

Mobility in Collaborative Alert Systems: Building Trust through Reputation

Manuel Gil Pérez^{1,*}, Félix Gómez Mármol²,
Gregorio Martínez Pérez¹, and Antonio F. Gómez Skarmeta¹

¹ Departamento de Ingeniería de la Información y las Comunicaciones,
University of Murcia, 30071 Murcia, Spain
Tel.: +34 868 887645, Fax: +34 868 884151
{[mgilperez](mailto:mgilperez@um.es),[gregorio](mailto:gregorio@um.es),[skarmeta](mailto:skarmeta@um.es)}@um.es

² NEC Europe Ltd., Kurfürsten-Anlage 36, 69115 Heidelberg, Germany
felix.gomez-marmol@neclab.eu

Abstract. Collaborative Intrusion Detection Networks (CIDN) are usually composed by a set of nodes working together to detect distributed intrusions that cannot be easily recognized with traditional intrusion detection architectures. In this approach every node could potentially collaborate to provide its vision of the system and report the alarms being detected at the network, service and/or application levels. This approach includes considering mobile nodes that will be entering and leaving the network in an ad hoc manner. However, for this alert information to be useful in the context of CIDN networks, certain trust and reputation mechanisms determining the credibility of a particular mobile node, and the alerts it provides, are needed. This is the main objective of this paper, where an inter-domain trust and reputation model, together with an architecture for inter-domain collaboration, are presented with the main aim of improving the detection accuracy in CIDN systems while users move from one security domain to another.

Keywords: Alert Systems; Mobility; Collaboration; Trust; Reputation.

1 Introduction

The open nature of mobile networks has promoted in recent years a high mobility of users among heterogeneous wireless networks, in which they can join to gain access. For example, students, teachers and research staff of a university can benefit of this mobility service in their academic campus without interceding in their wireless connectivity. They can in turn join other networks of the system later, thereby creating a distributed chain of interactions where they can enter and participate in each of those security domains that form the network.

Mobility of users represents a more global and new way of operation, but also introduces some challenges have to be tackled to control and, as a main

* Corresponding author.

goal, avoid their possible disruptive behaviors. Current and existing works measure and quantify the trustworthiness of users based on their reputation [1, 2]. That is, they estimate how good users are according to previous experiences or interactions they have had with the system in the past.

We focus mainly in this paper on controlling what mobile users do when they travel across heterogeneous networks and, specifically, in roaming users who can join and cooperate in the context of collaborative alert systems. These systems are built up from the inspection of a great amount of alerts produced in an individual fashion by each of the Intrusion Detection Systems (IDSs) [3]. Collaborative alert systems provide the basis for building a global knowledge of alerts, based on the cooperation of all members of the system, to increase the accuracy in detecting distributed threats from a more global point of view.

In this paper we expose a *Collaborative Intrusion Detection Network* (CIDN) designed to improve the accuracy when detecting distributed intrusions. To this end, a global and common alert system is necessary to achieve a high level overview of the entire system. This provides a way to know and anticipate the diverse actions an attacker can execute to compromise the system. This cooperative knowledge is then built upon a set of alerts exchanged among all the components of the network.

The building of this knowledge implies two kinds of communications with the aim of spreading out and sharing these alerts: an intra-domain exchange of alerts among all the detection units, either Host-based IDSs (HIDSs) or Network-based IDSs (NIDSs), into the same security domain; and an inter-domain communication among varying security domains that comprise the alert system to build the desired high level overview.

In this paper we focus on the process of building a common and cooperative knowledge of alerts among security domains, i.e., inter-domain communications, in the context of highly distributed environments. As detection units, especially *mobile HIDSs*, which can move from one security domain to others, we also propose a reputation mechanism to compute the trustworthiness each domain has on these detection units, regardless of the domain to whom they belong.

Reputation values will give security domains a way of assessing whether any of their detection units exhibits a correct or malicious behavior. Alerts provided from those malicious detection units, probably sent to diminish the performance of the system, or even due to a malfunctioning IDS, can be dropped to avoid confusion to others. This fact will improve the detection accuracy by rejecting false alarms received from malicious entities. Thus, the detection of certain malicious behaviors from roaming users is an essential requirement for the success in detecting distributed attacks.

The remainder of this paper is structured as follows. In Section 2, we motivate the problem under consideration. Then, we formulate in Section 3 the reputation mechanism designed to assess the mobile users' trustworthiness when they move across security domains. Section 4 describes the main components that comprise our system architecture. Next, Section 5 discusses the main related work and,

finally, Section 6 remarks the main conclusions and highlights the lines of current and future work.

2 Design of an Inter-domain Collaborative Alert System

This section describes the problem under consideration, which will be used throughout this paper to introduce and describe later the different components that compose the proposed system architecture.

How these components are distributed, and how they communicate each other, will be the basis for the definition of our inter-domain reputation mechanism. This mechanism aim to improve the detection accuracy in collaborative alert systems while roaming users move from one security domain to another.

As a sample scenario, let us suppose the system architecture depicted in Fig. 1. This example is composed by two different administrative domains, *A* and *B*. An administrative domain is managed by a single organizational authority, which consists of a certain number of hosts, e.g., workstations, servers or network devices.

Each administrative domain can in turn be split into more than one security domain. For example, the *Administrative Domain A* in Fig. 1 is split into two interconnected domains, named *Domain A₁* and *Domain A₂*. Each of them is managed internally by the qualified staff in the security field who defines the security policies, if any, under which the domain is governed.

In the scenario presented in Fig. 1, it is worth mentioning that there is an overlapping between two security domains. This overlapping may however not exist, so both domains would act as two networks totally independent of each other but belonging to the same administrate domain.

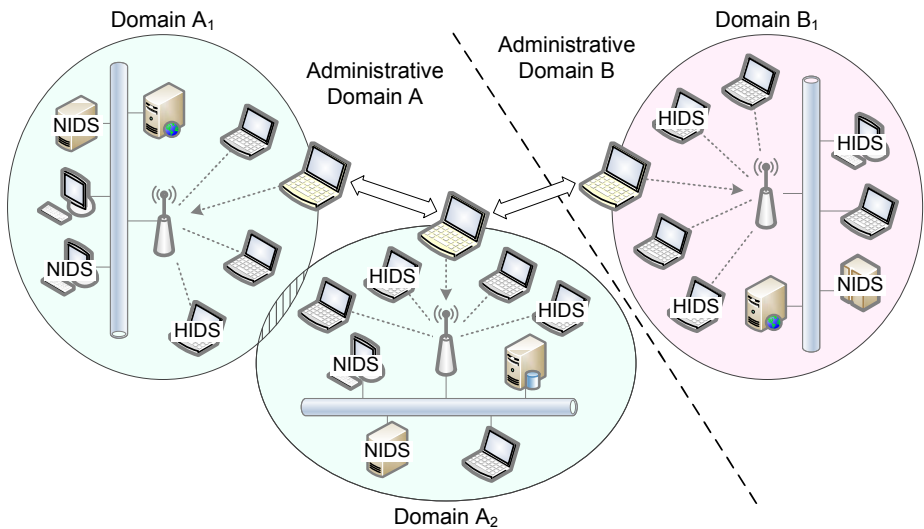


Fig. 1. Mobility of a user among security domains

Hosts connected to a given security domain can make use of it in several ways. They can vary from the ones installed by the domain's administrators to offer a specific service, e.g., Web servers or printing services, and even specific services to monitor and assess the fulfillment of both internal policies and users' behavior. With regard to the latter, we focus on the use of Intrusion Detection Systems, which enable hosts with detection capabilities to know whether an attacker is exploiting some vulnerability in other system components.

As shown in Fig. 1, some hosts have been labeled according to their skills to detect suspect or malicious activities when taking two different kinds of sources of information: at local host to monitor what is happening inside of it, by making use of a HIDS like the OSSEC tool [4]; and globally to inspect every single packet that goes through the network, by using a well-known NIDS such as Snort [5].

The distribution of these IDSs is as follows. We assume every NIDS is managed by a domain administrator who will have necessary skills to set it up. As NIDSs are detection units totally dependent of their security domain, they are configured according to the internal network distribution and the services their domain offers. These NIDSs have to be installed and deployed as static elements. That is, NIDSs cannot operate in mobile devices to detect anomalies in more than one security domain. Instead, HIDSs do not undergo this kind of restrictions since they only report malicious activities conducted internally in their device. As a consequence, HIDSs can seamlessly operate in mobile scenarios.

In Fig. 1, it can be seen a *mobile HIDS* that joins different wireless networks while it moves across security domains. In each of them, this HIDS can voluntarily join and collaborate (or disrupt) with such a domain by sending out alerts generated by it. This *mobile HIDS*, when joining to a particular domain, can come from another domain and it is not really a newcomer in the system.

Thus, the new security domain has to compute what its reputation value is, without forgetting what other domains *think* about such a detection unit. The calculation of reputation on *mobile HIDSs* takes place in a special unit for the domain, named *administration unit*, possibly the most trustworthy unit thereof.

Next sections detail how reputation values are computed, depending on whether *mobile HIDSs* are moving from one security domain to another, both belonging to the same administrative domain, or whether these movements imply more than one administrative domain.

3 Inter-domain Reputation System

As stated before, the problem we need to solve here is to assess the trustworthiness of the alerts generated by a *mobile HIDS* that is traveling across different security domains.

So let $HIDS_{\Omega_{j_1}}^i$ denote the i -th HIDS, $HIDS^i$, being currently at domain Ω_{j_1} , regardless the domain it actually belongs to.

When $HIDS_{\Omega_{j_1}}^i$ moves from domain Ω_{j_1} to Ω_{j_2} , the latter has to compute the reputation for such a *mobile HIDS* based upon two main sources of information; namely:

- the previous behavior records of $HIDS^i$ within domain Ω_{j_2} , if any; and
- the recommendations provided by other security domains where such $HIDS^i$ remained in the past, weighted by the reliability of each of those domains from the perspective of domain Ω_{j_2} .

More formally, the reputation of $HIDS^i_{\Omega_{j_1}}$ within the domain Ω_{j_2} , which is represented as $Rep_{\Omega_{j_2}}(HIDS^i_{\Omega_{j_1}}) \in [0, 1]$, would be computed as follows:

$$Rep_{\Omega_{j_2}}(HIDS^i_{\Omega_{j_1}}) = \alpha \cdot Rep_{\Omega_{j_2}}(HIDS^i_{\Omega_{j_1}}) + \beta \cdot \frac{\sum_{k=1}^n Rep_{\Omega_{j_k}}(HIDS^i_{\Omega_{j_1}}) \times T_{\Omega_{j_2}, \Omega_{j_k}}}{\sum_{k=1}^n T_{\Omega_{j_2}, \Omega_{j_k}}} \tag{1}$$

where $\alpha, \beta \in [0, 1]$ (fulfilling that $\alpha + \beta = 1$) represent the weights given to previous experiences of domain Ω_{j_2} with $HIDS^i$ (direct experiences), and the recommendations provided by other domains (indirect experiences), respectively. In turn, $T_{\Omega_{j_2}, \Omega_{j_k}}$ represents the trustworthiness or reliability on domain Ω_{j_k} given by domain Ω_{j_2} .

For the selection of the most appropriate value for the weights α and β , we propose two alternatives: a light but less accurate one; and a more accurate approach though requiring some more computation capabilities.

On one hand, α and β values can be predefined as static ones. That is, when bootstrapping the system, a value like $\alpha = 0.7$ and $\beta = 0.3$, for instance, can be set and never changed along the time. On the other hand, those values could be computed in a dynamic fashion as follows:

$$\begin{aligned} \alpha &\rightarrow \alpha_{\Omega_{j_k}}(HIDS^i) \\ \beta &\rightarrow \beta_{\Omega_{j_k}}(HIDS^i) = 1 - \alpha_{\Omega_{j_k}}(HIDS^i) \end{aligned} \tag{2}$$

Then, $\alpha_{\Omega_{j_k}}(HIDS^i)$ would be based on the number of alerts generated by the mobile $HIDS^i$ within domain Ω_{j_k} , so the more alerts it has generated in the past, the higher the weight domain Ω_{j_k} will give to its own experience with regards to $HIDS^i$, in contrast to the weight given to other domains' suggestions. In this sense, the direct experiences of $HIDS^i$ within domain Ω_{j_k} will have a higher importance as this detection unit collaborates more with this domain rather than the rest. On the contrary, $\beta_{\Omega_{j_k}}(HIDS^i)$ will have a higher weight when $HIDS^i$ provides more alerts to other domains than this one (indirect experiences).

Additionally, the temporal distribution of the alerts generated by $HIDS^i$ could be also taken into consideration in order to compute $\alpha_{\Omega_{j_k}}(HIDS^i)$. In this way, the most recent alerts will be more significant to represent the current behavior of a mobile $HIDS^i$ than the older ones. To this end, domain Ω_{j_k} has to assign a higher weight to the former ones as opposed to the weight given to the latter ones.

Thus, the calculation of $Rep_{\Omega_{j_k}}(HIDS^i)$ in a dynamic fashion provides a new way of tuning the reputation of $HIDS^i$ when it joins with domain Ω_{j_k} , instead

of assigning it a fixed reputation value as if this detection unit was a newcomer in the system.

Furthermore, it is worth mentioning that the mechanism for computing the amount of trust deposited by one domain in the recommendations provided by other domain, i.e., $T_{\Omega_{j_1}, \Omega_{j_2}}$, will be different depending on whether those two domains Ω_{j_1} and Ω_{j_2} are actually subdomains belonging to the same administrative domain, or they indeed represent different administrative domains.

Thus, if Ω_{j_1} and Ω_{j_2} are subdomains belonging to the same administrative domain, then $T_{\Omega_{j_1}, \Omega_{j_2}}$ will be, in some certain situations, directly equal to 1. In any case, when computing $T_{\Omega_{j_1}, \Omega_{j_2}}$, it will most probably have a higher value if Ω_{j_1} and Ω_{j_2} belong to the same administrative domain than if they do not.

Finally, the reputation of one HIDS, either mobile or not, within a domain is built up based on its behavior when generating alerts. Thus, a HIDS spreading out false or even malicious alerts, will end up having a low reputation and hence its alerts will most probably not be taken into consideration anymore. On the contrary, if the HIDS behaves properly and it is able (or collaborates) to detect critical threats, it will end up being considered as highly trustworthy and its generated alerts will be treated as reliable ones.

4 Architecture for Inter-domain Collaboration

This section outlines the proposed system architecture by means of the definition of the main functional blocks each administration unit should support.

This description will provide a clear vision and understanding on how the reputation mechanism presented above is integrated in those administration units to manage the reputation values deposited on the *mobile HIDSs*. Therefore, we also focus in this section on the trust and reputation management to strengthen administration units with this new sort of knowledge.

4.1 Main Functional Blocks

The four functional blocks that constitute a particular administration unit are depicted in Fig. 2. At a glance, we can identify each administration unit providing four main functionalities.

The first three ones (communication interface, intrusion detection engine and reaction system) are explained next, while the module to manage the trust and reputation values of *mobile HIDSs* is detailed later.

Communication module

The communication interface allows administration units to exchange data in two senses. First, they are capable of requesting reputation data to other security domains about a particular *mobile HIDS*, when it moves from one of those domains to the current one. Then, the current domain will be able to compute the reputation value of the *mobile HIDS* according to Equation (1).

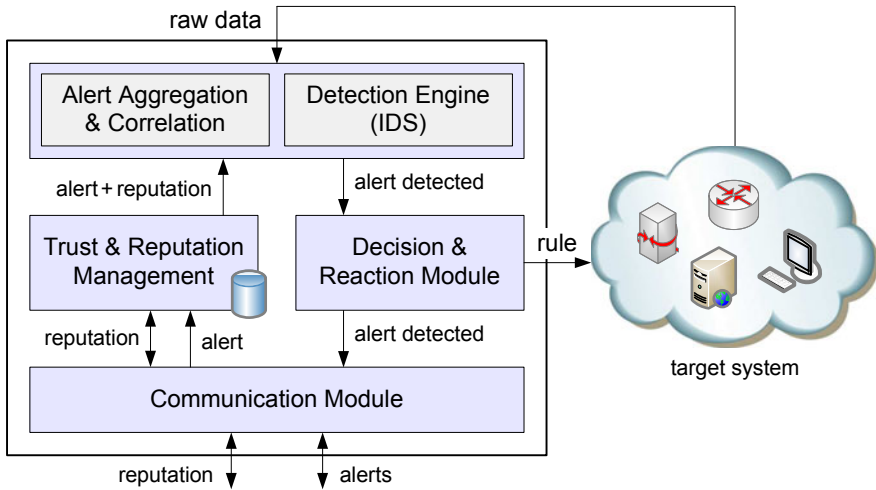


Fig. 2. Internal components deployed by an administration unit

Second, administration units will be also able to share the alerts detected by them with the rest of members of their domain. It aims at building a cooperative intra-domain knowledge of alerts, as well as sharing them with other security domains with which they have a close collaboration relationship.

Intrusion detection capabilities

As an administration unit can also offer detection capabilities, like any other detection unit of its domain, it is equipped with an intrusion detection module (first block at the top of Fig. 2) based on two different techniques: a *Detection Engine* to analyze and trigger alerts generated internally in the detection unit, by means of a well-known IDS solution such as Snort or OSSEC (depending on whether it is a NIDS or a HIDS, respectively); and an *Alert Aggregation & Correlation* submodule to generate higher level alerts that synthesize more complex intrusions by clustering isolated alerts in meaningful groups.

Both approaches are fed with raw data collected from the sources of information the detection unit handles, and with alerts generated by other IDSs of its domain. In the latter case, this module also receives the reputation value this unit has deposited on the one that generated the alert. With this information, this module can decide of using it or not depending on such a reputation value.

Decision and reaction system

This module enables the detection unit with reaction capabilities. It will enforce some security rules in the target system for remediating the damages caused by the unauthorized intrusion, provided this unit has the required credentials to do it. This module will also forward the alert detected by the intrusion

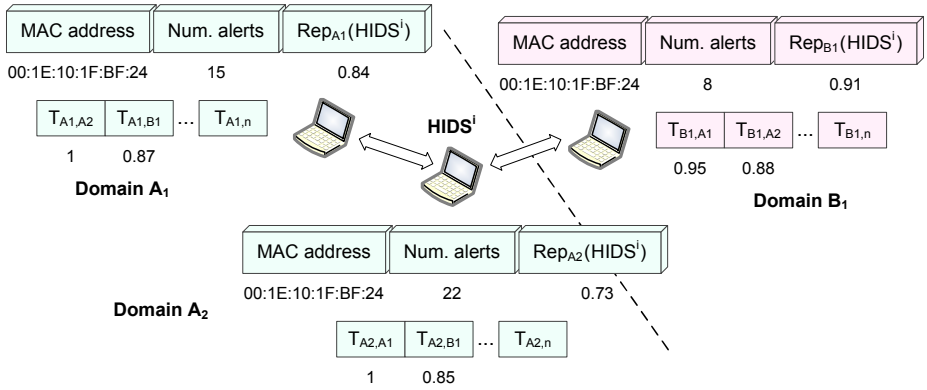


Fig. 3. Inter-domain reputation management for mobile HIDSs

detection module, either by the *Detection Engine* or by the *Alert Aggregation & Correlation* submodule, with the aim of sharing this knowledge with the rest of members of its community.

4.2 Trust and Reputation Management

An administration unit has to manage the reputation values of all IDSs of its domain, by taking their experiences with the domain in the past, especially those that can move from one security domain to another, i.e., the *mobile HIDSs*.

Let us suppose the sample scenario presented in Section 2. In such an example, a *mobile HIDS* joins several security domains and collaborates with them by means of sharing the alerts detected by it. Thus, this *mobile HIDS* will have a different reputation in each of them depending on its previous interactions with the domain (direct experiences) and past interactions conducted in other domains (indirect experiences), as explained in Section 3. Hence, it is necessary to maintain some historic data in each domain to compute these reputation values.

To that end, the *Trust & Reputation Management* module of Fig. 2 is required to deploy and maintain a repository for registering all detection units that have collaborated with the current domain. Fig. 3 shows the same scenario presented in Section 2, but extended with the information each administration unit has to maintain internally to compute reputation values.

The administration unit of each security domain will then store a list of IDSs, which have had any interaction with the domain before, with the following data for each detection unit:

- *MAC address*: unique identifier to distinguish this detection unit from the rest. Thus, this unit can be identified through the same MAC address in all the involved domains.
- *Number of generated alerts*: the amount of alerts is necessary to compute $\alpha_{\Omega_{j_k}}(HIDS^i)$, as explained in Equation (2).

- *Reputation value*: this attribute maintains updated the reputation value the domain has deposited on the current detection unit, which is computed by using Equation (1).

In addition to this information, each administration unit also has to store the trust its domain has on the rest of domains, i.e., $T_{\Omega_{j_1}, \Omega_{j_k}}$ from the perspective of domain Ω_{j_1} .

Bearing in mind the sample data shown in Fig. 3, let us suppose the mobile $HIDS^i$ is currently operating in domain A_1 , where its reputation is 0.84, and it then moves to the domain A_2 . In this new domain, where the mobile $HIDS^i$ had already participated in the past with 22 alerts, it has to compute the new reputation of mobile $HIDS^i$.

By applying Equation (1), this reputation value would be as follows:

$$Rep_{\Omega_{A_2}}(HIDS^i_{\Omega_{A_1}}) = 0,7 \cdot 0,73 + 0,3 \cdot \frac{0,84 \cdot 1 + 0,91 \cdot 0,85}{1 + 0,85} = 0,77$$

For this example, we have set static values for the weights on direct and indirect experiences; $\alpha = 0.7$ and $\beta = 0.3$, respectively.

As seen, the new reputation value is slightly higher in domain A_2 than before, from 0.73 to 0.77, since the behavior of such a detection unit in other domains was better than in domain A_2 . Then, the domain A_2 will update the reputation for the mobile $HIDS^i$ to this new value.

Following with this example, let us suppose the mobile $HIDS^i$ moves now from domain A_2 (where it has a new reputation value of 0.77) to domain B_1 . In this case, the mobile $HIDS^i$ travels from one administrative domain to another. The new reputation value of the mobile $HIDS^i$ in domain B_1 would be as follows:

$$Rep_{\Omega_{B_1}}(HIDS^i_{\Omega_{A_2}}) = 0,7 \cdot 0,91 + 0,3 \cdot \frac{0,84 \cdot 0,95 + 0,77 \cdot 0,88}{0,95 + 0,88} = 0,88$$

Now, the new reputation value is slightly lower (from 0.91 to 0.88) since this detection unit had a worse performance in other different domains.

5 Related Work

Enhancing security in highly distributed environments like mobile ad hoc networks (MANETs) or wireless sensor networks (WSNs) has been a subject of research for a number of years now. Recently, efficient and accurate trust and reputation management [2, 6, 7] has arisen as a novel and effective solution for certain security lacking environments.

Several works have been done so far in this field, dealing with MANET networks [1, 8, 9] and WSNs [10–12], among many other systems. In this way, for example, [1] presents a collaborative reputation model aimed to work in

MANETs in order to prevent selfish behavior of its users. The authors claim that all members have to contribute to the community life in order to be entitled to use its resources.

On the other hand, intrusion detection systems and, more specifically, collaborative intrusion detection networks have drawn as well the attention of a number of researchers and research groups from both academia and industry all around the world. These systems improve the efficiency of intrusion detection by exchanging low-level intrusion alerts between IDSs in order to produce a high level overview of the whole system [3].

Besides, the combination of trust and reputation management and CIDNs, or rather, the application of the former to the latter, is still at a preliminary stage. Only a few works like [13–15] address this issue. The first of them [13] proposes a trust-based framework for secure collaboration within a CIDN. In particular, each HIDS assesses how trustworthy other HIDSs are, based on its own experience with them. RADAR, which is presented in [14], consists of an anomaly detection system aimed to identify abnormal mesh nodes in wireless mesh networks through a reputation measurement characterizing and quantifying a node's behavior in terms of fine-grained performance metrics of interest. Finally, the authors of [15] show a trust-based selection mechanism for choosing the optimal information provider in an intrusion/fraud detection system, which works by analyzing the response of nodes (agents) to a set of prepared challenges inserted into the system.

Nevertheless, to the best of our knowledge, the work hereby presented is one of the first ones in the literature applying a reputation mechanism to assess the trustworthiness of the alerts generated by *mobile HIDSs* traveling across different domains in the context of collaborative intrusion detection networks.

6 Conclusion and Future Work

This paper presents a collaborative alert system where different domains can participate in building a cooperative knowledge of alerts. This fact will improve the accuracy in detecting distributed threats by sharing isolated alerts individually detected, thereby providing a high level overview of the entire system. To provide a better accuracy, the proposed Collaborative Intrusion Detection Network (CIDN) is strengthened with an inter-domain reputation mechanism, which is capable of computing the reputation of mobile HIDSs when they travel across security domains.

As a statement of direction, our intention is to continue working in some aspects that have remained open in this proposal. Among them, we indicate here two open issues in which we are interested: first, how to compute the trustworthiness or reliability one security domain has on another for calculating the reputation of a mobile HIDS; and, second, how a domain can estimate the weights α and β , according to its own experiences with the mobile HIDS and those that are provided by other domains.

Acknowledgment. This work has been partially supported by the SEMI-RAMIS EU-IST project (Secure Management of Information across multiple Stakeholders), with code CIP-ICT PSP-2009-3 250453, within the EC Seventh Framework Programme (FP7), by the Spanish MEC as part of the project TIN2008-06441-C02-02 SEISCIENTOS and by a Séneca Foundation grant within the Human Resources Research Training Program 2007 (code 15779/PD/10). Thanks also to the Funding Program for Research Groups of Excellence granted by the Séneca Foundation with code 04552/GERM/06.

References

1. Michiardi, P., Molva, R.: CORE: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In: Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, pp. 107–121 (2002)
2. Srinivasan, A., Teitelbaum, J., Liang, H., Wu, J., Cardei, M.: Reputation and Trust-based Systems for Ad Hoc and Sensor Networks. On Trust Establishment in Mobile Ad-Hoc Networks. John Wiley & Sons Ltd., Chichester (2007)
3. Zhou, C.V., Leckie, C., Karunasekera, S.: A survey of coordinated attacks and collaborative intrusion detection. *Computers & Security* 29, 124–140 (2010)
4. Trend Micro, Inc. OSSEC: An open source host intrusion detection system, <http://www.ossec.net>
5. Sourcefire, Inc. Snort: An open source network intrusion prevention and detection system, <http://www.snort.org>
6. Sun, Y.L., Yang, Y.: Trust establishment in distributed networks: Analysis and modeling. In: ICC 2007: Proceedings of the IEEE International Conference on Communications (June 2007)
7. Fernandez-Gago, M.C., Roman, R., Lopez, J.: A survey on the applicability of trust management systems for wireless sensor networks. In: SECPerU 2007: Proceedings of the Third International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing, pp. 25–30 (July 2007)
8. Buchegger, S., Boudec, J.Y.L.: A robust reputation system for P2P and mobile ad-hoc networks. In: Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems (June 2004)
9. Omar, M., Challal, Y., Bouabdallah, A.: Reliable and fully distributed trust model for mobile ad hoc networks. *Computers & Security* 28, 199–214 (2009)
10. Chen, H., Wu, H., Zhou, X., Gao, C.: Agent-based trust model in wireless sensor networks. In: SNPD 2007: Proceedings of the Eighth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, vol. 03, pp. 119–124 (August 2007)
11. Boukerche, A., Xu, L., El-Khatib, K.: Trust-based security for wireless ad hoc and sensor networks. *Computer Communications* 30, 2413–2427 (2007)
12. Gómez Mármol, F., Martínez Pérez, G.: Providing trust in wireless sensor networks using a bio-inspired technique. *Telecommunication Systems* 46, 163–180 (2011)

13. Fung, C., Zhang, J., Aib, I., Boutaba, R.: Trust management and admission control for host-based collaborative intrusion detection. *Journal of Network and Systems Management*, 1–21 (2010)
14. Zhang, Z., Ho, P.-H., Nat-Abdesselam, F.: Radar: A reputation-driven anomaly detection system for wireless mesh networks. *Wireless Networks* 16, 2221–2236 (2010)
15. Rehak, M., Staab, E., Pechoucek, M., Stiborek, J., Grill, M., Bartos, K.: Dynamic information source selection for intrusion detection systems. In: *AAMAS 2009: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, vol. 2, pp. 1009–1016 (May 2009)