering (DIMES) at the University of Calabria, Rende (CS, Italy). She graduated in Management Engineering and received a PhD in Operations Research from the above university. She has a seven-year working experience on knowledge management and quality assurance in research centers. She has also been a consultant for operations modeling and simulation in container terminals. Her current research interests include discrete-event simulation and optimum-seeking by simulation in complex systems. Her e-mail address is: rmazza@dimes.unical.it.