



ELSEVIER

PAID: A Probabilistic Agent-Based Intrusion Detection system

Vaibhav Gowadia, Csilla Farkas*, Marco Valtorta

Information Security Laboratory, Department of Computer Science and Engineering,
University of South Carolina, Columbia, SC 29208, USA

Received 9 August 2004; revised 16 May 2005; accepted 16 June 2005

KEYWORDS

Intrusion detection;
Network security;
Computer security;
Computer attack;
Agents;
Bayesian Networks

Abstract In this paper we describe architecture and implementation of a *Probabilistic Agent-Based Intrusion Detection (PAID)* system. The PAID system has a cooperative agent architecture. Autonomous agents can perform specific intrusion detection tasks (e.g., identify IP-spoofing attacks) and also collaborate with other agents. The main contributions of our work are the following: our model allows agents to share their *beliefs*, i.e., the probability distribution of an event occurrence. Agents are capable to perform *soft-evidential update*, thus providing a continuous scale for intrusion detection. We propose methods for modelling errors and resolving conflicts among beliefs. Finally, we have implemented a proof-of-concept prototype of PAID.

© 2005 Published by Elsevier Ltd.

30
31
32
33
34
35
36
37
38
39
40

Introduction

As the complexity of computer systems increases and attacks against them become more and more sophisticated, high-assurance intrusion detection techniques need to be implemented. During the last two decades, many strategies and methods for intrusion detection have been developed (for a survey see [Axelsson, 2000](#)).

The main goal of any IDS is to detect all intrusions and only intrusions in an efficient way. Correctness of an IDS is measured by the rate of false positives and false negatives over all events.

A false positive warning occurs when a non-intrusive event is labelled intrusive. A false negative warning occurs when an intrusive activity is not detected. Negative effects of false positives include false accusations, reduced system availability, and subsequent disregard of IDS warnings. The negative effects of false negatives include reduced trust in IDS and damages caused by the intrusions. For effective intrusion detection, it is necessary that IDSs reduce the number of misclassifications and find an acceptable balance between false positive and false negative rates. [Dacier \(2002\)](#) found that most of the false positives are generated due to under-specified attack signatures, intent-guessing signatures, or lack of abstraction. Therefore, it is important to specify signatures

41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56

* Corresponding author.

E-mail address: farkas@cse.sc.edu (C. Farkas).

precisely and develop IDSs that can process these signatures efficiently.

Moreover, network-based distributed attacks are difficult to detect because their detection requires coordination among different intrusion detection components or systems (Snapp and Brentano, 1991; Neumann and Porras, 1999). Failure to recognize these attacks leads to false negatives. Therefore, the development of models and protocols for information sharing among intrusion detection components is critical. IDSs are often categorized as distributed or centralized. Spafford and Zamboni (2000) defined centralized IDSs as those where the analysis of data is performed at a fixed number of locations, independent of how many hosts are monitored. Distributed IDSs are defined as those where the analysis of the data is performed at a number of locations proportional to the number of hosts monitored. Centralized IDSs are able to use all available audit data to form a decision, but they create large communication overhead, require a powerful central processor, and represent a single point of failure. To overcome this problem, distributed IDSs process audit data at multiple locations. Distributed IDSs, like DIDS (Snapp and Brentano, 1991) and AAFID (Balasubramaniyan et al., 1998; Spafford and Zamboni, 2000), can share filtered raw data or binary (i.e., yes/no) decisions among their components. However, they cannot share probability distributions of intrusion beliefs. Moreover, existing distributed IDSs do not support selective sharing of published data among peers. In this work, we propose a middle ground between centralized and distributed IDSs, where each IDS component shares its data or results only with those agents that subscribe for these results.

We believe that precise representation of attack signatures in probabilistic intrusion detection model requires: (1) ability to process observations (hard findings) and beliefs (probability distributions) about system parameters, and (2) flexibility in specifying threshold value of probability, above which an alarm is generated.

In this paper, we focus on distributed IDSs based on Bayesian technology and multiagent technology. IDSs based on Bayesian technology may allow sharing of raw data and results (probability distributions) among IDS components. A *Bayesian Network* (BN) is a graphical representation of the joint probability distribution for a set of discrete variables. The representation consists of a directed acyclic graph (DAG). Nodes of the DAG represent variables and edges represent cause-effect relations. The strength of each effect is modelled as a probability. These probabilities are represented

by a conditional probability table (CPT). CPTs specify conditional probability of the variable given its parents. For variables without parents, this is an unconditional distribution. Inference in Bayesian Network means computing the conditional probability for some variables given information (evidence) on other variables.

The traditional Bayesian inference (Jensen, 2001; Pearl, 1988) can be performed only with hard findings or observations as input. However, in intrusion detection scenarios modelled with Bayesian Networks, we often find that the input variables cannot be measured directly. Only a belief (probability distribution) in the current state of these variables may be computed. A typical example of such input is a probabilistic result computed by another IDS component. To accept results of other IDS components, existing IDSs (DuMouchel, 1999; Valdes and Skinner, 2000; Sebyala et al., 2002; Cho and Cha, in press) based on the traditional Bayesian inference technique have to coerce the result into a binary decision. Such coercions are performed by assuming occurrence (or non-occurrence) of represented event if the input probability is greater (or smaller) than a threshold value. We believe that IDSs that utilize such binary decisions have limited flexibility and have difficulty in removing false positives and false negatives.

We illustrate our above observation by a simple example given in Fig. 1. Fig. 1(a) shows a Bayesian Network that accepts two inputs: *B* (hard finding) and *C* (soft finding), and computes belief in *A*. Fig. 1(b) shows conditional probability tables for the example BN. Fig. 1(c) shows the likelihood of *A* being in "abnormal" state with respect to the likelihood of *C* being in "abnormal" state. Likelihood of *C* is computed with the two methods, with and without coercion. In both calculations we have assumed that *B* is observed to be in "abnormal" state. We observe that the likelihood graph of *A* is continuous when soft-evidential update is used. In this case the security officer has large flexibility in choosing a warning threshold for *A*. We also observe that the likelihood graph of *A* is not continuous with traditional probability update. Moreover, the likelihood of *A* has only discrete values that depend on the threshold set for *C*. We have developed a Bayesian Network-based technique that allows the IDS components to share results of their analysis in the form of beliefs. Such sharing enables our model to perform intrusion detection on a continuous scale.

Agents are software systems that function autonomously to achieve desired objectives in their environment. Recent research (Spafford and

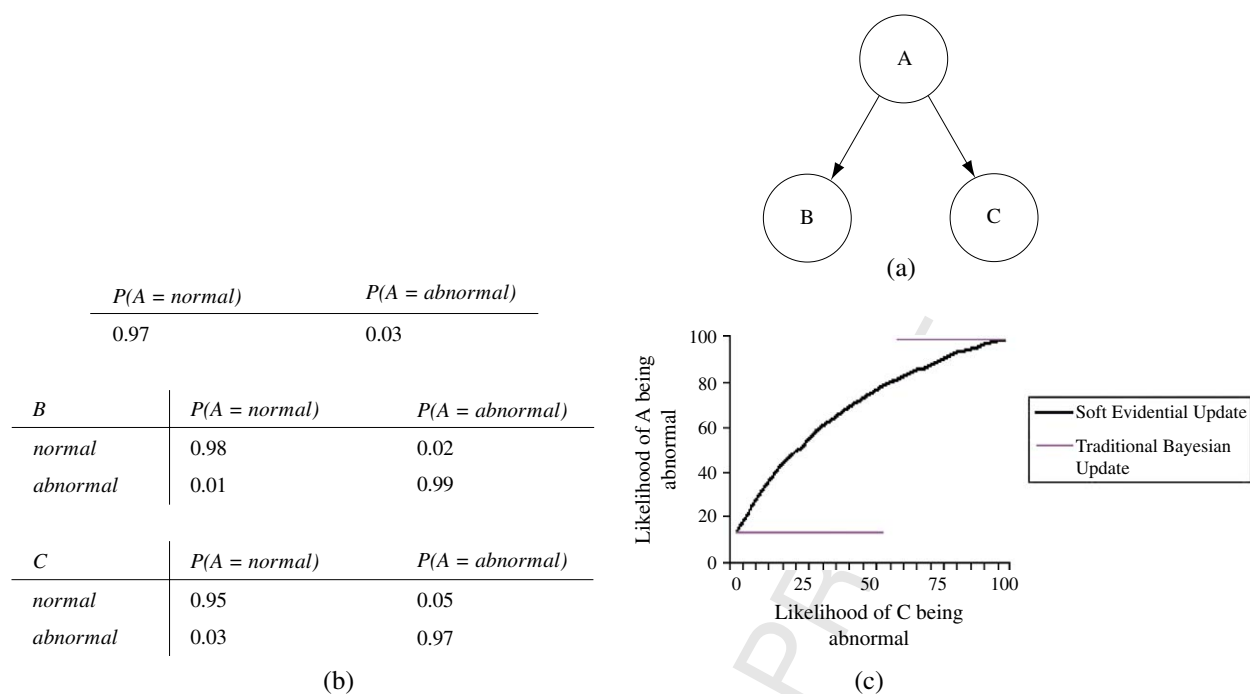


Figure 1 (a) Example BN, (b) CPTs for the example BN, and (c) likelihood of A being abnormal calculated with soft-evidential update and traditional probability update ($C_{threshold} = 60\%$).

169 Zamboni, 2000; Jansen et al., 1999; Carver et al.,
 170 2000; Helmer et al., 2003) shows that agent-based
 171 technology seems to be a promising direction for
 172 developing collaborative intrusion detection systems.
 173 We propose an agent-based, cooperative
 174 architecture where each IDS component is able to
 175 process its own data and to integrate local findings
 176 with the findings of other IDS components. Each
 177 agent acts as a wrapper for a Bayesian Network.
 178 Agents in our model utilize the communication
 179 protocols and languages (Bellifemine et al., 1999;
 180 FIPA, 2002a) developed by multiagent research
 181 community. In addition, they use Bayesian inference
 182 with soft-evidential update to support integration
 183 of beliefs and observations. We refer to this
 184 multiagent architecture as Agent Encapsulated
 185 Bayesian Network (AEBN) (Bloemeke and Valtorta,
 186 2002). Although, Bayesian Network-based archite-
 187 ctures have been considered for intrusion detection
 188 (DuMouchel, 1999; Valdes and Skinner, 2000; Bar-
 189 bara et al., 2001; Sebyala et al., 2002; Cho and Cha,
 190 in press), these models use traditional probability-
 191 update methods (Jensen, 2001; Pearl, 1988). They
 192 can effectively utilize only those parameters that
 193 result from an actual measure. This limitation
 194 often results in under-specified signatures. To solve
 195 this problem, in our model agents are enabled to
 196 share their beliefs (*soft findings*) in addition to
 197 measured values (*hard findings*).

More specifically, we propose an agent-based
 and cooperative architecture, called Probabilistic
 Agent-Based Intrusion Detection (PAID), to analyze
 system information and estimate intrusion proba-
 bilities. Agents in PAID accept facts and derived
 values as inputs. Agents may share their beliefs
 with, or request information (belief or data) from
 the other agents.

Our model uses three types of agents: system-
 monitoring agents, intrusion-monitoring agents
 and registry agents. System-monitoring agents
 are responsible for collecting, transforming, and
 distributing intrusion specific data upon request
 and evoking information collecting procedures.
 Each intrusion-monitoring agent encapsulates
 a Bayesian Network and performs belief update
 as described in Valtorta et al. (2002) using both
 facts (observed values) and beliefs (derived val-
 ues). Intrusion-monitoring agents generate proba-
 bility distributions (beliefs) over intrusion
 variables that may be shared with other agents.
 Each belief is called a soft finding. Soft findings can
 indicate that a system is in an abnormal state.
 Even in the absence of hard findings, soft findings
 can affect the probability of intrusion occurrence
 or attack against the monitored system. A proba-
 bilistic representation of attacks, using hard and
 soft findings makes our model capable of identify-
 ing variations of known intrusions. Coordination

227 between system-monitoring and intrusion-monitoring agents is provided by a registry agent.
 228 Within an IDS collaborative model, there may exist
 229 several registry agents that, upon failure, can
 230 compensate for each other. However, each intrusion-monitoring and system-monitoring agent is
 231 registered with a registry agent central to the
 232 monitoring agents.

235 Currently our model detects known intrusions by
 236 using well-documented patterns of attacks. Each
 237 intrusion-monitoring agent is looking for a particular
 238 intrusion pattern. If such a pattern is found, a possible intrusion is indicated. Distributed intrusion
 239 detection is achieved by enabling agents to share
 240 their beliefs. Depending on the level of collaborations and privacy concerns of the collaborating
 241 entities, each component may be able to build the
 242 full, global decision stage or only a partial one.

245 The organization of this paper is as follows. Next
 246 section gives a brief introduction to Bayesian
 247 Networks, and agent technology. Then, background information and related work followed by
 248 the description of the proposed framework (PAID)
 249 are given. Further, methodology for building BNs
 250 for intrusion detection is presented which is
 251 followed by implementation of the proposed intrusion detection framework. Finally, we conclude
 252 and recommend future research in last section.

255 Background

256 Bayesian Networks

257 A *Bayesian Network* (BN) is a graphical representation of the joint probability distribution for a set
 258 of discrete variables. The representation consists
 259 of a directed acyclic graph (DAG), prior probability
 260 tables for the nodes in the DAG that have no
 261 parents and conditional probability tables (CPTs)
 262

for the nodes in the DAG given their parents. As an
 263 example, consider the network in Fig. 2.

265 More formally, a Bayesian Network is a pair
 266 composed of: (1) a multivariate probability distribution over n random variables in the set $V =$
 267 V_1, \dots, V_n , and (2) a directed acyclic graph (DAG)
 268 whose nodes are in one-to-one correspondence
 269 with V_1, \dots, V_n . (Therefore, for the sake of convenience, we do not distinguish the nodes of a graph
 270 from variables of the distribution.)

273 Bayesian Networks allow specification of the
 274 joint probability of a set of variables of interest in
 275 a way that emphasizes the qualitative aspects of
 276 the domain. The defining property of a Bayesian
 277 Network is that the conditional probability of any
 278 node, given any subset of non-descendants, is
 279 equal to the conditional probability of that same
 280 node given the parents alone. The chain rule for
 281 Bayesian Networks (Neapolitan, 1990) given below
 282 follows from the above definition.

283 “Let $P(V_i | \pi(V_i))$ be the conditional probability of
 284 V_i given its parents. (If there are no parents for V_i ,
 285 let this be $P(V_i)$.) If all the probabilities involved
 286 are nonzero, then $P(V) = \prod_{v \in V} P(v | \pi(v))$ ”.

287 Three features of Bayesian Networks are worth
 288 mentioning. First, the directed graph constraints
 289 the possible joint probability distributions represented by a Bayesian Network. For example, in any
 290 distribution consistent with the graph of Fig. 2, D is
 291 conditionally independent of A given B and C . Also,
 292 E is conditionally independent of any subset of the
 293 other variables given C .

295 Second, the explicit representation of constraints
 296 about conditional independence allows
 297 a substantial reduction in the number of parameters to be estimated. In the example, assume that
 298 the possible values of the five variables are as
 299 shown in Fig. 2(b).

301 Then, the joint probability table $P(A, B, C, D, E)$
 302 has $2 \times 3 \times 2 \times 4 \times 4 = 192$ entries. It would be

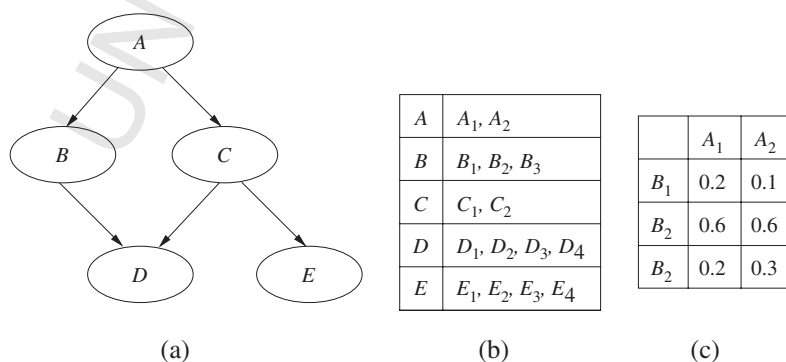


Figure 2 (a) An example Bayesian Network, (b) variable states, and (c) conditional probability table for B given A .

303 very difficult to assess 191 independent parame- 357
 304 ters. However, the independence constraints en- 358
 305 coded in the graph permit the factorization $P(A, B,$ 359
 306 $C, D, E) = P(A) \times P(B | A) \times P(C | A) \times P(D | B,$ 360
 307 $C) \times P(E | C)$ which reduces the number of param- 361
 308 eters to be estimated to $1 + 4 + 2 + 18 + 6 = 31$. 362
 309 The second term in the sum is the table for the 363
 310 conditional probability of B given A . This probability 364
 311 is shown in Fig. 2(c); note that there are only four 365
 312 independent parameters to be estimated since the 366
 313 sum of values by column is one. 367

314 Thirdly, the Bayesian Network representation 368
 315 allows a substantial (usually, dramatic) reduction 369
 316 in the time needed to compute marginals for each 370
 317 variable in the domain. The explicit representation 371
 318 of constraints on independence relations is ex- 372
 319 ploited to avoid the computation of the full joint 373
 320 probability table in the computation of marginals 374
 321 both prior and conditioned on observations. Limi- 375
 322 tation of space prevents the description of the 376
 323 relevant algorithms; see Jensen (2001) for a dis- 377
 324 cussion of the junction tree algorithm. 378

325 The most common operation on a Bayesian 379
 326 Network is the computation of marginal probabil- 380
 327 ities both unconditional and conditional upon 381
 328 evidence. Marginal probabilities are also referred 382
 329 as *beliefs* in the literature (Pearl, 1988). This 383
 330 operation is called probability updating, belief 384
 331 updating, or belief assignment. 385

332 We define evidence as a collection of findings. A 386
 333 (*hard*) *finding* specifies which value a variable is in. 387
 334 A *soft finding* specifies the probability distribution 388
 335 of a variable. These definitions of finding and of 389
 336 evidence may be generalized, for example, by 390
 337 allowing specifications of impossible configurations 391
 338 of pairs of variables (Cowell et al., 1999; Lauritzen 392
 339 and Spiegelhalter, 1988; Valtorta et al., 2002). 393
 340 However, applications rarely need the power of the 394
 341 more general definitions, and most Bayesian 395
 342 Network software tools support only the definition 396
 343 of (hard) evidence as a collection of (hard) findings 397
 344 given here. 398

345 Agent Encapsulated Bayesian Networks

346 Although there is no universally accepted defini- 400
 347 tion of agent, most authors agree that agents share 401
 348 the following properties: each agent is auton- 402
 349 omous, has a set of goals, and has a local model of 403
 350 the part of the world that affects the achievement 404
 351 of its goals, and has a way of communicating with 405
 352 other agents. In an Agent Encapsulated Bayesian 406
 353 Network (AEBN) (Bloemeke and Valtorta, 2002), 407
 354 each agent uses a single Bayesian Network (which 408
 355 is also called an AEBN) as its model of the world. 409
 356 The agents communicate via passing messages that 410

are distributions on variables shared between the 357
 individual networks. 358

The variables of each AEBN are divided into 359
 three groups: those about which other agents have 360
 better knowledge (input set), those that are used 361
 only within the agent (local set), and those for 362
 which the agent has the best knowledge and which 363
 the other agents may want to use (output set). The 364
 variables in the input set and the output set are 365
 shared with other agents. The variables in the 366
 local set are not. An agent subscribes to zero or 367
 more variables in the input set and publishes zero 368
 or more variables in the output set. 369

The mechanism for integrating the view of the 370
 other agents on a shared variable is to replace 371
 the agent's current belief (which is a probability 372
 distribution) in that variable with that of the 373
 communicating agent. The update of a probability 374
 distribution represented by a Bayesian Network 375
 upon receipt of a belief is called a soft-eviden- 376
 tial update and is explained in detail by Valtorta 377
 et al. (2002). In this work, we have used the Big 378
 Clique algorithm for soft-evidential update, im- 379
 plemented in the BC–Hugin system (Kim et al., 380
 2004). 381

When a publisher makes a new observation, it 382
 sends a message to its subscribers. The subscribers 383
 in turn adjust their internal view of the world and 384
 send their published values to their subscribers. 385
 Assuming that the graph of agent communication 386
 (which we simply call agent graph) is a directed 387
 acyclic graph (DAG), equilibrium is reached, and 388
 a kind of global consistency is assured because the 389
 belief in each shared variable is the same in agents 390
 that subscribe to that variable. 391

The restriction that an agent has correct and 392
 complete knowledge of the variables it publishes 393
 forces unidirectional communication, and it may 394
 seem excessive. However, there is a good reason to 395
 insist on this requirement. The alternative (i.e., 396
 to allow bidirectional communication between 397
 agents) requires that the agent graph be a tree, 398
 as shown in Xiang (2002). Most agent-based sys- 399
 tems demonstrate acyclic graph communication 400
 model. For example, it is possible to have multiple 401
 views of the same parameter. That is, two agents 402
 may publish variables that correspond to their 403
 measurement (or belief) of the same parameter. 404
 Moreover, nothing prevents another agent from 405
 integrating the published values of these two 406
 agents, thus obtaining a new (and possibly more 407
 accurate) view of the parameter. 408

Table 1 summarizes some features of AEBNs and 409
 other related representation formalisms. AEBNs 410
 have very good scalability and shared variables 411
 are independent of variables in descendant BNs. 412

Table 1 Agent Encapsulated Bayesian Networks and related representation formalisms

Name	Granularity	Topological restrictions	Constraints on independence relations	Purpose	Scalability
Bayesian Network (Jensen, 2001)	Individual variable	DAG of variables	Local Markov condition (<i>d</i> -separation)	Efficient representation of multivariate probability distribution	Poor
Multiply Sectioned Bayesian Network (MSBN) (Xiang, 2002)	Bayesian Network (BN)	Tree (of BNs)	<i>d</i> -Separation on composition of BNs;	Efficient distribution of computation among processors	Good: distributed computation, if tree decomposition is possible
Multiple Entity Bayesian Networks (MEBN) (Laskey et al., 2001)	Bayesian Network Fragments (BNFragments)	DAG (of BNFragments)	<i>d</i> -Separation on composition of BNs; encapsulation	Distributed representation of Bayesian Networks	Mediocre: representation decomposed, computation centralized
Agent Encapsulated Bayesian Networks (AEBN) (Bloemeke and Valtorta, 2002)	Bayesian Network (BN)	DAG (of BNs)	Shared variables independent of variables in descendent BNs given parent BNs; encapsulation	Construction of interpretation models by collaborating agents	Very good: distributed computation, distributed representation
Decentralized Sensing Networks (DSN) (Utete, 1998)	Sensor	Undirected graph (of sensors)	None: non-probabilistic approach	Distributed sensing and data fusion	Poor: rumor problem is unsolvable in DSNs

413 Therefore, we chose to use the AEBN organization
414 for the work described in this paper.

415 We now briefly overview related research work
416 on intrusion detection with the help of Bayesian
417 networks or agent technology.

418 Bayesian Networks based intrusion 419 detection

420 IDSs using Bayesian Networks have been proposed
421 by many researchers (DuMouchel, 1999; Valdes and
422 Skinner, 2000; Barbara et al., 2001; Sebyala et al.,
423 2002; Cho and Cha, in press). However, these IDS
424 models use only hard findings in their Bayesian
425 models. We now briefly overview their IDS archi-
426 tectures.

427 DuMouchel (1999) proposed an anomaly detec-
428 tion technique using the Bayes classifier. They
429 keep a profile of commands issued by each user
430 and compute command transition probabilities.
431 Their IDS detects abnormal behavior based on the
432 observed command transitions.

433 Valdes and Skinner (2000) proposed an adaptive
434 model that detects attacks using probability theory.

Their architecture analyzes the traffic from a given 435
client's TCP sessions. This analysis is done by 436
Bayesian inference at periodic intervals in a session, 437
and the interval is measured in number of events or 438
elapsed time. Between inference intervals, the 439
system state is propagated according to a Markov 440
model. After each inference, the system may give 441
alerts for suspicious sessions. 442

Sebyala et al. (2002) have incorporated Bayesian 443
Network in their IDS as anomaly detector. They 444
keep a profile of CPU and memory utilization by 445
proxylets in active networks. They use a Bayesian 446
Network to compute state (good or bad) of 447
proxylet. A proxylet is in bad state if the CPU 448
and memory utilization is anomalous. 449

Cho and Cha (in press) proposed a technique to 450
detect anomalies in web sessions. A web session 451
consists of sequence of page requests. Anomalous 452
request in given web session may correspond to 453
request for secured pages without accessing the 454
login page, or repeated access to a same page. 455
Their model utilizes Bayesian parameter estima- 456
tion technique (Friedman and Singer, 1998) to 457
compute probability that a user may request 458

459 certain pages in given sequence. A web session
460 may consist of multiple sub-sessions. To combine
461 the anomaly scores of these sub-sessions, they
462 suggest use of either maximum value (for high
463 sensitivity) or average value (for low sensitivity).

464 In the above models, Bayesian inference is
465 performed when a hard evidence is received like
466 a command sequence, TCP parameters, CPU or
467 memory utilization, or a page request. None of
468 the above models describe Bayesian model for
469 attacks when the input observation is a probability
470 distribution over the states of a system parameter.
471 For example, there is 80% chance of a DOS attack on
472 the file server, or there is a 70% chance for command
473 sequence to be anomalous. Unlike existing models,
474 our model allows accepts beliefs (probability dis-
475 tributions) as input to the Bayesian models.

476 Agent-Based Intrusion Detection systems

477 Agent-based systems require a communication in-
478 frastructure. Agents in our system communicate
479 with each other by sending messages in the Agent
480 Communication Language specified by FIPA (FIPA,
481 2002a,b). JADE (Bellifemine et al., 1999) is a soft-
482 ware framework to aid the development of agent
483 applications in compliance with the FIPA specifi-
484 cations for inter-operable intelligent multiagent
485 systems. The purpose of JADE is to simplify de-
486 velopment while ensuring standard compliance
487 through a comprehensive set of system services
488 and agents. To achieve such a goal, JADE offers
489 a distributed agent platform, directory facilitator
490 (DF), and library of interaction protocols. The
491 agent platform includes an agent management
492 system that allows monitoring and logging of agent
493 activities and performs life-cycle operations
494 (start, suspend, and terminate) on agents. Inter-
495 action protocols (e.g., request, query, subscribe,
496 etc.) are used to design agent's interaction, pro-
497 viding a sequence of acceptable messages and
498 semantics for those messages.

499 Agent communications can be divided into two
500 categories: communication among agents at the
501 same host and communication among agents at
502 different hosts. Balasubramaniyan et al. (1998)
503 examine these methods in the context of intrusion
504 detection.

505 Spafford and Zamboni (2000) and Balasubrama-
506 niyan et al. (1998) presented a framework called
507 AAFID in which autonomous agents report their
508 findings to entities called transceivers. Each host
509 has a unique transceiver that collects information
510 from all other agents on its host machine. Agents
511 also perform data reduction and send data to

512 monitors that oversee the operation of several
513 transceivers. Monitors have the capability to de-
514 tect events that may be unnoticed by the trans-
515 ceivers. In mobile agent-based systems, like the
516 ones presented by Helmer et al. (2003) and Asaka
517 et al. (1999), mobile agents collect, integrate, and
518 analyze data from different components of a dis-
519 tributed system. The agent's findings are recorded
520 in a database and/or reported to the users.

System design goals

521

522 One of the design goals of our IDS is to enable it to
523 function as a stand-alone system or to support
524 existing IDSs. Our main goal is to improve upon
525 existing IDS technologies by allowing flexible in-
526 formation sharing among system components in
527 a way that the shared data are easily incorporated
528 in the analysis of the components. Our model
529 supports the calculation of intrusion probabilities
530 on a continuous scale of [0, 1]. A probability of
531 zero means it is certain no intrusion has occurred,
532 and one means that an intrusion has definitely
533 occurred. For each intrusion type there is an
534 associated variable that represents the probability
535 of that intrusion. Each Bayesian network is able to
536 modify its own belief (probability distribution over
537 an intrusion variable) and to import or export
538 beliefs from or to other Bayesian networks. These
539 input variables are accepted during all states of
540 processing.

541 Analysis of distributed attacks on a large net-
542 work may require monitoring of numerous hosts
543 and large volumes of network traffic. Thus a large
544 amount of data is generated that must be ana-
545 lyzed. Our model supports local analysis of col-
546 lected data and sharing of results (and partial
547 results). We also allow agents to share probability
548 distributions (beliefs) of intrusion occurrences and
549 system states. This "belief sharing" carries more
550 information than sharing a binary decision and also
551 has a lower overhead than raw data sharing.

552 Each intrusion-monitoring site or network may
553 have different sensitivity and selectivity require-
554 ments. Our model allows security officers to
555 customize these parameters according to the local
556 requirements. This customization does not affect
557 the probability distribution values shared among
558 the agents.

559 Finally, we address some of the issues related to
560 reliability and ease of maintenance. Based on the
561 distributed nature of our model and the possibility
562 of replicated Bayesian Networks for monitoring
563 intrusion, our model remains functional even if
564 some of the IDS network nodes are unavailable.

565 Since the detection of an intrusion is based on
 566 several parameters, including local findings and
 567 findings from several other agents. The misleading
 568 data from compromised agents are small if the
 569 number of non-compromised agents is large and
 570 the number of compromised agents is small. Each
 571 Bayesian Network is responsible to monitor a par-
 572 ticular intrusion; therefore the modification of an
 573 intrusion pattern will affect only those networks
 574 that monitor the intrusion. Similarly, protection
 575 against new types of attacks can be added easily to
 576 the model.

577 Probabilistic Agent-Based Intrusion 578 Detection

579 In our model, we use agent graphs to represent
 580 intrusion scenarios. Each agent is associated with
 581 a set of input variables, a set of output variables,
 582 and a set of local variables. The agent at each
 583 node of the graph encapsulates a Bayesian Net-
 584 work. Nodes of the Bayesian Network are variables
 585 that represent suspicious events, intrusions, or
 586 system and network parameter values. A variable
 587 can have any number of states, and the belief in
 588 the variable is the distribution on its states. The
 589 encapsulated Bayesian Network is used to model
 590 intrusion scenarios. It is also able to incorporate
 591 measurement errors and handle multiple beliefs
 592 on input variables.

593 PAID architecture

594 The PAID architecture uses agent technology to
 595 collect and analyze data and to distribute infor-
 596 mation among the PAID components. PAID supports
 597 three types of agents: (1) system-monitoring
 598 agents, (2) intrusion-monitoring agents, and (3)
 599 registry agents.

- (1) *System-monitoring agents*: The system-monitoring agents perform either online or offline processing of log data, communicate with the operating system, and monitor system resources. These agents publish their output variables (facts and beliefs derived from observations) that can be utilized by other agents.
- (2) *Intrusion-monitoring agents*: Each intrusion-monitoring agent computes the probability for a specific intrusion type. These agents subscribe to variables and/or beliefs published by the system-monitoring agents and other intrusion-monitoring agents. The probability values for each agent are updates, calculated

according to the values of input variables and beliefs.

- (3) *Registry agent*: The registry agent maintains information about the published variables and monitored intrusions for each system-monitoring or intrusion-monitoring agent. It is required that all agents of PAID must register with the registry agent. The registry agent also maintains the location and current status of all the registered agents. Agent status is a combination of two parameters *alive* and *reachable*. The status of a communication link between any two agents is determined by attempting to achieve a reliable UDP communication between them. The registry agent is used to find information (e.g., name and location) about agents who may supply required data. The PAID architecture can support multiple registry agents also as described later in section 'Scalability and complexity analysis'. For simplicity, we describe the examples with a single registry agent.

Agent communication

635

The interactions among the components of PAID
 are shown in Fig. 3. The messages are sent in XML
 syntax (Bray et al., 2001) among the agents. These
 messages correspond to registration requests, in-
 formation requests and other agent actions. A
 brief overview of agent actions and the corre-
 sponding messages is given below.

636
637
638
639
640
641
642

1. *Registration of an agent with the registry agent*: Each agent in PAID must register with

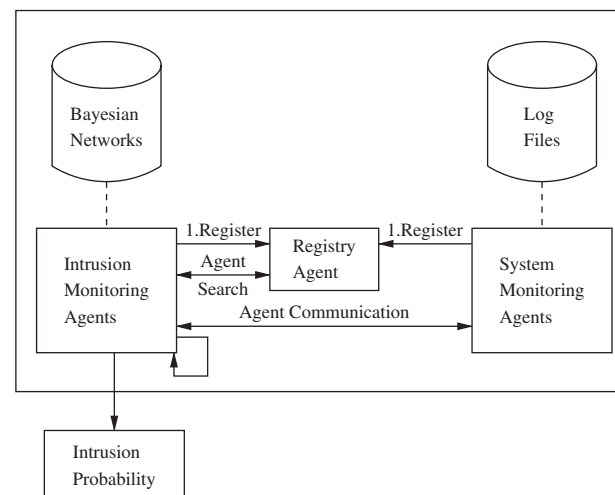


Figure 3 Probabilistic Agent-Based Intrusion Detection (PAID).

the registry agent. A registration message includes the registering agent's agent-id, IP address, list of published variables and their possible states, digital signature, and digital certificate. The registry agent issues an acknowledgment message upon successfully entering the new agent in its database.

2. *Information request by an agent about other agents:* Each intrusion-monitoring agent has a set of input variables (determined from the encapsulated Bayesian Network). To find agents capable of providing required input data, the intrusion-monitoring agent sends a search request to the agent registry. The search request includes the requester's agent-id, IP address, and the required input variables. The message is digitally signed by the requester.
3. *Registry agent's reply to an information request:* Upon receiving a search request, the registry agent verifies that the request is legitimate before searching its database to determine which agents can supply the requested variables and the status of these agents. The message from the registry agent to the requester includes the requested variable name, the agent-id of the agent publishing the variable, its IP address, and status. The message is digitally signed by the registry agent.
4. *Request for belief subscription:* Upon receiving the list of agents capable of providing the required input from the registry, the subscribing agent sends requests directly to these agents. A subscription request consists of the requester's agent-id, requester's IP address, requested input variable name, the duration of subscription time, the desired time interval between subsequent updates, a request-id, and the timestamp of the request. The message is digitally signed by the requester.
5. *Belief-update messages:* Upon receiving a belief subscription request the publishing agent sends regular updates within the agreed intervals and duration of the subscription. The message contains the request-id, the sender's id, and the probability distribution of the requested variable. The message is signed by the publisher.

update of the active agents. We use secret key encryption for message content to reduce encryption overhead. Message and agent authentication is guaranteed by public-key cryptosystem and the use of digital certificate. Each message is digitally signed by the sending agent. In addition, we require that agents authenticate themselves to the registry by their digital certificates.

Status probes of registered system-monitoring agents and network links are periodically performed by the registry agent. Responses to the probing messages carry information about the state of the system-monitoring and intrusion-monitoring agents. The status of a communication link between any two agents is determined by attempting to achieve a reliable UDP communication between them. Compromised agents can be identified by periodically launching attacks over the monitored network and verifying that the expected results are generated. This approach was proposed by [Dacier \(2002\)](#).

Scalability and complexity analysis

The factors affecting the scalability of our model are the costs of data transfer, belief updates and registry operations (i.e., register, deregister, and query). During normal operation, agents share their beliefs; thus, PAID has a low bandwidth requirement. Sharing of data or partial data is required only to analyze suspicious events.

[Pearl \(1988\)](#) has shown that belief update can be performed in linear time in trees and (more generally) singly connected networks. Unfortunately, belief update in general Bayesian Networks is NP-hard ([Cooper, 1990](#)). The computational complexity of the algorithm found to be the best in practice, the junction tree algorithm, is exponential in a graphical parameter called the tree-width of the Bayesian Network. This negative result holds, even for some notions of approximation and for many restrictions on the structure of the Bayesian Network. Despite these negative theoretical results, update in most Bayesian Networks, using the junction tree algorithm [Lauritzen and Spiegelhalter \(1988\)](#) is very fast because most practical Bayesian Networks compile (after an intermediate step that converts them into an undirected graph) into a junction tree where the largest clique is small. The process is described in detail in the literature, for example in [Neapolitan \(1990\)](#). More precisely, the computational complexity of the junction tree algorithm, which is widely found to be the fastest algorithm in

693 Communication security and reliability

694 Reliable and secure communication is achieved by
695 using commercially available encryption techniques
696 to achieve communication security and authentica-
697 tion. Reliability is supported by periodic status

698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718

719

720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749

750 practice is exponential in a graphical parameter
751 called the treewidth of the Bayesian Network.

752 PAID can provide scalability by supporting mul-
753 tiple registries. Each subnet may have its own
754 agent registry. The agent-registries can forward
755 requests and replies to neighboring registries
756 based on the IP address of the receiving agent.
757 Dynamic routing algorithms for IP networks (Moy,
758 1997; Perlman, 1992) are applicable for this
759 purpose.

760 Modelling Bayesian Networks for PAID

761 To assess the probability of an intrusion, we use
762 Bayesian Networks. Modelling a domain with
763 Bayesian Networks involves two major steps. First,
764 a domain expert needs to specify the qualitative
765 structure of the network, which depends solely on
766 the independence relation among the variables of
767 the domain of interest. Second, the numerical
768 parameters need to be assessed; these parameters
769 are the prior probabilities of variables that have no
770 parents, and the conditional probabilities of every
771 other variable given its parents. The graph and the
772 probabilities uniquely and completely specify the
773 joint probability of the variables in the domain of
774 interest.

775 We now present methodologies for modelling
776 Bayesian Networks for attack patterns, system
777 parameters, incorporating errors, and resolving
778 conflicts.

779 Bayesian Network building methodology

780 There are two methods of building Bayesian Net-
781 works for a particular application domain. The first
782 method consists of asking the domain expert to
783 construct the network (DAG) and assess the prior
784 and conditional probabilities manually. This is how
785 we build our networks. The second method builds
786 the network from data. There are several algo-
787 rithms available to accomplish this learning task.
788 These are: BIFROST (Lauritzen and Spiegelhalter,
789 1988), K2 (Cooper and Herskovits, 1992), and CB
790 (Singh and Valtorta, 1993, 1995). The prior and
791 conditional probabilities can also be computed
792 from data. The models are validated by comparison
793 with the performance of an expert (Spiegelhalter
794 et al., 1993; Neapolitan, 2004). We plan to extend
795 our model to incorporate these algorithms to build
796 Bayesian Networks.

797 We now illustrate our method of building Bayes-
798 ian Network model for attack patterns, with an
799 example of a Mitnick attack.

Modelling computer attacks with Bayesian Networks: an example 800 801

802 The Mitnick attack (see Fig. 4) is difficult to
803 identify due to its distributed nature. In a network
804 vulnerable to Mitnick attack, the victim host
805 authenticates a trusted host using an IP address
806 only. The trust relationship between the victim
807 and trusted hosts implies that the users logged in
808 on the trusted host or applications running on the
809 trusted host can access resources on the victim
810 host without secure authentication. Mitnick attack
811 exploits the weakness of IP based authentication
812 systems and a flaw in TCP packet sequence number
813 generation algorithm. An attacker launches a dis-
814 tributed denial of service attack on the trusted
815 host making it temporarily unavailable. The at-
816 tacker is then able to gain access to the victim host
817 by pretending to be a user from the trusted host.
818 Hence, the identification of a Mitnick attack
819 requires evidence of both IP spoofing and DOS
820 attacks on different machines in the victim net-
821 work. Soft and hard findings detected on the
822 victim's network can be used to identify the
823 attack. We now examine the Mitnick attack in
824 detail and model possible findings and their de-
825 pendencies with a Bayesian Network model.

826 In preparation for the attack, the intruder
827 installs malicious programs (zombies) on many
828 vulnerable computers over the Internet. Mean-
829 while, the intruder gathers information about the
830 real victim. This information will allow the attacker
831 to successfully guess TCP sequence numbers of
832 the victim host. At a specific time, the attacker
833 activates the zombies to launch a denial of service
834 (DOS) attack against the host trusted by the victim.
835 As a result, the trusted host is unable to reply to
836 packets sent by the victim host. Under such
837 a situation, the intruder tries to open a TCP
838 connection with the victim host by spoofing the IP
839 address of the trusted host. The victim host sends

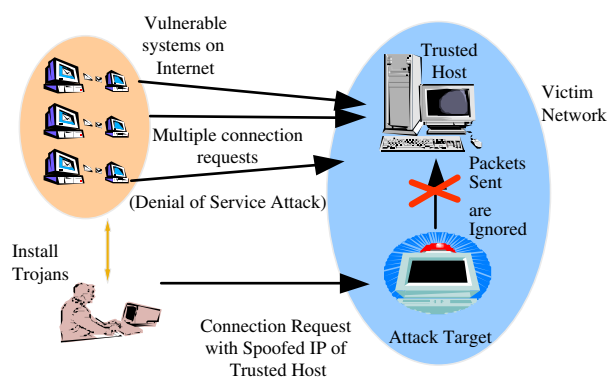


Figure 4 Mitnick attack.

840 an SYN-ACK packet as a reply to the trusted host.
 841 The trusted host drops this packet due to the DOS
 842 attack. The attacker now sends an ACK packet to
 843 the victim with appropriate sequence number.
 844 Upon receipt of this packet, the victim computer
 845 assumes that the other party is the trusted host.
 846 The attacker is now in a position to use the services
 847 provided by the victim host. We observe that the
 848 belief of a Mitnick attack's occurrence depends on
 849 the belief in occurrence of an IP-spoofing attack
 850 and a DOS attack on a trusted server. We add these
 851 dependencies to the qualitative structure (Fig. 6)
 852 of Mitnick attack's Bayesian model.

853 During an IP-spoofing attack, the network fire-
 854 wall and/or routing nodes on a network may
 855 discover incoming IP-packets with local source IP
 856 addresses. There is a high probability of an IP-
 857 spoofing attack when such packets are observed.
 858 This finding can be detected directly by a network
 859 device. The output of network device may say that
 860 either such packet was observed, or was not
 861 observed. Therefore, we model this finding as
 862 a hard finding (called *Local-SrcOnExternalInter-*
 863 *face*) supporting the hypothesis of an IP-spoofing
 864 attack. Another possible evidence of an IP-spoofing
 865 attack is an abnormal variation of the TTL value in
 866 the IP header of packets from a trusted host. This
 867 variation can be detected against recorded TTL
 868 values from the same host. However, such an
 869 observation is subjective and there is no clear
 870 demarcation between normal and abnormal val-
 871 ues. We model this type of finding as a soft finding,
 872 and represent them by probability distributions
 873 over normal and abnormal states.

874 In a situation when a server is under a TCP-SYN
 875 attack, the server receives a greater number of
 876 SYN packets than the number of connections it can
 877 handle. An increased ratio of SYN to SYN-ACK
 878 packets when observed along with decreased out-
 879 going data packets increases the probability of
 880 a TCP-SYN DOS attack and distinguishes it from the
 881 normal busy hours of the day.

882 During a Mitnick attack, applications at a victim
 883 host are not able to open new connections with the
 884 trusted host, but the victim host is still receiving

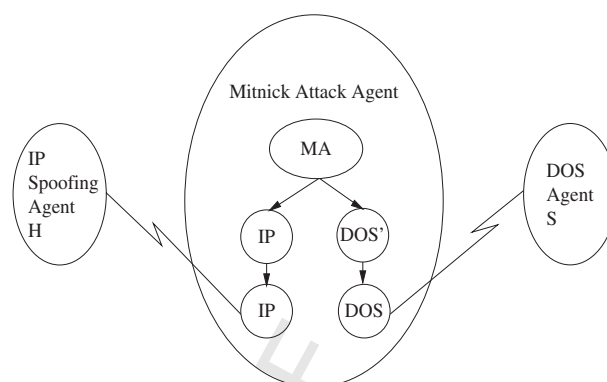


Figure 5 Multiagent system for Mitnick attack.

885 packets with the source IP address of the trusted
 886 host. The former observation by itself may not be
 887 sufficient to make an inference about the Mitnick
 888 attack, but is very useful when combined with
 889 other information. We model this observation as
 890 a soft finding called *OneWayCommunication*.

891 Fig. 5 shows the agent graph to model a simple
 892 Mitnick attack. There are three intrusion-monitor-
 893 ing agents: IP Spoofing Agent, Mitnick Attack
 894 Agent, and DOS Agent. The Mitnick Attack
 895 Agent subscribes to the beliefs of the IP Spoofing
 896 Agent and the DOS Attack Agent. Note that the IP
 897 Spoofing Agent and DOS Attack Agent may further
 898 subscribe to beliefs published by other agents as
 899 suggested by the Bayesian Network of Mitnick
 900 attack shown in Fig. 6. For simplicity, we do not
 901 show these subscriptions.

902 In addition to the facts and beliefs received
 903 from other agents, an agent may have local
 904 variables that are used only by this agent. For
 905 example, variable *S* is a local variable of the DOS
 906 Attack Agent and *H* is a local variable of the IP
 907 Spoofing Agent. In addition, the DOS Attack Agent
 908 subscribes to the variables corresponding to the
 909 number of received SYNs (connection requests)
 910 and number of sent SYN-ACKs (connections han-
 911 dled) in a given time period. These values are
 912 obtained from system-monitoring agents. Local
 913 and input variables are used to calculate the
 914 probability distribution of a DOS attack.




Figure 6 Bayesian Network for Mitnick attack.

915 Based on the value of S , belief about a DOS
 916 attack is computed. We distinguish between three
 917 states: low, medium and high probability. If all
 918 connection requests are handled, S is equal to zero
 919 or has a very small value, thus the probability of
 920 the attack is either zero or low. When the system is
 921 under attack, with the result that many connec-
 922 tion requests cannot be handled, the probability of
 923 the attack is high. Actual values that differentiate
 924 between high and low states can be determined by
 925 applying data mining techniques on log data.

926 Belief calculation by system-monitoring 927 agents

928 System-monitoring agents perform simple process-
 929 ing and querying on log files and compute beliefs
 930 on variables they publish. These agents use the
 931 *method of counts* (Jensen, 2001; Cowell et al.,
 932 1999) to estimate the prior marginal probability of
 933 a variable being in a certain state by dividing the
 934 number of cases in which the variable is in that
 935 state by the total number of cases. The method of
 936 counts estimates the prior conditional probability
 937 of a variable being in a certain state given that its
 938 parents are in a certain configuration by dividing the
 939 number of cases in which the child variable is in
 940 that state and the parent variables are in that
 941 configuration by the number of cases in which the
 942 parent variables are in that configuration.

943 Modelling errors in measurement

944 The calculation of a belief depends on factors such
 945 as accuracy of measurement and conflicts among
 946 beliefs reported by various agents. In our model, if
 947 an agent is not able to accurately determine the
 948 state of a published variable, the agent publishes
 949 a probability distribution (belief) over the possible
 950 states of the variable. The publishing agent de-
 951 termines this distribution by incorporating mea-
 952 surement errors. Errors in the measurement of
 953 a variable state are modelled within an agent with
 954 help of the Bayesian Network shown in Fig. 7. This
 955 is achieved by representing the state of a variable
 956 with a belief or soft finding. The parent node S
 957 represents the actual value of interest. The prior
 958 distribution of the actual values is $P(S)$. The
 959 measured value is represented by variable S_{obs} .
 960 The measurement error is modelled by the condi-
 961 tional probability $P(S_{obs} | S)$. In the absence of
 962 error, this is a diagonal matrix. The magnitude of
 963 non-diagonal entries is directly proportional to the
 964 measurement errors. In the special case of a 2×2
 965 matrix, the two entries on the main diagonal




Figure 7 Incorporating error in measurement of variable.

quantify the specificity and sensitivity of the 966
 measurement, and the other entries quantify the 967
 false positive and false negative ratios (Vomlel, 968
 2004). When the actual value is propagated to 969
 parent node S , we get a probability distribution 970
 over different states of the variable. The agent 971
 can publish this distribution as its belief on the 972
 state of the measured variable. 973

Conflict resolution

974

Conflicts among beliefs on a state of variable, due 975
 to information provided by multiple agents on the 976
 same underlying quantity, can be resolved using 977
 soft-evidential update. For example, let A_1 and A_2 978
 be two agents that measure a variable v . The 979
 values measured by them are B_1 and B_2 , respec- 980
 tively. We design a Bayesian Network as shown in 981
 Fig. 8. The computed posterior probability of v 982
 effectively fuses the information provided by the 983
 two agents in the context specified by variable CR. 984

This approach requires estimating the prior 985
 probabilities of B and CR. In most practical uses 986




Figure 8 Conflict resolution.

of the Bayesian Network, the value of CR is known, so the assessment of the prior probability of CR does not need to be accurate. The prior probability of B needs to be more accurate than CR. It is normally possible to estimate B by using counts of the values of B in past cases. A similar technique (based on counts) can be used for the conditional probability tables $P(B_1 | v, CR)$ and $P(B_2 | v, CR)$. See Jensen (2001) and Cowell et al. (1999) for a discussion of the technique in general and Valdes and Skinner (2000) for an application of the technique in an intrusion detection scenario. For the situation involving complete cases, i.e., cases in which all variables are observed, the technique consists simply of replacing the (prior or conditional) probability of interest with the corresponding observed frequency (in the case of prior probabilities) or with a ratio of frequencies (in the case of conditional probabilities). In the more interesting and realistic case in which some variables are not observed, a similar approach, called *fractional updating* (or one of its improved variants) is used, see Jensen (2001) and Cowell et al. (1999) for details. To apply the technique in intrusion detection situation, we require that cases be labelled by attack type.

In special cases, B_1 and B_2 are statements that v is in a particular value. In general, they are probability distributions representing each agent's belief that the variable v has a particular value. The unique feature of the AEBN approach is to allow such general situations, whereas other approaches require the beliefs of the two agents to be hard findings. The process of updating v in the presence of the probability distributions on B_1 and B_2 is called soft-evidential update. In this work, we have used the Big Clique algorithm for soft-evidential update, implemented in the BC–Hugin system (Kim et al., 2004).

1026 Implementation

In this section, we describe our implementation of the proposed architecture from two different perspectives. First, we explain the developer's view of the system, and then we describe how the user (System Administrator) can interact with the IDS.

1032 Developer's perspective

The PAID system uses a behavior-based agent model. In this model, agents are characterized by certain behaviors. A behavior class describes the action that an agent will perform during its life

time. Domain specific behaviors are developed by extending class *Behavior* defined in JADE API. These behaviors may be either one shot behaviors or cyclic behaviors. Once a behavior completes its task, it may change its state to inactive by setting the instance variable *done* to true. The underlying agent management system in the agent platform (JADE is this case) invokes agents' active behaviors in each simulation cycle. We now describe the constituent modules of the PAID system:

- Main IDS agent (*IDS*): A singleton agent to supervise the working of the entire system and provide results. *IDS* agent provides the administrative interface. It also controls other tasks in the PAID system including creation and termination of system-monitoring and intrusion-monitoring agents. *IDS* agent exhibits *StartAgentsBehavior* and *StopAgentsBehavior*.
- System-Monitoring Agents (*SMAgent*): A class representing the system-monitoring agents in the *IDS*. This class is responsible for registering itself with the JADE DF, and for executing *PublishingBehavior* and a custom behavior to query log files or measure system performance. The name of custom behavior class is determined by the main *IDS* agent from the system-audit configuration file and is invoked during runtime with the help of Java Reflection Class API.
- Intrusion-Monitoring Agents (*IMAgent*): A class representing the intrusion-monitoring agents responsible for detecting intrusions. This class is responsible for registering itself with the JADE DF. A Bayesian Network model of intrusion to be monitored by this agent is provided as an argument to this agent on startup. From the input intrusion model, the agent determines required input beliefs and queries the directory facilitator to locate agents publishing those beliefs. This agent then subscribes to beliefs of other agents and updates its belief on intrusion periodically. In other words, intrusion-monitoring agents exhibit *SubscriptionBehavior*, *BeliefUpdateBehavior*, and *PublishingBehavior*.

Administrator's perspective

The administrative interface provided by our implementation is shown in Fig. 9. As described earlier in section 'PAID architecture', the PAID system contains several system-monitoring agents and intrusion-monitoring agents that utilize the directory facilitator (DF) provided by JADE as the

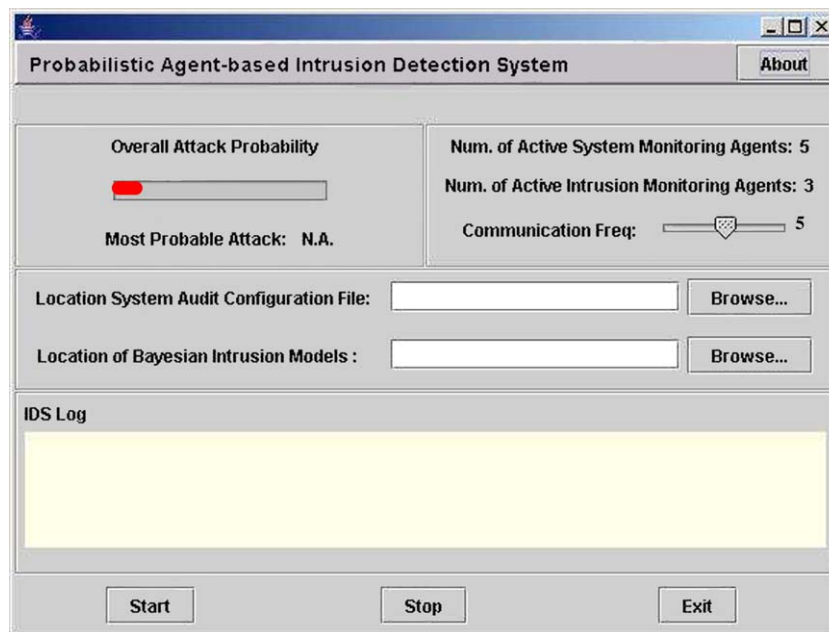


Figure 9 Administrative interface.

1089 registry agent. The interface allows the adminis-
 1090 trator to choose communication frequency among
 1091 these agents. To start the IDS, administrator must
 1092 specify a system-audit configuration file and di-
 1093 rectory where Bayesian network models for various
 1094 intrusions are stored. The system-audit configura-
 1095 tion file provides bootstrap information (name of
 1096 Behavior class and input log files) for system-
 1097 monitoring agents. Multiple intrusion-monitoring
 1098 agents are started, each with a different Bayesian
 1099 network model as input. The output of the PAID
 1100 system is the overall probability for the host
 1101 computer to be under attack. This value is graph-
 1102 ically shown in the user interface. When IDS
 1103 identifies a probable attack, it brings up detailed
 1104 probability information of that attack. For exam-
 1105 ple, Fig. 10 shows detailed information for Mitnick
 1106 attack.

1107 The PAID system goes through the following four
 1108 phases:

(1) *Initialization phase*: All the agents register themselves with the agent registry on boot-up. JADE provides APIs for enabling the agents to register themselves with the AMS and DF agents for the system. This provides every agent with a globally unique identifier, the Agent-ID (AID), through which the other agents can interact by taking advantage of the white page services provided by the JADE AMS. In addition, each agent has to provide its service description during registration, which the JADE DF uses to provide yellow page services to other agents.

(2) *Analysis phase*: After initialization, agents enter analysis phase. In this phase agents execute *SubscriptionBehavior*, *BeliefUpdateBehavior*, and *PublishingBehavior*. These behaviors are cyclic, i.e. they are repeated indefinitely after every few seconds determined by the communication frequency set for the session.

(3) *Resetting phase*: The Administrative interface allows the user to reset all the agents by stopping all agents with the *Stop* button and starting the system again. When the system is stopped, all agents are deregistered and terminated. The administrator may then start all the agents again by pressing the *Start* button, or exit the system. This feature can be useful if some agents terminate abnormally during execution and need to be restarted.

(4) *Termination phase*: When the administrator *Exits* the system, all the agents are deregistered and the IDS shuts down.

Conclusions

In this paper, we demonstrated the feasibility of probabilistic intrusion detection technique using soft-evidential updates. We developed and implemented an intrusion detection architecture called Probabilistic Agent-Based Intrusion Detection (PAID). The advantages of our framework over existing models follow.

1142

1143
 1144
 1145
 1146
 1147
 1148
 1149

The screenshot shows a window titled "Node Values" with several sections for configuring attack probabilities:

- DataVsAckPacket:** Data pkt: 0.98975, Ack pkt: 0.01025
- DecreaseInDataOut:** yes: 0.98404, no: 0.01596
- DOSAttack:** Radio buttons for No Evidence, good: 0.23738, bad: 0.76262
- varInTTL:** yes: 0.92645, no: 0.07355
- OneWayCommunication:** Yes: 0.99307, No: 0.00693
- IPspoofing:** Radio buttons for No Evidence, yes: 0.03801, no: 0.96199
- LocalSrcOnExtIntf:** Radio buttons for No Evidence (selected), yes: 0.13041, no: 0.86959
- MitnickAttack:** Radio buttons for No Evidence, yes: 0.77986, no: 0.22014

At the bottom of the window are two buttons: "Initialize" and "Propagate".

Figure 10 Calculation of attack probability using Big Clique algorithm.

1150 PAID requires low volume of data sharing over
 1151 network in contrast to centralized data analysis.
 1152 Although communication overhead is higher than
 1153 in IDS that allow only binary decision sharing, the
 1154 improved processing power makes PAID more suit-
 1155 able for sophisticated intrusion detection. PAID
 1156 also provides a continuous scale to represent event
 1157 probabilities. This feature allows easy exploration
 1158 of the trade-off between sensitivity and selectivity
 1159 that affects the rate of false positive and false
 1160 negative decisions.

1161 The current version of PAID was illustrated in
 1162 misuse detection mode, but the same principles
 1163 can be applied for anomaly-based intrusion de-
 1164 tection. Distributed intrusion detection is achieved
 1165 by allowing each agent to cooperate with others
 1166 and to build full or partial, global intrusion graphs.
 1167 Distributed processing not only increases efficiency
 1168 but also eliminates single point of failure.

1169 A proof-of-concept prototype of our model has
 1170 been developed using agents developed with Java
 1171 and C alone. At present we are migrating the
 1172 complete agent model to JADE framework. We are
 1173 planning to improve and fine-tune our current
 1174 model to address agent trust management and
 1175 dynamic agent-activation protocols.

Acknowledgments

1176

This material is based upon work supported by
 National Science Foundation under Grant No. IIS-
 0237782. This work was supported in part by the
 Advanced Research and Development Agency
 (ARDA), an entity of the U.S. Government. Any
 opinions, findings, conclusions, or recommenda-
 tions expressed in this material are those of the
 authors and do not necessarily reflect the views of
 the U.S. Government.

1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185

References

1186

- Asaka M, Taguchi A, Goto S. The implementation of IDA: an intrusion detection agent system. In: Proceedings of 11th annual FIRST conference on computer security incident handling and response. Brisbane, Australia; 1999.
- Axelsson S. Intrusion detection systems: a taxonomy and survey. Tech. Rep. 99-15. Göteborg, Sweden: Department of Computer Engineering, Chalmers University of Technology; 2000.
- Balasubramanian J, Garcia-Fernandez J, Isacoff D, Spafford E, Zam-Boni D. An architecture for intrusion detection using autonomous agents. Coast 98-05. West Lafayette, IN: Department of Computer Science, Purdue University; 1998.

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200

- 1201 Barbara D, Wu N, Jajodia S. Detecting novel network intrusions
1202 using bayes estimator. In: Proceedings of first SIAM confer-
1203 ence on data mining; 2001.
- 1204 Bellifemine F, Poggi A, Rimassa G. JADE – a FIPA compliant
1205 agent framework. In: Proceedings of the fourth international
1206 conference and exhibition on the practical application of
1207 intelligent agents and multi-agents. London; 1999.
- 1208 Bloemeke M, Valtorta M. The rumor problem in multiagent
1209 systems. Tech. Rep. 2002–006. Columbia, SC: Department
1210 of Computer Science and Engineering, University of South
1211 Carolina; 2002.
- 1212 Bray T, Paoli J, Sperberg-McQueen C, Maler E. Extensible
1213 Markup Language (XML) 1.0 specification. W3C Recommendation,
1214 Retrieved October 16, 2002 from, <[http://www.
1215 w3.org/TR/2000/REC-xml-20001006](http://www.w3.org/TR/2000/REC-xml-20001006)>; 2001.
- 1216 Carver C, Hill J, Surdu J, Pooch U. A methodology for using
1217 intelligent agents to provide automated intrusion response.
1218 In: Proceedings of the IEEE systems, man, and cybernetics
1219 information assurance and security workshop. WestPoint,
1220 NY; 2000.
- 1221 Cho S, Cha S. SAD: web session anomaly detection based on
1222 parameter estimation. Computers and Security, in press.
- 1223 Cooper G. The computational complexity of probabilistic
1224 inference using Bayesian networks. Artificial Intelligence
1225 1990;42(2–3):393–405.
- 1226 Cooper GF, Herskovits E. A Bayesian method for the induction of
1227 probabilistic networks from data. Machine Learning 1992;
1228 9(4):309–47.
- 1229 Cowell RG, Lauritzen SL, David AP, Spiegelhalter DJ, Nair V,
1230 Lawless J, et al. Probabilistic networks and expert systems.
1231 Springer-Verlag New York, Inc.; 1999.
- 1232 Dacier M. Design of an intrusion-tolerant intrusion detection
1233 system. Tech. Rep. D10. IBM Zurich Research Laboratory;
1234 2002.
- 1235 DuMouchel W. Computer intrusion detection based on Bayes
1236 factors for comparing command transition probabilities.
1237 Tech. Rep. 91. National Institute of Statistical Sciences;
1238 1999.
- 1239 FIPA. FIPA ACL Message structure specifications. Retrieved
1240 October 16, 2002 from, <[http://www.fipa.org/specs/
1241 fipa00061](http://www.fipa.org/specs/fipa00061)>; 2002a.
- 1242 FIPA. FIPA-OS Developers guide. Retrieved October 16, 2002
1243 from, <[http://fipa-os.sourceforge.net/docs/Developers-
1244 Guide.pdf](http://fipa-os.sourceforge.net/docs/Developers-Guide.pdf)>; 2002b.
- 1245 Friedman N, Singer Y. Efficient Bayesian parameter estimation
1246 in large discrete domains. In: Proceedings of neural in-
1247 formation processing systems (NIPS 98). MIT Press; 1998.
- 1248 Helmer G, Wong JSK, Honavar V, Miller L, Wang Y. Lightweight
1249 agents for intrusion detection. Journal of Systems and
1250 Software 2003;67(2):109–22.
- 1251 Jansen W, Mell P, Karygiannis T, Marks D. Applying mobile
1252 agents to intrusion detection and response. Tech. Rep. 6416.
1253 Gaithers-burg, MD: Computer Security Response Center,
1254 National Institute of Standards and Technology; 1999.
- 1255 Jensen F. Bayesian networks and decision graphs. Springer
1256 Verlag; 2001.
- 1257 Kim Y, Valtorta M, Vomlel J. A prototypical system for soft
1258 evidential update. Applied Intelligence 2004;21(1):81–97.
- 1259 Laskey KB, Mahoney SM, Wright E. Hypothesis management in
1260 situation-specific network construction. In: Proceedings of
1261 the 17th annual conference on uncertainty in artificial
1262 intelligence (UAI-01). Seattle, WAAugust 2001. p. 301–9.
- 1263 Lauritzen S, Spiegelhalter DJ. Local computations with proba-
1264 bilities on graphical structures and their application to
1265 expert systems. Journal of the Royal Statistical Society,
1266 Series B (Statistical Methodology) 1988;50(2):157–224.
- 1267 Moy J. OSPF Version 2. Internet draft, RFC-2178; 1997.
- Neapolitan RE. Probabilistic reasoning in expert systems: theory 1268
and algorithms. New York, NY: John Wiley and Sons; 1990. 1269
- Neapolitan RE. Learning Bayesian networks. Upper Saddle River, 1270
NJ: Pearson Prentice Hall; 2004. 1271
- Neumann P, Porras P. Experiences with EMERALD to date. In: 1272
Proceedings of the first USENIX workshop on intrusion
1273 detection and network monitoring. Santa Clara, CA; 1999. 1274
- Pearl J. Probabilistic reasoning in intelligent systems: networks 1275
of plausible inference. San Mateo, CA: Morgan-Kaufmann
1276 Publishers; 1988. 1277
- Perlman R. Interconnections: bridges and routers. Reading, MA: 1278
Addison-Wesley Professional; 1992. 1279
- Sebyala AA, Olukemi T, Sacks L. Active platform security 1280
through intrusion detection using nave Bayesian network
1281 for anomaly detection. In: Proceedings of London commu-
1282 nications symposium; 2002. 1283
- Singh M, Valtorta M. A new algorithm for the construction of 1284
Bayesian network structures from data. In: Proceedings of
1285 the ninth annual conference on uncertainty in artificial
1286 intelligence (UAI-93). Washington, DC; July 1993. p. 259–64. 1287
- Singh M, Valtorta M. Construction of Bayesian belief networks 1288
from data: a brief survey and an efficient algorithm.
1289 International Journal of Approximate Reasoning February
1290 1995;12(2):111–31. 1291
- Snapp S, Brentano J. DIDS (Distributed Intrusion Detection 1292
System) – motivation, architecture, and an early prototype.
1293 In: Proceedings of the 1991 national computer security
1294 conference; 1991. 1295
- Spafford EH, Zamboni D. Intrusion detection using autonomous 1296
agents. Computer Networks 2000;34(4):547–70. 1297
- Spiegelhalter D, Dawid A, Lauritzen SL, Cowell R. Bayesian 1298
analysis in expert systems. Statistical Science 1993;8(3):
1299 219–83. 1300
- Utete SW. Local information processing for decision making 1301
in decentralized sensing networks. In: Proceedings of
1302 the 11th international conference on industrial and
1303 engineering applications of artificial intelligence and
1304 expert systems (IEA/AIE-98). Benicassim, Castellon, Spain;
1305 1998. p. 667–76. 1306
- Valdes A, Skinner K. Adaptive, model-based monitoring for 1307
cyber attack detection. In: Proceedings of recent advance in
1308 intrusion detection. LNCS, No. 1907; 2000. p. 80–92. 1309
- Valtorta M, Kim Y-G, Vomlel J. Soft evidential update for 1310
probabilistic multiagent systems. International Journal of
1311 Approximate Reasoning January 2002;29(1):71–106. 1312
- Vomlel J. Probabilistic reasoning with uncertain evidence. 1313
Neural Network World. International Journal on Neural and
1314 Mass-Parallel Computing and Information Systems 2004;
1315 14(5):453–6. 1316
- Xiang Y. Probabilistic reasoning in multiagent systems: a graph- 1317
ical models approach. Cambridge: Cambridge University
1318 Press; 2002. 1319
- Vaibhav Gowadia** is a doctoral student at the Department of 1318
Computer Science and Engineering and Research Assistant in 1319
Information Security Laboratory at University of South Carolina, 1320
Columbia, SC. His current research interests include information 1321
systems security, intrusion detection, security protocols, 1322
threshold cryptography, and Semantic Web security. He 1323
obtained his Master of Science degree in Computer Engineering 1324
at the University of South Carolina in 2003, and Bachelor of 1325
Technology degree in Instrumentation and Control Engineering 1326
at the National Institute of Technology, Jalandhar, India in 2000. 1327
- Csilla Farkas** is an Assistant Professor at the Department of 1326
Computer Science and Engineering, and Director of the In- 1327
formation Security Laboratory at the University of South 1328

Carolina. She received a B.S. degree in Geological Sciences from the Eotvos Lorand University (1985), Hungary, a B.S. in Computer Science from SZAMALK (1989), Hungary, and a B.S. in Computer Science (1993) and Ph.D. in Information Technology (2000) from George Mason University, VA. Dr. Farkas' research interests include information security, data inference problem, economic and legal analysis of cyber crime, and security and privacy on the Semantic Web. She is a recipient of the National Science Foundation Career award. The topic of her award is "Semantic Web: Interoperation vs. Security – A New Paradigm of Confidentiality Threats." Dr. Farkas actively participates in international scientific communities as program committee member and reviewer.

Marco Valtorta is an Associate Professor at the Department of Computer Science and Engineering at the University of South Carolina. He obtained a Laurea in Electrical Engineering from the Politecnico di Milano in 1980, an M.A. in Computer Science from Duke University in 1984, and a Ph.D. in Computer

Science from Duke University in 1987. Between 1985 and 1988, he was a project officer for ESPRIT at the Commission of the European Communities in Brussels, where he supervised projects in the Advanced Information Processing area. As a faculty member at the University of South Carolina since 1988, he has conducted research funded by ARDA, SPAWAR, DARPA, the Office of Naval Research (ONR), the U.S. Department of Agriculture (DOA), CISE (an Italian laboratory controlled by ENEL, the state electricity company), and the South Carolina Law Enforcement Division. He is the author of over 30 refereed publications, an associate editor of the International Journal of Approximate Reasoning, and a member of the editorial boards of Applied Intelligence and of the International Journal of Applied Management and Technology. His research interests are in the areas of normative reasoning under uncertainty (especially Bayesian Networks, influence diagrams, and their use in stand-alone and multiagent systems), heuristics for problem solving, and computational complexity in artificial intelligence.

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

UNCORRECTED PROOF