

A Trust-based Approach for Secure Data Dissemination in a Mobile Peer-to-Peer Network of AVs

BHARAT K. BHARGAVA¹, PELIN ANGIN¹, ROHIT RANCHAL¹, RANJITKUMAR SIVAKUMAR¹, MARK LINDERMAN² AND ASHER SINCLAIR²

¹Purdue University, ²Air Force Research Laboratory

Mobile peer-to-peer networks of aerial vehicles (AVs) have become significant in collaborative tasks including military missions and search and rescue operations. However, the nature of the communication between the nodes in these networks makes the disseminated data prone to interception by malicious parties, which could cause serious harm for the designated mission of the network. In this paper, we propose an approach for secure data dissemination in a mobile peer-to-peer network, where the data disclosed to a particular node in the network depends on the trustworthiness of that node as well as the matching of policies of the data source and destination. We demonstrate the use of active bundles for protecting sensitive data as they are sent from one node to another, on the simulation we developed for data dissemination in a mobile peer-to-peer network of AVs. We also discuss filtering techniques for dissemination of sensitive data in such networks.

Keywords: Aerial vehicle, trust, security, data dissemination, active bundles, data filtering

1. INTRODUCTION

Aerial Vehicles (AVs) have assumed increasingly important roles in the military context as well as for search and rescue operations since their invention. The quality and trustworthiness of the data collected by AVs are of utmost importance since actions taken based on the data could have serious consequences. Data dissemination between AVs is a way to optimize the use of resources and improve the quality of the collected data to attain a complete picture of the task environment. While sharing of the collected data facilitates achieving the common task, uncontrolled sharing could lead to leakage of information to adversaries, impairing mission accomplishment. Data is broadcast over a shared communication media in wireless Vehicular Ad hoc Networks, so it is possible for a malicious node to intercept or modify data, or to inject erroneous data. The open nature of a VANET thus renders communication security a great challenge as in described in [Blum et al., Zarki et al., Hubaux et al., Parno et al., Aijaz et al.].

In this paper we propose a scheme for secure dissemination of data in a mobile peer-to-peer network of AVs, which filters data to be shared with an AV based on the dynamically determined trustworthiness of that AV and the policies of the source and destination AVs.

2. RELATED WORK

VANET (Vehicular Ad-hoc Network) communication is an important emerging research area. Most of the previous research in this field focused on the development and enhancement of communication protocols (such as [Yang et al.]). The open nature of a VANET, which renders communication security a great challenge, calls for work in secure data dissemination techniques. There have been some system proposals addressing the security problem. Raya et al. proposed a system [Raya 2006, Raya 2007] that allows vehicles to securely communicate using base stations placed for the infrastructure setup, key distribution and revocation. The public key infrastructure (PKI) deployment for this system is a large-scale and potentially costly procedure since it requires

Authors' addresses: Bharat Bhargava, Pelin Angin, Rohit Ranchal, Ranjitkumar Sivakumar, Purdue University, Dept. of Computer Science, 305 N. University St., West Lafayette, IN, U.S.A.; Mark Linderman, Asher Sinclair, Air Force Research Laboratory, Rome, NY, U.S.A.

large-scale testing after deployment to ensure operation under real-world VANET conditions. Further the solution really only aims ensuring authentication (of pseudonyms) but a security infrastructure must be aimed at establishing the authenticity of message contents for safety and security. Some security architectures have been proposed to address specific attacks. Fonseca et al. [Fonseca 2007] address the problem of privacy in a VANET with the help of infrastructures (base station and certification authorities) and pseudonym use. Furgel and Lemke [Furgel 2004] describe the security of the electronic systems in a vehicle that are actually responsible for transporting or generating the data before it is sent. Golle et al. [Golle 2004] focus on detecting and correcting malicious data in VANETs. Another proposal is the use of a PKI and a virtual infrastructure by Blum and Eskandarian in [Blum 2004], where cluster-heads are responsible for reliably disseminating messages after digitally signing them. However this can only counter the intentionally caused intelligent collisions. Also, this approach creates bottlenecks at cluster-heads in addition to high security overhead.

Trust-based systems for peer-to-peer data sharing have been studied by researchers including [Selcuk et al., Ooi et al., Aberer et al.]. However, most of these proposals rely on the history of interactions with a particular node for determining its trustworthiness, which could be misleading if the context of the node changes.

3. PROPOSED DATA DISSEMINATION SCHEME

In order to provide secure data dissemination between AVs, we propose a trust-based approach, where the data sent to a particular AV depends on the context of the interaction (such as normal conditions/emergency/disaster) and the level of trustworthiness of the AV as determined by a set of factors such as the history of interactions with that AV, location, distance, authentication level etc. While the trust level of an AV determines whether data will be shared with it, the granularity/quality of the data is determined by the sharing policy of the data owner. Data is transformed / filtered based on the rules specified in the dissemination policy of the data owner before being shared with the correspondent AV to minimize the risk of leakage of confidential information to malicious or unauthorized parties. The main components of the proposed data dissemination scheme are active bundles, dynamic trust calculation and data filtering, described in the following sections.

3.1 Data Protection Mechanism: Active Bundle

An Active Bundle (more detailed description can be found in [Lilien 2006, Ben-Othmane 2009, Ben-Othmane 2010]), the structure of which is shown in Figure 1 below, is a data protection mechanism that encapsulates sensitive data with metadata and a virtual machine.

Sensitive data constitutes content to be protected from privacy violations, data leaks and unauthorized dissemination. Metadata describes the active bundle and its privacy policies. The metadata includes (but is not limited to) the following components: (a) provenance metadata; (b) integrity check metadata; (c) access control metadata; (d) dissemination control metadata; (e) life duration value. The virtual machine (VM) manages and controls the program enclosed in a bundle. The main VM functions include (a) enforcing bundle access control policies through apoptosis (self-destruction) or data filtering (b) enforcing bundle dissemination policies; and (c) validating bundle integrity [Lilien 2006, Ben-Othmane 2009, Ben-Othmane 2010]. With the assumption that there is a secure communication channel between the AB and the trusted third parties, below we provide a description of the lifecycle of an active bundle (AB).

Initialization of an AB: An owner of sensitive data constructs an AB by putting together data, metadata (including access control and dissemination control metadata), and adding a virtual machine. After this stage, the AB becomes an active entity (since it has its own virtual machine) that can perform the remaining steps of this algorithm.

Building an AB: The steps taken in the process of building an AB are as follows:

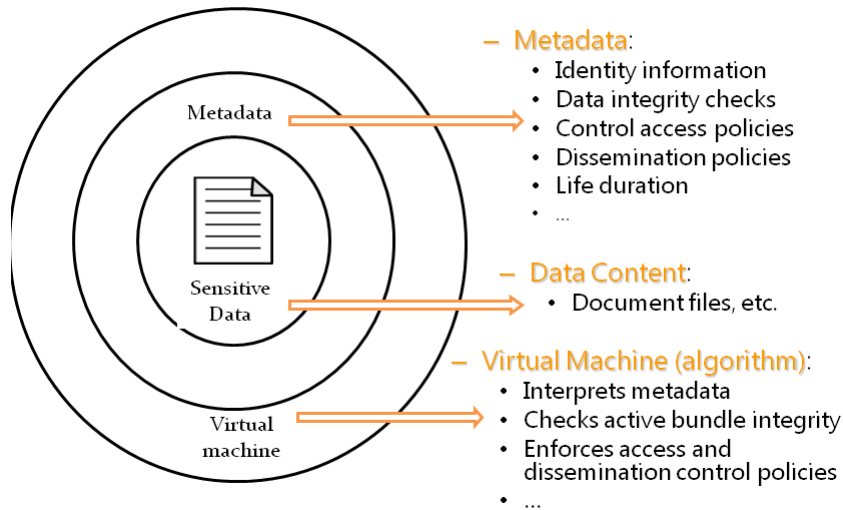


Figure 1: Active bundle structure.

1. The AB gets two pairs of public/private keys from a Security Service Agent (SSA) (trusted third party), where the first pair of keys is used for encrypting the data included in it and the second pair of keys is used for signing/verifying the data signature. The reason for having two key pairs is to prevent attackers from modifying AB’s sensitive data and signing it again with the private decryption key of the data owner.
2. The AB sends a request to SSA asking it to record the AB’s security information. The AB’s identity data includes its name, a decryption key, and the trust level that a host must satisfy to use the AB. The goal is to keep the decryption keys and other auxiliary data for ABs in a trusted location. The decryption keys are given only to hosts that are eligible to access the AB.
3. The AB computes a hash value for sensitive data and signs them using the signature key. The signature certifies that sensitive data is from its owner.
4. The AB encrypts sensitive data using the encryption key.

Enabling an AB: After arriving at the destination host, AB enables itself. The steps of the enabling algorithm seen in fig. 2 are as follows:

- Stage 1: AB sends a request to SSA asking for the security information on AB and the host’s trust level.
- Stage 2: AB checks if the host’s trust level is lower than the minimal trust level required for AB access. If so, the AB apoptosizes; otherwise, it executes the next step.
- Stage 4: AB checks integrity of its sensitive data. It computes the hash value for sensitive data and it verifies the AB’s signed hash value by comparing it to the computed hash value. If verification fails, AB apoptosizes; otherwise, the AB decrypts the data.
- Stage 7: AB enforces its privacy policies.
- Stage 8: AB provides the output to the host.
- Stage 9: AB sends audit information to an Audit Services Agent. This information includes AB’s name, the host’s identity, and the name of the event being audited (“the move to the host”).

The main challenge with AB approach is assuring that a visited host executes the AB’s VM code correctly. We assume that any host receiving an AB has a Trusted Platform Module (TPM) which assures correct execution of the VM code. Further the AB’s VM code runs in a restricted sandbox on the host which allows fine-grained control over the actions that code within the VM is permitted to take. [Lilien 2006, Ben-Othmane 2009, Ben-Othmane 2010]

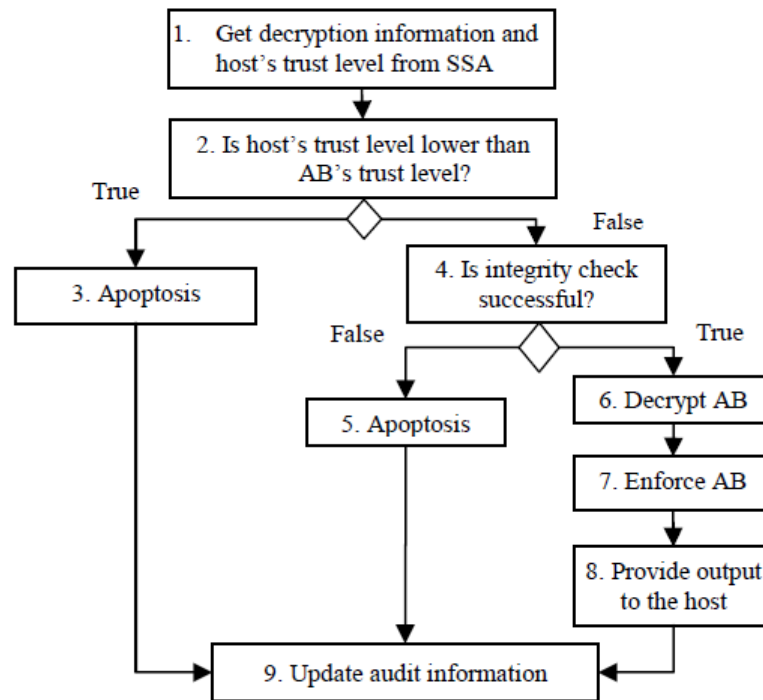


Figure 2: A UML activity diagram for enabling of an active bundle.

present a detailed description on this approach. In the scenario of data sharing between AVs discussed in this paper, the data item to be shared is packed in an active bundle along with policy information and the bundle becomes active on the receiver AV (the data item in the active bundle is encrypted to prevent unauthorized access). Upon activation, the trust level of the AV is calculated (by communicating with a third-party trust server, which can be a ground controller) and if the threshold trust condition is met, the virtual machine of the active bundle applies the policy, performing the necessary filtering on the data before it is shared with the receiver AV. If the threshold trust condition is not met, the active bundle destroys itself, therefore preventing leakage of any sensitive data.

3.2 Data Filtering

The active bundle mechanism, more specifically the virtual machine of the active bundle, can be used to filter many different types of data in different ways based on the access rights of the data receivers. While for image data, the filtering could make use of various image processing algorithms available in the literature as well as providing a zoomed in/out version of the data, filtering on relational databases would make use of query languages to provide a restricted view. Fig. 3 below demonstrates a simple example of filtering on complex data consisting of an image, geographical coordinates of the location the picture was taken at and the name of the location. In this example, the data owner's policy is to share the complete data (the leftmost picture) with an AV with the highest trust level, filter the geographical coordinates out when sharing the data with an AV of medium trust value (middle picture) and filter out both the location name and the geographical coordinates when sharing with an AV of low trust value (the rightmost picture).

Below are techniques for filtering image data collected by AVs. These techniques are preferred because they allow for automatic filtering of images without human intervention.

Low Dynamic Range Rendering: This method applies the reverse of high dynamic range



Figure 3: Example of data filtering.

rendering [Reinhard et al.] on an image to degrade image quality and hide details. For filtering, this procedure can be utilized by changing the intensity values of the pixels in an image such that the contrast ratio of the image is lowered by making pixels take on intensity values from a narrower range than that of the original image. This approach would force very dark pixels to blend in with relatively lighter ones around and light pixels to blend in with relatively darker ones to help hide the full details of the image while preserving the edges. For example, the image seen on the left side of figure 4 below can be transformed to the image on the right side using this process. As we see in this figure, details like the color and some features like the faces of the children and fine grain details of the objects around them are hidden after the transformation, while the basic shapes and edges from the original image are preserved.



Figure 4: Example of low dynamic range rendering.

Pattern Recognition and Blurring: Another method that can be used for filtering images is the recognition of specific patterns in the image to block out those high sensitivity areas. This approach would require the filtering algorithm to be supplied with a list of images that should appear blurred or as a rectangle of uniform color (e.g. completely black) if detected as a sub-image in any of the images to be filtered. For the detection of the specific patterns, a simplistic and efficient approach would be to use Haar features [Papageorgiou et al.], which considers adjacent rectangular regions at a specific location in a detection window, sums up the pixel intensities in these regions and calculates the difference between them. Recognition of the pattern being searched would then be achieved by comparing this difference calculated for the pattern image and the different sub-regions of the image being filtered.

Data Equivalence Techniques: The techniques as mentioned in [Annamalai et al.] can be used to transform an image such that the information content of the image remains the same while the fine grain details change. The most interesting among these techniques is the information content level data equivalence, which is defined by what the image represents as opposed to how it is represented. For example, if we have an aerial image, part of which has an aircraft of a

specific model, and we need to hide the details of which model the aircraft is, we can replace the image of the aircraft with that of another aircraft of a different model in a such a way that it preserves the uniformity of the complete picture. By transforming the image this way, we will be preserving the information content while hiding the specific details.

3.3 Dynamic Trust Calculation

The trust calculation component of the proposed scheme works like a reputation based system, where the trustworthiness of a node in the network is evaluated based on various dynamic parameters. The parameters include the history of transactions, bandwidth of communication, physical distance of AVs, the context of communication, and authentication level of the AV. These parameters vary with the scenario in which the AVs communicate, and have different weights assigned to each of them, resulting in an equation with the form:

$$P = \beta_1 \times F_1 + \beta_2 \times F_2 + \dots + \beta_n \times F_n$$

where n is the number of different factors considered in the evaluation of trust, each F_i is a real value in a specific range (with a high value signifying a higher trust value based on that factor) and each β_i is the weight associated with the corresponding factor F_i (the higher the weight, the more important that factor is for trustworthiness). The sum of all β_i 's is 1. The computed net value is mapped onto a real scale to determine the trust level of the destination AV (the same calculation can be used by the receiver AV to assess the trustworthiness of the sender node, hence validity of the received data). The computed trust value is used to determine whether it is safe to share data with the destination AV as well as the degree of filtering to be applied on the data before sharing. The dynamicity of the trust calculation (as opposed to techniques basing trust solely on the history of interactions) is especially important in this scheme, as a once trusted node can become compromised due to its changing context (Here context includes factors such as location, situations like attacks, war etc).

For trust calculation, an approach similar to [Can 2007] is used, where the effect of the history of transactions with a particular AV fades with time. The trust value for an AV a at time t is calculated using the equation:

$$T_a(t) = \alpha \times T_a(t-1) + (1-\alpha) \times P$$

where $\alpha < 0.5$ and P is value of trust as determined by the other parameters mentioned above. In the absence of any previous communication with a particular AV, trusted peers can be queried (if no global trust server can provide this information) as in [Can 2007] to learn whether they had previously communicated with the AV in question and a trust value can be obtained from them. If no previous history about the AV exists, the equation is adjusted to only take into account the other trust parameters.

4. DATA DISSEMINATION MODELS

The secure data dissemination scheme proposed can be used in the two main peer-to-peer data sharing models described below.

4.1 Direct Link

The direct link model is the basic peer-to-peer data sharing model, where AVs discover each other through broadcast of ALIVE messages as mentioned in [Meka et al.] and initiate data transfer without involvement of any third party nodes. In this model, the (data) source AV creates an active bundle including the sensitive data and dissemination policies and sends it to the peer AV. Once the active bundle reaches the destination, it can calculate the trust level of the destination either by contacting a trusted third party (such as the ground controller) or solely using local

parameters (as discussed in section 3.3) and run the rest of the active bundle enabling steps described in the previous section.

4.2 Publish-Subscribe

The publish-subscribe model is different from the direct link model in that it requires a third-party (such as a ground controller) called the Information Broker (as seen in Figure 5 below) to mediate the data dissemination between AVs. In this model, a publisher node registers an active bundle with the data content and dissemination policies at the Information Broker (IB) and a subscriber node subscribes with the IB for receiving data with particular content. The IB in this case is responsible for executing the trust calculation algorithm and the determining the matching between the policies of the publisher and subscriber nodes. This dissemination model can be implemented using publish-subscribe techniques for P2P networks described in [Pham et al.].

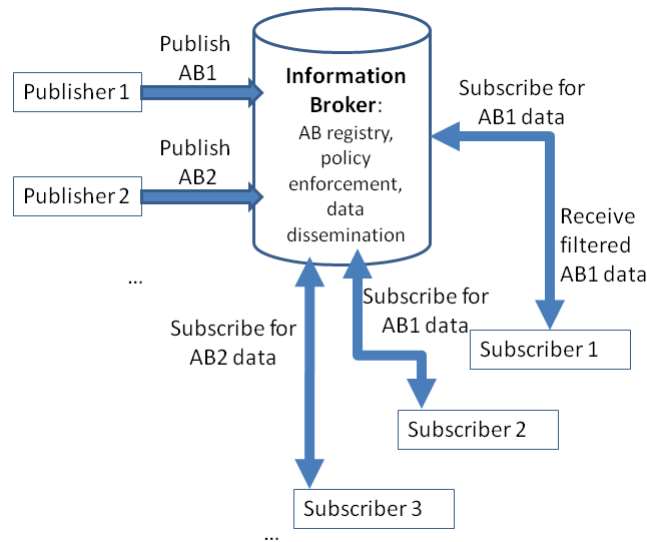


Figure 5: Publish-Subscribe data dissemination model.

5. AV P2P NETWORK SIMULATION

We simulated data dissemination in a small mobile peer-to-peer network of AVs (following an approach similar to that of Meka et al. [Meka 2010]) in an environment where a communication link is formed opportunistically between any two AVs if they are within a threshold distance of each other. The bandwidth of the communication link between the AVs is inversely proportional to the distance between the AVs as well as the number of communication links they have formed with other AVs. Each AV has a policy associated with the data items it owns, which is used to determine whether to share the data item with another AV and the level of filtering to be applied to the data before it is shared. The simulation implemented in Java consists of the modules below.

5.1 Main Module

The main module is responsible for coordinating events for the entire simulation as well as the graphical display of the progress of the simulation. The main events in the simulation are the movement of AVs, formation of communication links, bandwidth updates and data transfer between AVs. Figures 6a and 6b below present the graphical interface for the simulation. As the simulation proceeds, periodically one of the AVs chosen randomly among the ones in the scene

moves to a location with random displacement in the geographical coordinates (displacement being upper bounded by a threshold value). This displacement causes formation of communication links with the AVs that are now within a threshold distance and tear down of connection with AVs that are too far to communicate with. The recently moved AV then initiates a data transfer to one of the AVs in its neighborhood. Figures 6.a, 6.b, 6.c and 6.d below provide a step-by-step explanation of a sample data transfer between two AVs.

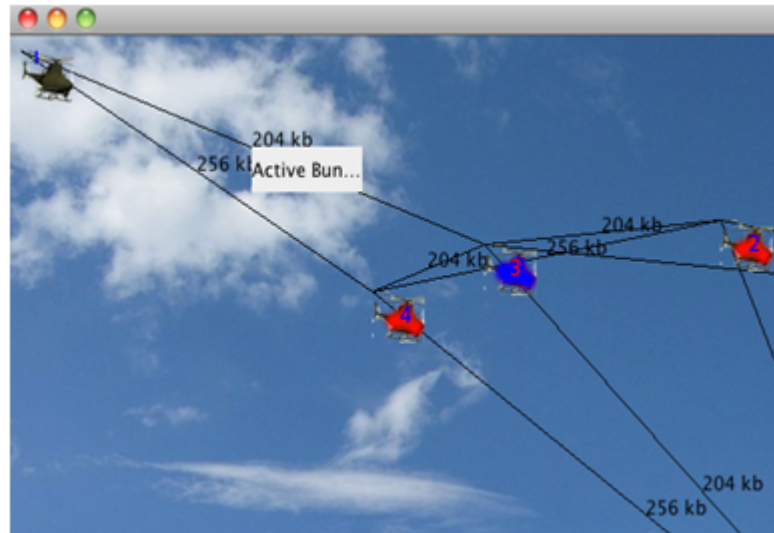


Figure 6.a: AV network. Data transfer is initiated from AV3 to AV1. Available bandwidths are displayed on the lines connecting pairs of AVs.

The details of the data transfer are controlled by the trust calculation and active bundle modules as described below.

5.2 Trust Calculation Module

This module is responsible for calculating the trustworthiness of a specific AV based on a set of parameters (the type/color of the AV, distance of the AV to the sender, bandwidth of the communication link between the AVs, and the authentication level of the AV). The context of the AV is also taken into consideration in this calculation to assign a higher trust value in case of certain emergencies such as fire or earthquake, as information sharing becomes vital in those situations and the sender AV would need to disseminate the information as fast as possible with less concern about the possibility of leakage. This module can be considered as a trusted-third party (such as a ground controller), which an AV needs to contact to verify the credentials of another AV it is communicating with. The module outputs a real trust value between a preset lower and upper bound.

5.3 Active Bundle Module

The active bundle module is responsible for the protection of sensitive data during the data sharing process between two AVs as described earlier. This module was implemented using the Jade mobile agent framework (<http://jade.tilab.com/>).

6. CONCLUSION

In this paper, we proposed an approach for secure data dissemination in a mobile peer-to-peer network of AVs, where the data disclosed to a particular node in the network depends on the trustworthiness of that node as well as the matching of policies of the data source and destination.



Figure 6.b: This figure shows the policy of data sharing at the top, original data in the middle and the virtual machine status at the bottom. In this example, the policy is based on the trust level of the AV: If it is above 2.5, the original data is shared; if it is below 2.5 but above 2.3, minimal filtering is applied; if it is between 2.3 and 2.0 greater filtering is applied and if the trust level is below the threshold of 2.0, no data is shared, in which case the active bundle destroys itself.



Figure 6.c: The trust level of the receiver AV is calculated as 2.099, which is higher than the threshold trust level, but not high enough to share the original data.




Figure 6.d: Data is transformed by the virtual machine of the active bundle according to the policy and the transformed data is shared with the receiver node. As seen here, the data shared provides a narrower view of the environment than the original image.

We described the use of active bundles for protecting sensitive data as it travels between AVs and discussed two possible models of secure data dissemination. The proposed data protection

scheme was implemented as a simulation using the Jade mobile agent framework, which enables easy integration into real-world systems. Future work will involve performance evaluation of the proposed scheme on both models of data dissemination in terms of speed and effectiveness at protecting sensitive data from malicious parties. We will also investigate the feasibility of an approach not relying on a trusted third party for trust calculation.

REFERENCES

- ABERER, K., AND DESPOTOVIC, Z.. Managing trust in a peer-2-peer information system. In *The 10th International Conference on Information and Knowledge Management, Atlanta, GA, 2001*.
- AIJAZ, A., BOCHOW, B., DOTZER, F., FESTAG, A., GERLACH, M., KROH, R., AND LEINMULLER, T.. Attacks on Inter-Vehicle Communication Systems - An Analysis. In *The 3rd International Workshop on Intelligent Transportation, 2006*.
- ANNAMALAI, M., AND BHARGAVA, B.. Defining Data Equivalence for Efficient Access of Images in a Distributed Environment. In *The 2nd World Conference on Integrated Design and Process Technology, 1996*.
- BEN-OTHTMANE, L.. Active Bundles for Protecting Confidentiality of Sensitive Data throughout Their Lifecycle. Ph.D. dissertation, Dept. of Computer Science, Western Michigan University, Kalamazoo, MI, 2010.
- BEN-OTHTMANE, L., AND LILIE, L.. Protecting Privacy in Sensitive Data Dissemination with Active Bundles. In *The 7th Annual Conference on Privacy, Security and Trust, Saint John, NB, Canada, 2009*.
- BLUM, J., AND ESKANDARIAN, A.. The Threat of Intelligent Collisions. *IT Professional*. 6, 1 (2004), 24-29.
- CAN, A.B.. Trust and Anonymity in Peer-to-peer Systems. Ph.D. dissertation, Dept. of Computer Science, Purdue University, West Lafayette, IN, 2007.
- FONSECA, E., FESTAG, A., BALDESSARI, R., AND AGUIAR, R.. Support of Anonymity in VANETs - Putting Pseudonymity into Practice. In *IEEE Wireless Communications and Networking Conference, 2007*.
- FURGEL, I., AND LEMKE, K.. A review of the digital tachograph system. In *The Workshop on Embedded Security in Cars (ESCAR'04), Bochum, Germany, 2004*.
- GOLLE, P., GREENE, D., AND STADDON, J.. Detecting and correcting malicious data in VANETs. In *The 1st ACM Workshop on Vehicular Ad Hoc Networks, Philadelphia, PA, 2004*.
- HUBAUX, J., CAPKUN, S., AND LUO, J.. The Security and Privacy of Smart Vehicles. *IEEE Security and Privacy*. 2, 3 (2004), 49-55.
- LILIE, L., AND BHARGAVA, B.. A Scheme for Privacy-preserving Data Dissemination. *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*. 36, 3 (May 2006), 503-506.
- MEKA, H., MADRIA, S., KUMAR, M., LINDERMAN, M., AND CHAKRAVARTHY, S.. Efficient Simulation Architecture for Routing and Replication in Mobile Peer to Peer Network of UAVs. In *The 11th International Conference on Mobile Data Management, Kansas City, MO, 2010*.
- OUI, B., LIAU, C., AND TAN, K.. Managing trust in peer-to-peer systems using reputation-based techniques. In *The 4th International Conference on Web Age Information Management, Chengdu, China, 2003*.
- PAPAGEORGIOU, C.P., OREN, M., AND POGGIO, T.. A general framework for object detection. In *The International Conference on Computer Vision, 1998*.
- PARNO, B., AND PERRIG, A.. Challenges in Securing Vehicular Networks. In *The 4th Workshop on Hot Topics in Networks, College Park, MD, 2005*.
- PHAM, C., AND TRAN, D.A.. Publish/Subscribe Techniques For P2P Networks. In *Al-Sakib Khan Pathan, Mukaddim Pathan and Hae Young Lee, ed., Advancements in Distributed Computing and Internet Technologies: Trends and Issues. IGI Global, Hershey, PA, 2012*.
- RAYA, M., AND HUBAUX, J.. Securing vehicular ad hoc networks. *Journal of Computer Security*. 15, (2007), 39-68.
- RAYA, M., PAPADIMITRATOS, P., AND HUBAUX, J.. Securing Vehicular Communications. *IEEE Wireless Communications Magazine, Special Issue on Inter-Vehicular Communications*. 13, 5 (2007), 8-15.
- REINHARD, E., WARD, G., PATTANAİK, S., AND DEBEVEC, P.. High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting. Morgan Kaufmann, Westport, Connecticut, 2005.
- SELÇUK, A.A., UZUN, E., AND PARIENTE, M.R.. "A reputation-based trust management system for p2p networks. In *The 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID), Chicago, IL, 2004*.
- YANG, X., LIU, J., ZHAO, F., AND VAIDYA, N.. A vehicle-to-vehicle communication protocol for cooperative collision warning. In *The First Annual International Conference on Mobile and Ubiquitous Systems, Boston, MA, 2004*.
- ZARKI, M.E., MEHROTRA, S., TSUDIK, G., AND VENKATASUBRAMANIAN, N.. Security issues in a future vehicular network. In *The European Wireless Conference, 2002*.

Bharat Bhargava is a professor of the Department of Computer Science with a courtesy appointment in the School of Electrical and Computer Engineering at Purdue University. Professor Bhargava is conducting research in security and privacy issues in distributed systems. This involves host authentication and key management, secure routing and dealing with malicious hosts, adaptability to attacks, and experimental studies. Related research is in formalizing evidence, trust, and fraud. Applications in e-commerce and transportation security are being tested in a prototype system. Based on his research in reliability, he is studying vulnerabilities in systems to assess threats to large organizations. He has developed techniques to avoid threats that can lead to operational failures. The research has direct impact on nuclear waste transport, bio-security, disaster management, and homeland security. These ideas and scientific principles are being applied to the building of peer-to-peer systems, cellular assisted mobile ad hoc networks, and to the monitoring of QoS-enabled network domains.

In the 1988 IEEE Data Engineering Conference, he and John Riedl received the best paper award for their work on "A Model for Adaptable Systems for Transaction Processing." Professor Bhargava is a Fellow of the Institute of Electrical and Electronics Engineers and of the Institute of Electronics and Telecommunication Engineers. In 1999, he received the IEEE Technical Achievement Award for a major impact of his decade long contributions to foundations of adaptability in communication and distributed systems. He has been awarded the charter Gold Core Member distinction by the IEEE Computer Society for his distinguished service. He received Outstanding Instructor Awards from the Purdue chapter of the ACM in 1996 and 1998. He has graduated the largest number of Ph.D students in CS department and is active in supporting/mentoring minority students. In 2003, he was inducted in the Purdue's Book of Great Teachers. He serves on seven editorial boards of international journals. He also serves the IEEE Computer Society on Technical Achievement award and Fellow committees. Professor Bhargava is the founder of the IEEE Symposium on Reliable and Distributed Systems, IEEE conference on Digital Library, and the ACM Conference on Information and Knowledge Management.



Pelin Angin is a Ph.D. Student at the Department of Computer Science at Purdue University. She received her BS degree in Computer Engineering at Bilkent University, Turkey in 2007. Her research interests lie in the fields of Mobile-Cloud Computing, Cloud Computing Privacy and Data Mining. She is currently working under the supervision of Professor Bharat Bhargava on leveraging mobile-cloud computing for real-time context-awareness and development of algorithms to address the associated mobile-cloud computing privacy issues.



Rohit Ranchal is a Ph.D student in Computer Science at Purdue University, West Lafayette, IN. He received his MS in Computer Science from Purdue University in 2011 and B. Tech in Information Technology from DAV Institute of Engineering and Technology, Jalandhar, India in 2009. His research interests include security metrics, secure information dissemination, security, privacy, trust in Cloud Computing, and Service Oriented Architecture (SOA).



Ranjitkumar Sivakumar is a Master's Student at Department of Computer Science at Purdue University. He has completed his Bachelors of Engineering in Computer Science and Engineering with Distinction at Anna University, India. His research interests lie in Databases, Cloud Computing and Information Security and he is working under the supervision of Professor Bharat Bhargava. He is going to work with Amazon.com upon completion of his Master's degree.



Asher Sinclair is a senior Program Manager at AFRL's Information Directorate working in the Enterprise Information Management Branch at the Rome Research Site. His work history includes enterprise systems management, service-oriented architectures, information-centric quality of service, and application and network-level security. He has contributed to more than 16 technical papers and conference proceeding publications. He holds a bachelor's degree in Computer Information Systems from the State University of New York and a master's degree in Information Management from Syracuse University.



Mark Linderman is a Principal Researcher in Information Management Technologies in the Air Force Research Laboratory Information Directorate where he leads the Information Management technology area. From 2000-2004, he was the Technical Lead of the AFRL Information Directorate Joint Battlespace Infosphere (JBI) Program. The JBI program addresses a broad spectrum of information management challenges from access control, the role of metadata, dissemination, persistence and destruction of information. Dr. Linderman is active in the Technical Cooperation Program (TTCP) that fosters collaboration of research, development and experimentation among the United States, the United Kingdom, Australia, and Canada. He chaired of the C3I Group Technical Panel on Information Management from its inception until 2009. Prior to joining the JBI project, Dr. Linderman was the Technical Advisor for the Advanced Computing Architectures Branch, and he was active in several projects involving the application of HPCs to signal processing challenges. Dr. Linderman made several contributions to the field of embedded High Performance Computing (HPC), including the demonstration of a highly efficient Space-Time Adaptive Processing implementation, embedded signal processor simulation, design and testing, and high performance computing. Dr. Linderman joined AFRL in 1994 after completing his M.Eng. and Ph.D. in electrical engineering from Cornell University. He holds a BSEE degree from the University of Delaware.