

Towards a flexible trust establishment framework for sensor networks

Efthimia Aivaloglou · Stefanos Gritzalis · Charalabos Skianis

Published online: 9 September 2007
© Springer Science+Business Media, LLC 2007

Abstract Wireless sensor networks highly depend on the distributed cooperation among network nodes. Trust establishment frameworks provide the means for representing, evaluating, maintaining and distributing trust within the network, and serve as the basis for higher level security services. In this paper, we propose a trust establishment framework targeted sensor networks that can uniformly support the needs of nodes with highly diverse network roles and capabilities, by exploiting the pre-deployment knowledge on the network topology and the information flows. The framework allows for flexibility by combining aspects from alternative approaches on trust establishment on common evaluation metrics, and enables controlled trust evolution based on the network pre-configuration.

Keywords Trust · Trust establishment · Wireless sensor networks

1 Introduction

In the context of ambient intelligence systems, wireless sensor networks is a technology that can enable the provision of unobtrusive and context aware applications and services.

Sensor networks are composed of inexpensive, small and resource constrained sensor nodes, densely spread over sensing fields, that capture diverse types of contextual information related to their environment and make it available to sensor applications and services in other networks and application platforms. The application space that each sensor network application is designated for influences the contexts and types of information that is captured, the types of sensor nodes that are utilised, and the services that can be provided to the end users.

The security and integrity of the data and the communications within sensor networks is an essential requirement for the end applications and services to be reliable. Securing sensor networks generally entails ensuring the confidentiality and integrity of the data communicated, providing the means for node authentication and access control, along with lower level security issues like secure routing and node grouping. The additional requirement for trust management is set because sensor networks highly depend on the distributed cooperation among network nodes, and the assessment of trust relationships within the network could serve as the basis for higher level security solutions, such as trusted key exchange or secure routing.

The notion of trust, as used in different research areas like trusted computing, trusted platforms, trusted code and trust management, has received various interpretations [1]. Throughout this work, we study the in-network trust relationships that can exist between network entities. A trust relationship between a trust issuer and a trust target is the result of the trust establishment process, which includes the specification of valid types of evidence, and its generation and evaluation. Trust is transitive if it can be extended beyond the two parties between whom it was established, allowing for the building-up of trust paths between entities

E. Aivaloglou (✉) · S. Gritzalis · C. Skianis
Information and Communication Systems Security Laboratory,
Department of Information and Communication Systems
Engineering, University of the Aegean, Samos, Greece
e-mail: eaiv@aegean.gr

S. Gritzalis
e-mail: sgritz@aegean.gr

C. Skianis
e-mail: cskianis@aegean.gr

that have not directly participated in a process of trust evaluation.

The trust evaluation requirements and challenges posed by sensor networks are substantially different from the case of traditional wired networks. The existence of trusted third parties used as intermediaries for establishing trust relationships cannot be taken for granted, and trust relationships change frequently due to the dynamic topology. To tackle those challenges for the case of ad hoc networks in general, distributed trust establishment frameworks have been proposed for representing, evaluating, maintaining and distributing trust. However, considering the application of those frameworks in the case of sensor networks, they are either found too computationally complex or they do not exploit the pre-deployment knowledge that will usually exist in sensor network deployments.

The trust establishment framework proposed in this work is targeted specifically for sensor networks. Its main objective is that it should be applied uniformly throughout the sensor network, being able to support through proper configuration from simple nodes that have very restricted role, computational capabilities and should only trust the nodes they are pre-configured to trust, to highly adaptive nodes and gateways to other networks. Its novel characteristics include enabling the exploitation of pre-deployment knowledge in order to adjust the supported trust characteristics for each node, allowing for the adjustment of trust degradation according to the distance from pre-established trust relationships, and combining aspects from certificate-based and behavior-based trust establishment in a unified framework.

The rest of the paper is organised as follows: Sect. 2 discusses the challenges of trust establishment and the motivation for our work. Section 3 presents the related work on trust establishment on ad hoc and sensor networks. The proposed framework, the metrics that it uses and the pre-configuration it requires are introduced and analysed in Sect. 4. Section 5 evaluates the framework against the requirements initially set. Finally, Sect. 5 concludes the paper and suggests future directions.

2 Challenges, requirements and motivation

The characteristics of sensor networks both at node and network level pose unique challenges in the design of security solutions. Sensor nodes have constrained memory, computation and communication capabilities, and limited energy supplies. Sensor networks inherit the infrastructureless nature of ad hoc networks, characterised by dynamic topology and membership changes, and lack of centralised monitoring and management points that could be used as trust managing authorities. Trust establishment schemes for sensor networks should thus support distributed and cooperative

trust evaluation, using mechanisms that entail acceptable resource consumption.

Unlike the general case of ad hoc networks, in the case of sensor networks pre-deployment knowledge on the roles of the network nodes and their trust associations will usually be available. Moreover, depending on the application space and the role of each node in the network, both its capabilities and its trust evaluation requirements can be highly diverse. Diversity can be identified in the roles of the nodes, which can be from simple sensor nodes to cluster heads and gateways to other networks, in their computational capabilities, in the type of information that they collect, in their mobility and the possibility of their regular maintenance. More importantly, depending on the application domain of the deployments, diversity exists in the level of distrust that the nodes should exhibit during the network lifecycle towards unknown parties.

Moreover, some sensor nodes may be clustered by deployment so that the trust relationships within the cluster may be assumed long-term and stable. Within predefined clusters like body sensor networks, for example, trust relationships between the nodes do not need to be continuously evaluated. Trust establishment frameworks for sensor networks can exploit the pre-deployment knowledge that will usually be available in the deployments, by allowing for the pre-configuration of stable trust relationships. At the same time, the frameworks should provide the means for restricting the set of external trusted parties of the predefined clusters, according to the level of distrust that they should exhibit, through parameterised trust evolution and spreading.

The main objective of this work is define a trust establishment framework for sensor network deployments of different purposes and application domains, that can be applied uniformly throughout the network, and can support through proper configuration from simple nodes that have very restricted role, computational capabilities and should only trust the nodes they are pre-configured to trust, to highly adaptive nodes and supervision nodes. The main objectives that we thus set for our trust establishment framework are:

1. Support and exploit the diversity in the roles and the capabilities of the nodes in the deployments by allowing for flexibility in the trust establishment process.
2. Be decentralised, not based on on-line trusted parties. Instead, it should support distributed, cooperative evaluation.
3. Support pre-established and stable trust relationships within clusters.
4. Support nodes that should exhibit varying levels of distrust towards unknown parties.

3 Related work

The trust establishment frameworks that have been proposed for ad hoc and sensor networks can be classified into two categories, namely certificate-based and behavior-based, according to their scope, purpose and type of evidence that trust evaluation is based on [2]. Certificate-based frameworks aim to define mechanisms for pre-deployment knowledge on the trust relationships within the network, usually represented by certificates, to be spread, maintained and managed either independently or cooperatively by the nodes. Trust decisions are mainly based on the provision of a valid certificate, that proves that the target node is considered trusted either by a certification authority or by other nodes that the issuer trusts. In behavior-based frameworks, each node performs trust evaluation based on continuous monitoring of the behavior of its neighbors, in order to evaluate how cooperative they are. Trust is evaluated both independently by each node based the statistical data that is being continuously accumulated, and cooperatively through sharing recommendations and spreading reputation.

The main challenge confronted by certificate-based frameworks for ad hoc networks is the lack of pre-established infrastructure, which hinders the use of on-line certification authorities. In the framework proposed in [3], trust is represented by certificates signed by off-line certification authorities, whose public keys the trustors maintain locally in order to verify the signatures. Hubaux et al. [4] propose a distributed public key management scheme, where trust is evaluated using certificate chains similarly to the “web of trust” approach of the PGP model, with the difference that each node maintains locally a subset of the trust graph. In the mobile certification authority framework, presented by Yi and Kravets [5], secret sharing mechanisms are used to distribute trust to an aggregation of nodes that can collaboratively provide certification authority services. The distributed trust establishment framework proposed by Eschenauer et al. [6] takes a broader view on the inputs required for trust decisions by accepting as trust evidence not only certificates and public keys, but also information like identities, locations, or independent security assessments.

Trust in behavior-based frameworks is formulated as a combination of the direct trust value to the target node, which is evaluated independently by the trust issuer based on previous interactions and network traffic monitoring metrics, and the indirect trust value derived from the recommendations of neighboring nodes. In the reputation-based framework for sensor networks [7], a watchdog mechanism is used for monitoring the behavior of neighboring nodes in terms of data forwarding and raw sensing data consistency, and a Bayesian formulation is proposed for representing node reputation and trust evolution. Huang et al. [8] developed a trust evaluation model targeted for sensor networks,

where the Dempster-Shafer Theory of Evidence is proposed for combining recommendations. Confidence values are assigned along with the recommendations in [9], where trust and confidence values are mapped in a trustworthiness composite metric, and [10], where the trust inference problem is formulated as a shortest path problem on a weighted directed graph and theory of semirings is being used.

It is our belief, however, that both the behavior-based and the certificate-based frameworks that have been proposed are better targeted for ad hoc than for sensor networks. The main reasons are that they do not exploit the pre-deployment knowledge that will usually be available in sensor network deployments, and they do not allow for pre-established, stable trust relationships. None of the behavior-based frameworks includes any bias with respect to the identity of the node under evaluation. From the certificate-based frameworks, this requirement could be satisfied by the framework proposed in [6] through introducing identity related bias in the trust metrics and policies of the nodes, and [4], through appropriate selection of the locally stored subsets of the trust graph.

Moreover, the computational complexity of the certificate-based and the energy requirements of the behavior-based trust evaluation frameworks raise concerns related to their applicability on resource constrained sensor nodes. The former category utilises asymmetric cryptography, that is considered too expensive for sensor nodes [11, 12]. However, Elliptic Curve Cryptography, that has recently emerged as an attractive alternative to traditional public key generation, is considered to be efficient enough to be attained and executed on resource-constrained sensor nodes, mainly due to the fact that it can offer equivalent security with smaller key sizes [12]. The frameworks in the latter category are resource consuming in terms of computation, memory and energy, since they require the radio on each node to be continuously on, and the trust values of the neighboring nodes to be stored and continuously updated as interactions occur.

4 Trust establishment in the diverse environment of sensor networks

4.1 Overview

The differences in scope and purpose between the two categories of trust establishment frameworks as discussed in Sect. 3, show that they should not be viewed as alternative approaches, but as supplementary. In our framework, we adopt aspects both from certificate-based and behavior-based trust establishment, in order to benefit both from the representation of pre-deployment trust relationships as certificates and from the continuous behavior-based evaluation of trust.

Table 1 Trust establishment evidence and evaluation

Trust relationship between i, j	Evidence	Evaluation
Pre-established	Stored $T_{ij} \leq 1, R_{ij} \leq 1$	Not required
Hierarchical, trust managing authority x	Stored $T_{ix} \geq T_{threshold}$, stored $R_{ix} \geq R_{threshold}$, stored public key of x , signed certificate of j	Validation of certificate $\Rightarrow T_{xj} = 1$ used as a recommendation
Distributed, set N_i of neighboring nodes and supervision nodes	Stored $T_{ix} \geq T_{threshold}$, stored $R_{ix} \geq R_{threshold}$, $T_{xj}, \forall x \in N_i$	Combination of recommendations

Our framework enables the use of certificates signed by offline trust managing authorities [3] for trust establishment by a subset of the network nodes. For certificate validation to be performed locally, each node needs to store the public keys of the trust managing authorities that it is pre-configured to trust and that issued the certificates of the target nodes. Trust associations can thus be evaluated between nodes that are associated with common trust managing authorities that issue the certificates for particular deployments.

Behavior-based trust evaluation can be performed only by a subset of the nodes, called supervision nodes. Those are designated to monitor the network traffic, evaluate the nodes within their range according to their behavior in network and data level, and spread information on their direct trust [7, 8] for the target nodes. The results of behavior-based evaluation are thus provided as a network service by the supervision nodes, so as not to consume the resources of the entire network.

As discussed in Sect. 2, the main objective of our trust establishment framework is to support the diversity in the roles and the capabilities of the nodes in the deployments by allowing for flexibility in the trust establishment process. The trust associations between any trust issuer i and any trust target j that the framework supports can be:

1. Established prior to deployment through storing locally at each node information on its trust associations.
2. Established as hierarchical trust relationships so that each node j is considered trusted by node i if it holds a valid certificate that i can verify using the stored public key of an offline trust managing authority that it has a trust association with.
3. Established by a cooperative procedure, where i asks for recommendations for j from nodes that it has a trust association with.
4. Evaluated and made available by supervision nodes that perform behavior-based trust evaluation and i has a trust association with.

The parties that may be involved in the trust establishment procedure are thus offline trusted third parties whose public key is locally stored for signature verification, other sensor nodes, cluster heads or gateways, and supervision nodes that perform behavior-based trust evaluation. Table 1 describes the supported trust evidence for each type of trust evaluation. Once the trust relationship of node i with node j needs to be determined, the options on Table 1 can be used. If a trust relationship is not already established either before deployment or as a result of a previous trust establishment procedure, node i first attempts to establish a hierarchical and then a distributed trust relationship.

4.2 Trust evaluation metrics

Any node i can determine its trust relationship with any node j either by its stored list of trust associations or by evaluation based on recommendations from third parties. For generality, we take a view of a signed certificate from an offline trust managing authority as a recommendation with the highest trust value. Third parties providing recommendations are thus considered other sensor nodes, supervision nodes, and offline trust managing authorities.

A trust association stored locally at node i and referring to node j contains two metrics, namely the trust metric T_{ij} and the transition metric R_{ij} . Both of those metrics should have values above a certain threshold for i to accept recommendations from j for other nodes. The first is the trust value $T_{ij} \in [-1, 1]$ of node i for j , provided by a function that can uniformly calculate the trust value based on the recommendations from the third parties. This function is common both for hierarchical and for cooperative trust establishment. Provided the high importance of the pre-deployment knowledge that exists in sensor network deployments, T_{ij} can be equal to 1 only for trust relationships established prior to deployment between i and either other nodes or trust managing authorities. Using N_i as the set of trusted nodes that i receives recommendations T_{xj} for node j , for the evaluation of T_{ij} a function can be formulated as:

$$T_{ij} = t(T_{ix}, R_{ix}, T_{xj}, \forall x \in N_i). \quad (1)$$

There exist several choices for the function $t(\cdot)$, that should satisfy the requirement that for nodes $x \in N_i$ where $R_{ix} = 0$, T_{xj} will not be used for the evaluation, even if $R_{threshold} = 0$. An example simple rule is the weighted average of the recommendations:

$$T_{ij} = \frac{\sum_{x \in N_i} T_{ix} \cdot R_{ix} \cdot T_{xj}}{|N_i|}. \quad (2)$$

The transition metric $R_{ij} \in [-1, 1]$ is the second part of a trust association, used to indicate a weight that node i will assign to future recommendations from node j . An example of the usability of a separate metric is that, during the initial configuration of a node in a cluster, it can be greater than zero only for the cluster head, so that i accepts recommendations only from it and not from the other nodes that it trusts. This metric is also used as the means to control trust evolution and spreading according to the level of distrust that each node should exhibit during its lifetime towards unknown parties. The level of distrust is represented as a degradation parameter $d_i \in [0, 1]$, used for the calculation of R_{ij} . Setting $d_i = 1$ indicates that trust should not degrade according to the number of steps from a node that i is pre-configured to trust. Setting $d_i = 0$ should make $R_{ij} = 0$, $\forall j \in N_i$, and thus i should not calculate recommendations from nodes except the ones it is pre-configured to.

The transition metric R_{ij} is evaluated by a function accepting as parameters d_i and the transition values of the nodes in N_i :

$$R_{ij} = r(R_{ix}, d_i, \forall x \in N_i). \quad (3)$$

The function $r(\cdot)$ should enforce the degradation of the value R_{ij} in relation to R_{ix} according to d_i . A possible $r(\cdot)$ can be formulated as:

$$R_{ij} = \max(R_{ix}) \cdot d_i, \quad x \in N_i \Rightarrow T_{xj} \geq T_{ij}. \quad (4)$$

This function uses for the computation of R_{ij} the maximum transition value, from the nodes whose recommendations are greater than or equal to the trust value computed for j .

4.3 Pre-configuration parameters

The proposed framework can uniformly support through proper configuration from simple nodes that have very restricted role, computational capabilities and should only trust the nodes they are pre-configured to trust, to highly adaptive nodes and supervision nodes. This is achieved through the parameterisation of each node i during pre-deployment, that should include:

- Setting the pre-established trust associations through assigning values T_{ij} and R_{ij} for the nodes j that node i should trust and receive recommendations from, based on the pre-deployment knowledge of the network structure.

- Balancing the parameters $T_{threshold}$ of the minimum positive trust value, $R_{threshold}$ of the minimum transitivity value, and the degradation parameter d_i . The last two parameters are the ones that eventually determine the maximum allowed distance from pre-established trust relationships that the node can establish during the network lifetime.

For nodes that have strictly defined roles in the network or have limited computational capabilities, the set of pre-established relationships that recommendations are accepted from should be restrained through setting $R_{ij} < R_{threshold}$. An example sensor network scenario where this kind of configuration could be applied is that of a body sensor network, where sensor nodes are collecting physiological data that the cluster head or gateway to a B3G network is aggregating and transmitting. In this scenario, it should be allowed only for the cluster head c to expand the trust relationships in the cluster. For this reason, for the sensor nodes only R_{ic} should be set above the threshold, and d_i should be set to zero. The initial trust associations of the cluster head could allow it for more flexibility, according to its role and the computational capabilities.

5 Evaluation against requirements

The requirements that were initially set for the proposed framework were intended to address the perceived trust establishment needs of real-world sensor network deployments. Depending on the application space of each sensor network, diversity is expected to exist in the roles and the capabilities of the nodes. In order to support this diversity and fulfil the first objective in Sect. 2, the trust establishment process was designed to be flexible by providing alternative options for trust evaluation and combining them on common evaluation metrics. Moreover, in order for the framework to be applicable to resource constrained nodes, it enables the restriction of the alternative options for trust establishment through the proper configuration based on pre-deployment knowledge on the network topology and the information flows. The framework can uniformly support from highly adaptive nodes to static and restricted nodes, that will never during the network lifetime need to perform certificate validations or combinations of recommendations.

Regarding the second objective that was initially set, the framework supports both hierarchical trust establishment based on offline trust managing authorities and cooperative trust establishment based on recommendations from other nodes and supervision nodes. It fulfils the third requirement by supporting pre-established and stable trust relationships within clusters according to the initial configurations. Finally, it enables control over the evolution of trust relationships through the use of the degradation metric. This metric

represents the level of distrust that each node should exhibit towards unknown parties, thus its value depend on the application domain and the role of the node.

Evaluated against the supported trust characteristics identified in [2], the proposed framework does not include support for uncertain evidence, since it does not support assignment of confidence values to evidence supplied for trust evaluation, including the recommendations. This would be beneficial especially for the recommendations provided by supervision nodes. The framework supports controlled trust transitivity, since the trust values from third parties are weighted according to the trust relationship the requester has with the third party. Trust revocation, which is characterised as controlled if either trust is revoked only by trusted third parties or some mechanism exists to protect from defaming attacks, is not supported by the framework yet.

The evaluation of the complexity and computational requirements of the framework highly depends on the type of each node and its pre-configuration. It is considered that high computational power would be required to perform public key operations and certificate validations, or to continuously monitor surrounding nodes and re-evaluate trust relationships based on every event monitored. The first would only be required by highly adaptive nodes that are pre-configured to support hierarchical trust establishment, while the latter should only be performed by supervision nodes. In the actual sensor network deployments, however, it is expected that the nodes that would be designated for those roles would usually be computationally more powerful than the sensor nodes.

6 Conclusions and future work

The trust establishment framework for sensor networks proposed in this work fulfills its main objectives that it should be applied uniformly throughout various sensor network deployments, and that it should support through proper configuration the diverse characteristics and needs of sensor nodes. It both allows for flexibility and for restriction of the supported trust characteristics by allowing for configuration based on pre-deployment knowledge on the network topology and the information flows.

It is incomplete, however, regarding an important aspect of the trust management problem, the revocation of trust relationships. Particularly for the case of sensor networks, that are susceptible to node misbehavior, this is a non trivial issue. Malicious nodes may perform defaming attacks against legitimate nodes to spread bad reputation, either by directly spreading false evidence or by pretending to be victims of defaming attacks themselves to make a legitimate node look malicious [11]. The design and integration of distributed and controlled trust revocation mechanisms in the framework is our future work direction.

Acknowledgements This paper describes work undertaken in the context of the e-SENSE project, Capturing Ambient Intelligence for Mobile Communications through Wireless Sensor Networks (www.ist-e-SENSE.org). e-SENSE is an Integrated Project (IP) supported by the European 6th Framework Programme, contract number: 027227.

References

- Gollmann, D. (2006). Why trust is bad for security. In *Electronic notes in theoretical computer science* (Vol. 157, pp. 3–9). Elsevier.
- Aivaloglou, E., Gritzalis, S., & Skianis, C. (2006). Trust establishment in ad-hoc and sensor networks. In *Lecture notes in computer science: Vol. 4347. Proceedings of CRITIS'06 1st international workshop on critical information infrastructure security* (pp. 179–194).
- Davis, C. R. (2004). A localized trust management scheme for ad hoc networks. In *3rd International conference on networking (ICN'04)* (pp. 671–675).
- Hubaux, J. P., Buttyán, L., & Capkun, S. (2001). The quest for security in mobile ad hoc networks. In *MobiHoc '01: Proceedings of the 2nd ACM international symposium on mobile ad hoc networking & computing* (pp. 146–155). New York: ACM Press.
- Yi, S., & Kravets, R. (2003). Moca: Mobile certificate authority for wireless ad hoc networks. In *Proceedings of 2nd annual PKI research workshop*.
- Eschenauer, L., Gligor, V. D., & Baras, J. S. (2002). On trust establishment in mobile ad-hoc networks. In *Security protocols workshop* (pp. 47–66).
- Ganeriwai, S., & Srivastava, M. B. (2004). Reputation-based framework for high integrity sensor networks. In *Proceedings of the 2nd ACM workshop on security of ad hoc and sensor networks (SASN'04)* (pp. 66–77). ACM Press.
- Huang, L., Li, L., & Tan, Q. (2006). Behavior-based trust in wireless sensor network. In *APWeb workshops* (pp. 214–223). Berlin: Springer.
- Zouridaki, C., Mark, B. L., Hejmo, M., & Thomas, R. K. (2006). Robust cooperative trust establishment for manets. In *SASN '06: Proceedings of the fourth ACM workshop on security of ad hoc and sensor networks* (pp. 23–34). New York: ACM Press.
- Theodorakopoulos, G., & Baras, J. S. (2004). Trust evaluation in ad-hoc networks. In *Workshop on wireless security* (pp. 1–10).
- Shi, E., & Perrig, A. (2004). Designing secure sensor networks. *Wireless Communication Magazine*, 11(6), 38–43.
- Arazi, B., Elhanany, I., Arazi, O., & Qi, H. (2005). Revisiting public-key cryptography for wireless sensor networks. *IEEE Computer*, 38(11), 103–105.



Eftimia Aivaloglou holds a Diploma in Information and Communication Systems Engineering from the University of the Aegean, Greece, and an MSc in Advanced Computer Science from the Department of Computer Science, University of Manchester, UK. She is currently a Ph.D. candidate in the Department of Information and Communication Systems Engineering at the University of the Aegean. She is working on the field of information and communication systems security, and her research focuses on security, trust and privacy in wireless ad hoc and sensor networks.



Stefanos Gritzalis (www.icsd.aegean.gr/sgritz) holds a BSc in Physics, an MSc in Electronic Automation, and a PhD in Informatics all from the University of Athens, Greece. Currently he is an Associate Professor, the Head of the Department of Information and Communication Systems Engineering, University of the Aegean, Greece and the Director of the Laboratory of Information and Communication Systems Security (Info-Sec-Lab). He has been involved in

several national and EU funded R&D projects in the areas of Information and Communication Systems Security. His published scientific work includes several books on Information and Communication Technologies topics, and more than 140 journal and national and international conference papers. The focus of these publications is on Information and Communication Systems Security. He has served on program and organizing committees of national and international conferences on Informatics and is an editorial advisory board member and reviewer for several scientific journals. He was a Member of the Board (Secretary General, Treasurer) of the Greek Computer Society. He is a member of the ACM and the IEEE. Since 2006 he is a member of the 'IEEE Communications and Information Security Technical Committee' of the IEEE Communications Society.



Charalabos Skianis is currently Assistant Professor in the Department of Information and Communication Systems at the University of the Aegean in Samos, Greece. He holds a PhD degree in Computer Science, University of Bradford, United Kingdom and a BSc in Physics, Department of Physics, University of Patras, Greece. His work is published in journals, conference proceedings and as book chapters and has also been presented in numerous

conferences and workshops. He acts within Technical Program and Organizing Committees for numerous conferences and workshops (e.g., IFIP Networking 2006, IEEE Globecom 2006, IEEE ICC 2006) and as a Guest Editor for scientific journals (e.g., IEEE Networks magazine). He is at the editorial board of journals (e.g., IEEE Wireless Communications), a member of pronounced professional societies (senior member of IEEE) and an active reviewer for several scientific journals. He is an active member of several Technical Committees within the IEEE ComSoc [TC II; TC CSIM-Editor in Chief for the forthcoming eNewsletter; TC ComSoft], and member of IEEE BTS; IEEE TVT and IEEE CS.