Decision Networks for Modeling and Analysis of Attack/Defense Scenarios in Critical Infrastructures

D. Codetta-Raiteri, L. Portinale, R. Terruggia
Computer Science Institute, DiSIT
University of Piemonte Orientale
Alessandria, Italy

Abstract

We propose to exploit Decision Networks (DN) for the analysis of attack/defense scenarios. We show that DN extend both the modeling and the analysis capabilities of formalisms based on Attack Trees, which are the main reference model in such a context. Uncertainty can be addressed at every system level and a decision-theoretic analysis of the risk and of the selection of the best countermeasures can be implemented, by exploiting standard inference algorithms on DN.

Introduction

Security risk assessment and mitigation are important activities that must be performed "intelligently" and under uncertainty, to safely maintain critical infrastructures like computer systems and networks. The classical approach is to predefine a set of attack scenarios based on the knowledge of the systems and networks. Such scenarios are very often described and modeled through *Attack Trees* (AT) (Schneier 2000) where attacks can be represented in a tree structure, with the goal as the root node, and different ways of achieving that goal as multi-level hierarchical structures based on Boolean operators. Leaves represent basic attacks; these are specific operations an attacker can put in place, in order to pursue the ultimate goal, the latter represented by the top node. AT do not include defense mechanisms, so extensions have been proposed. In *Attack Countermeasure Trees* (ACT) (Roy, Kim, and Trivedi 2012) each countermeasure is the logical AND of two other constructs, called "detection" and "mitigation" events: a countermeasure is active when the attack has been both detected and mitigated. A scenario modeled through ACT is essentially based on the Boolean semantics; even if it is possible to introduce probabilistic parameters and to compute probabilistic indices, a sound decision-theoretic analysis is not directly supported, as well as the modeling of uncertainty at every arbitrary level.

These possibilities can be easily accounted for, by using *Decision Networks* (DN) (Jensen and Nielsen 2007) which extend modeling situations and patterns as defined by AT. This results in either well-known interactions mechanisms like noisy-AND/OR or in more general probabilistic dependencies (Langseth and Portinale 2007; Codetta-Raiteri, Montani, and Portinale 2010) DN allow the analyst to adopt a rational decision making approach, concerning the assessment of specific countermeasures in terms of expected utility or costs. Standard inference on DN can be used to compute posterior probability, given a set of observed evidence, for any variable of interest in the scenario. We can compute several indices, such as the *Birnbaum Importance* (BI) (Meng 2000) of attacks, and the *Return on Investment* (ROI) (Roy, Kim, and Trivedi 2012) of countermeasures. Finally, the determination of the suitable set of defense mechanisms can be naturally formulated as a decision problem, solved through DN inference. We consider a case study concerning the *Border Gateway Protocol* (BGP): we show how to derive a DN from an AT and how to perform quantitative and decision-theoretic analyses exploiting DN.

The case study

We consider an attack/defense scenario concerning a BGP session (Roy, Kim, and Trivedi 2012). BGP is used to exchange routing information across the Internet. An attacker prevents two peers from exchanging routing information by repeatedly causing a BGP session in "Established" state to reset. The BGP session can be reset by injecting a spoofed TCP (*Transmission Control Protocol*) or BGP message into the router message stream. Such spoofed packets can often be detected by methods such as the *Inter-domain packet filter* (IDPF) and mitigated by adding an MD5 (*Message-Digest* algorithm) based authentication for packets from the source host of the spoofed packet. Building a valid TCP/BGP packet requires a valid TCP sequence number (obtained by TCP sequence number prediction). During the initial stages of a TCP sequence number attack, a spoofed packet from an attacker is usually followed by the original packet from the authentic source. Detecting such duplicate packets can be a giveaway for on-going TCP sequence number attacks. Dropping compromised connections and initiating a new connection to destination with a different route will mitigate such attacks. Spoofed TCP message with RST flag set will cause a connection to reset. Spoofed BGP messages (Open, Notification or Keepalive messages) received by the BGP speaker in the "Connect" or "Active" states will cause the router to reset resulting in a denial of service. The BGP speaker can also be compromised by gaining physical or logical (hijacking a router management session) ac-

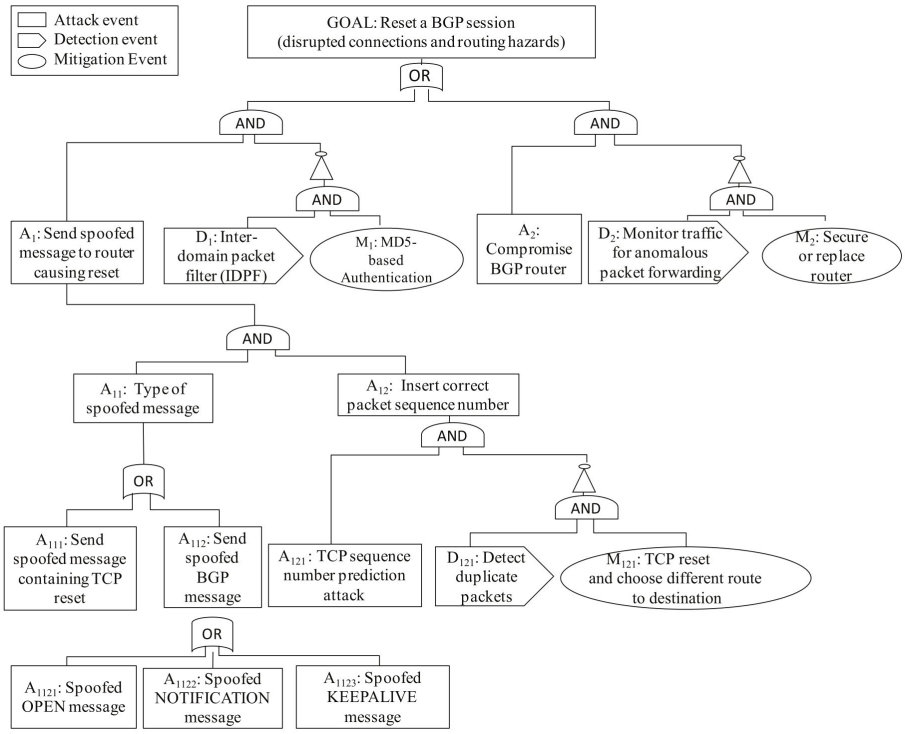

Figure 1: ACT for BGP attack/defense scenario (Roy, Kim, and Trivedi 2012)

cess to the router. Usually router hijacking is characterized by anomalous packet forwarding which can be detected by traffic monitoring at the router and mitigated by securing or replacing the router.

From ACT to DN

The ACT of Fig. 1 formalizes such a scenario and is completed by: the probability of occurrence of each atomic attack (attack leaf node), the probability of a successful detection (detection nodes), the probability of a successful mitigation (mitigation nodes), the security investment cost of countermeasures (both detections and mitigations), the impact or cost of each atomic attack. The ACT and all its parameters are taken from (Roy, Kim, and Trivedi 2012).

DN generation. We define the following rules to construct a DN from an ACT:

- for each atomic attack node A_i with probability p_{A_i}, create a binary chance node X_{A_i} (graphically represented as oval) with values *true* (occurrence of attack) and *false*; set the *Conditional Probability Table* (CPT) of X_{A_i} such that $P[X_{A_i} = true] = p_{A_i}$;
- for each countermeasure (i.e. a pair $CM = \langle D, M \rangle$ with D detection event and M mitigation event), create a binary decision node X_{CM} (rectangle) with values *active* (attack detected and mitigated) and *inactive*;
- for each attack event A output of a gate G (representing the Boolean function g) with inputs $A_1, \ldots A_k$, create a binary deterministic node X_E (double-boarded oval), set $X_{A_1} \ldots X_{A_k}$ as parent of X_E and set the deterministic function of X_E according to g;
- for each attack event A output of a gate G (Boolean function g) with input attack events $A_1, \ldots A_k$ and input countermeasure $CM = \langle D, M \rangle$ (with probability of detection p_D and probability of mitigation p_M), create a binary chance node X_E, set $X_{A_1} \ldots X_{A_k}, X_{CM}$ as parent of X_E and set the CPT for $X_E = true$ in the following way: entries corresponding to $X_{CM} = inactive$ are set according to the truth value of $g(X_{A_1}, \ldots X_{A_k})$ i.e. either 1 when *true* or 0 when *false*; entries corresponding to $X_{CM} = active$ are set as $g(X_{A_1}, \ldots X_{A_k})\ (1 - p_D\ p_M)$.

DN model. The DN corresponding to the ACT is reported in Fig. 2. We used GENIE (http://genie.sis.pitt.edu) to build the model and to perform every computation that is reported hereinafter; probabilities of attacks are annotated near chance nodes representing basic attacks, while probabilities of success of countermeasures are near decision nodes. Deterministic variable nodes have the corresponding Boolean function reported by the node itself.

Example. Tab. 1.a reports the CPT for node *TSNAS* given the parents *TSNA* and *DDRR*. Given that countermeasure *DDRR* is composed by the detection event *Detect Duplicate Packets* and by the mitigation event *TCP reset and different route to destination* (Fig. 1), with probability of detection $p_D = 0.8$ and probability of mitigation $p_M = 0.5$ respectively, the probability of countermeasure success is $p_{DDRR} = 0.8 \cdot 0.5 = 0.4$; thus, when the attack occurs (*TSNA=true*), if the countermeasure is activated (*DDRR=active*) there is a 60% chance of having the attack unmitigated and successful (CPT in Tab. 1.a).

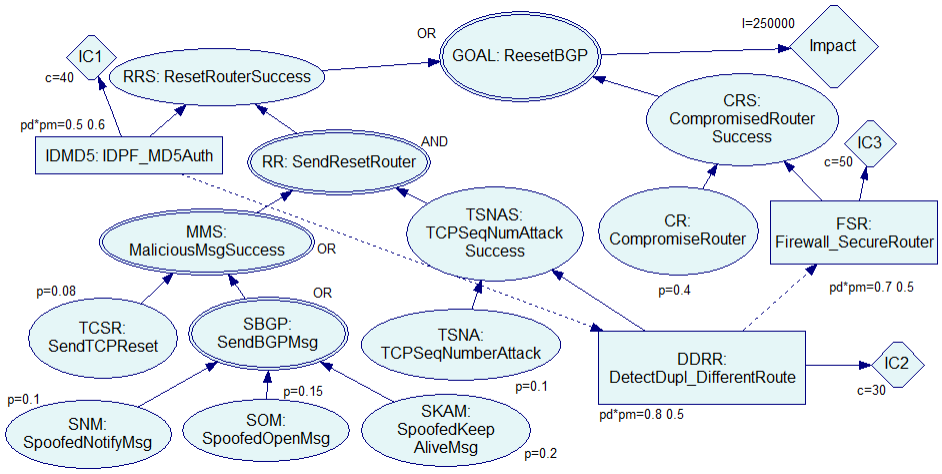

Figure 2: DN for the BGP scenario

Table 1: **a)** Modeling countermeasure success **b)** Adding uncertainty to GOAL in BGP scenario

a)

TSNA	**true**		**false**	
DDRR	**inactive**	**active**	**inactive**	**active**
TSNAS=true	1	0.6	0	0
TSNAS=false	0	0.4	1	1

b)

RRS	**true**		**false**	
CRS	**true**	**false**	**true**	**false**
GOAL=true	1	0.9802	1	0.01
GOAL=false	0	0.0198	0	0.99

Noisy gates. The structure created with this method represents a main skeleton over which more features can be added or modified. For example, in case of "noisy gates" (which are not representable in the ACT), we can adapt the corresponding CPT to account for this additional uncertainty.
Example. Consider the case where there is a 2% probability that a spoofed malicious message does not reach the router; in addition, suppose we want to model additional uncertainty, by introducing also a small chance (e.g. 1% probability) of the router being reset for some unmodeled causes. In such a case the type of node *ResetBGP* can be changed from deterministic to a noisy-OR chance node with leak (Jensen and Nielsen 2007); the noisy-OR parameters are then
p_1=*P[Goal=true—RRS=true, CRS=false]*=0.98
p_2=*P[Goal=true—CRS=true, RRS=false]*=1
p_{leak}=*P[Goal=true—RRS=false,CRS=false]*=0.01
resulting in the CPT of Tab. 1.b.

Quantitative Analysis

Probability of attacker's goal. This is the probability that an attacker will actually pursue the goal, given some initial specification in terms of probability of basic attacks and presence of countermeasures.
Example. We compute the probability of a successful attack given that we implemented only the IDPF detection with MD5 authentication (i.e. evidence inserted as $IDMD5 = active$ and the other countermeasures set to $inactive$), by performing the query $P[Goal|IDMD5 = active, DDRR = inactive, FSR = inactive] = 0.418$, resulting in more than 40% probability of being vulnerable to the attack. We notice that, given that countermeasure are not 100% effective, the attacker can reach the goal (with about 27% of probability) even in presence of all countermeasures: $P[Goal|IDMD5 = active, DDRR = active, FSR = active] = 0.274$
Importance. Such quantities can be defined in different ways and are usually identified with the aim of prioritizing defense mechanisms (i.e. countermeasures) to counteract attack events. An importance measure that can be adapted from reliability theory is the *Birnbaum Importance* (BI) (Meng 2000) which measures the change in the probability of the attacker's goal caused by a change in the probability of the attack of interest: $BI(A_i) = P[Goal = true|A_i = true] - P[Goal = true|A_i = false]$.
Example. Fig. 3.a shows BI under different sets of countermeasures. BI points out that compromising the router (attack *CR*) is the most important attack. This is due to the fact that unmitigating the attack will definitely cause the occurrence of the goal; moreover, BI puts in evidence that such an attack is more important in case it is not defended by the suitable countermeasure.
Risk. This corresponds to compute the *expected impact* of a particular attack/defense scenario. Since the impact I_{Goal} is measured as the amount of damage provided by the success of the attacker's goal, the risk is defined as $R = p_{Goal} I_{Goal}$ being p_{Goal} the probability of success of the goal (Roy, Kim, and Trivedi 2012). R is computed relatively to a particular context, usually a specific set of active countermeasures. In the DN framework, R can be computed by adding a value node (*Impact* in Fig. 2) to the goal node and setting the active and inactive countermeasures as evidence to decision nodes. The value function on the *Impact* node can be determined either by a direct estimation of the damage of a successful attack (occurrence of the goal) or by exploiting some heuristic approach synthesizing the impact on the goal, starting from

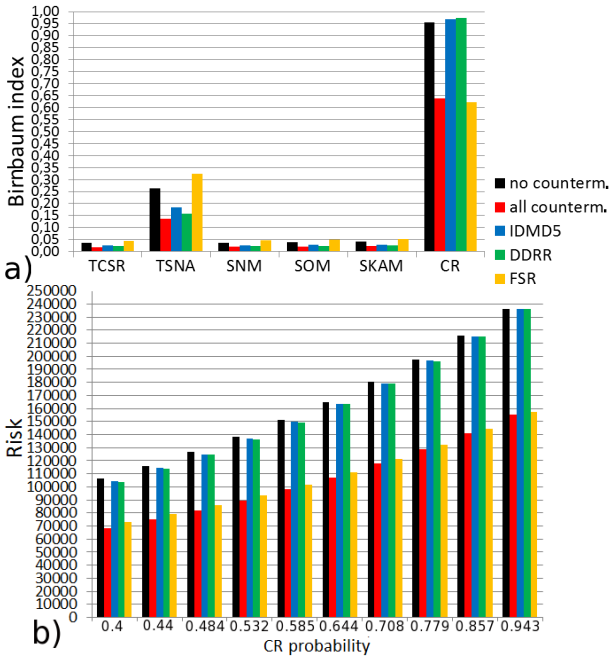

Figure 3: **a)** BI measure of basic attacks **b)** Risk evaluation w.r.t. probability of *CR* attack

Table 2: Best set of countermeasures

Obs. Att.	IDMD5	DDRR	FSR	Tot. Exp. Cost
None	✓	✓	✓	68515
CR=false	✓	✓		4658
CR=true	✓	✓	✓	164225
TSNA=true	✓	✓	✓	99072
TSNA=false			✓	65050
TCSR=true	✓	✓	✓	72890
TCSR=false	✓	✓	✓	68135

local impact estimation of each basic attack's impact (Roy, Kim, and Trivedi 2012). Since the goal's impact represents a damage (i.e. a cost), a negative utility can be used in the corresponding value node.

Example. Assuming $I_{Goal} = 250000$, Fig. 3.b reports the risk value (expected impact) with respect to different sets of active countermeasures, by varying the probability of *CR* attack (the most important one as noticed before). We confirm also from risk evaluation that activating countermeasures not related to a router hijacking attack is not useful at reducing the global risk, while the presence of an active firewall and an alternative routing strategy can provide a risk reduction.

Investment. When measuring the impact deriving from a set of countermeasures, the investment cost in setting up such defense mechanisms should be taken into account as well. To this aim, value nodes *IC1, IC2, IC3* (with negative utility values) are added in the DN to decision nodes. The *total expected cost* can be computed by considering a cost function which is the sum of the cost nodes and the impact node. This can be useful to evaluate the best set of countermeasures to activate given a set of observed attacks.

Example. Considering a cost of $40, 30$ and 50 for activating countermeasures *IDMD5, DDRR* and *FSR* respectively, Tab. 2 reports the results of the computation of the best set of countermeasures to activate, depending on observations concerning nodes *CR, TSNA* and *TCSR*. We notice that having countermeasure *FSR* inactive is a good option (the best one), only when we are sure that *CR* attack has not occurred.

ROI. An interesting aspect related to the selection of the best countermeasures concerns the so called *Return on Investment* (ROI) index (Roy, Kim, and Trivedi 2012). It represents the percentage of investment gain w.r.t. the investment cost. It is defined by comparing a status-quo situation (a set of countermeasures CM_{i-1}) and a target one (another set CM_i), differing from the status-quo in terms of a set of investments. By denoting with R_i the risk associated with CM_i and by C_i the cost of implementing CM_i from CM_{i-1}, then the ROI index is defined as:

$ROI(i) = \frac{R_{i-1} - R_i - C_i}{C_i}$

Example. Consider the situation CM_0 (no active countermeasure) corresponding to a risk (expected impact) of 106554; in case we do not have any evidence about attacks, solving the DN suggests that the best (in terms of risk) decision is CM_1 (all the countermeasures are activated) corresponding to a risk of 68395. Since implementing CM_1 has an investment cost $C_1 = 120$ we compute the ROI as $\frac{R_0 - R_1 - C_1}{C_1} = \frac{106554 - 68395 - 120}{120} = 316.99$, meaning that for each unit of investment, we get back about 317 units.

Conclusions

We have proposed DN as a reference model for the analysis of attack/defense scenarios. The advantages can be considered from both the modeling and the analysis point of view: uncertainty at every level of the scenario can be captured, probabilistic indices can be computed through standard inference, and a decision theoretic approach can be exploited to select the best set of countermeasures to activate. This has been shown through a case study concerning BGP.

References

Codetta-Raiteri, D.; Montani, S.; and Portinale, L. 2010. Supporting reliability engineers in exploiting the power of dynamic bayesian networks. *International Journal of Approximate Reasoning* 51(2):179–195.

Jensen, F., and Nielsen, T. 2007. *Bayesian Networks and Decision Graphs (2nd ed.)*. Springer.

Langseth, H., and Portinale, L. 2007. Bayesian networks in reliability. *Reliability Engineering and System Safety* 92(1):92–108.

Meng, F. 2000. Relationships of Fussell–Vesely and Birnbaum importance to structural importance in coherent systems. *Reliability Engineering and System Safety* 67:55–60.

Roy, A.; Kim, D.; and Trivedi, K. 2012. Scalable optimal countermeasure selection using implicit enumeration on attack countermeasure trees. In *International Conference on Dependable Systems and Networks*.

Schneier, B. 2000. *Secrets and Lies: digital security in a networked world*. J. Wiley.