On Data-Centric Trust Establishment in Ephemeral Ad Hoc Networks

Maxim Raya*, Panagiotis Papadimitratos*, Virgil D. Gligor†, Jean-Pierre Hubaux*
*School of Computer and Communication Sciences
EPFL, Switzerland
Email: {maxim.raya,panagiotis.papadimitratos,jean-pierre.hubaux}@epfl.ch
†Department of Electrical and Computer Engineering
Carnegie Mellon University, USA
Email: gligor@cmu.edu

***Abstract*—We argue that the traditional notion of trust as a relation among entities, while useful, becomes insufficient for emerging data-centric mobile ad hoc networks. In these systems, setting the data trust level equal to the trust level of the data-providing entity would ignore system salient features, rendering applications ineffective and systems inflexible. This would be even more so if their operation is ephemeral, i.e., characterized by short-lived associations in volatile environments. In this paper, we address this challenge by extending the traditional notion of trust to data-centric trust: trustworthiness attributed to node-reported data per se. We propose a framework for data-centric trust establishment: First, trust in each individual piece of data is computed; then multiple, related but possibly contradictory, data are combined; finally, their validity is inferred by a decision component based on one of several evidence evaluation techniques. We consider and evaluate an instantiation of our framework in vehicular networks as a case study. Our simulation results show that our scheme is highly resilient to attackers and converges stably to the correct decision.**

I. Introduction

In all traditional notions of trust, data trust (e.g., trust in the identity or access/attribute certificates) was based exclusively on a priori trust relations established with the network entities producing these data (e.g., certification authorities, network nodes) [9], [16], [17]. This was also the case when trust was derived via fairly lengthy interactions among nodes, as in reputation systems [4], [8], [18], [27]. Moreover, any new data trust relationships that needed to be established required only trust in the entity that produced those data. All trust establishment logics proposed to date have been based on entities (e.g., "principals" such as nodes) making statements on data [4], [7], [9], [12], [16], [17], [24], [25]. Furthermore, traditional trust relations evolved generally slowly with time: once established, they lasted a long time and changed only after fairly lengthy operations (e.g., certificate revocation or monitoring and then voting-off of peers).

These observations indicate that existing trust notions are *entity-centric* and *slow to change*. However, several emerging mobile networking systems are heavily, if not entirely, *data-centric* in their functionality and operate in *ephemeral* environments. In such scenarios, it is more useful to establish trust in data rather than in the nodes reporting them. For example, in vehicular networks, node identities are largely irrelevant; rather, safety warnings and traffic information updates, along with their time freshness and location relevance, are valuable. At the same time, interactions with data reporters do not rely on any prior association, and encounters are often short-lived, especially due to high mobility.

In such scenarios, and unlike in traditional trust establishment schemes, the trust level associated with data is not the same as that of the node that generated the data. More specifically, in the vehicular networks example, vehicles will have different preset node trust levels (e.g., police cars are more trustworthy than private vehicles), but (i) different events reported by the same vehicle may have different levels of trust (due to distance to the event, timeliness of the report, vehicle equipment level) that may differ from that of the vehicle itself; (ii) the same event reported by multiple vehicles with different preset node trust levels has to be associated with a single trust level that would, of course, differ from some of the levels of the reporting vehicles; and (iii) an event reported by a vehicle requires corroboration by other vehicles and hence its level of trust would differ from that of the reporting vehicle.

In other words, the following question arises naturally: how can these emerging systems be effective and trustworthy when their basic operational requirements are not satisfied by existing trust notions? To address this challenge, we advocate a clean-slate approach. We propose data-centric trust establishment: *data trustworthiness should be attributed primarily to data per se, rather than being merely a reflection of the trust attributed to data-reporting entities*.

The logic we propose extends the traditional notions of trust and methods of trust establishment in several ways.

This work is partially funded by the EU project SEVECOM (http://www.sevecom.org).

Virgil Gligor's research was supported in part by the US Army Research Laboratory and the UK Ministry of Defence Agreement Number W911NF-06-3-0001 and by the US Army Research Office under Contract W911NF-07-1-0287 at the University of Maryland. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, US Army Research Office, the U.S. Government, the UK Ministry of Defence, or the UK Government.

First, unlike traditional trust, a priori trust relationships in entities (nodes) represent only one of the default parameters for establishing data trust. For example, our logic, while using nodes' statements on data, does not rely exclusively on such statements. Instead, it takes into account dynamic factors, such as location and time, as well as the number and type of the statements on data, to derive data trust relations. Second, beyond the traditional time-invariant or slow-evolving trust notions, data-centric trust relations are by definition ephemeral and have to be established and re-established frequently, based on network and perceived environment changes. For example, an event report (e.g., accident report, weather report) that must be believed by recipient nodes in real-time cannot last longer than the lifetime of the event or of the network that generated this event. Multiple rounds of node interactions are typically not possible in such networks. Third, trust does not stem from a single source of data (e.g., a certification authority) and generally it is application-dependent (in contrast to entity-centric trust when, for example, multiple applications use certificates for their access control and authentication policies).

We derive trust in data (e.g., reported event) from multiple pieces of evidence (e.g., reports from multiple vehicles). Then, our logic weighs each individual piece of evidence according to well-established rules and takes into account various trust metrics, such as time freshness and location relevance, defined specifically in the context of an application. Then, data and their respective weights serve as inputs to a decision logic that outputs the level of trust in these data. We evaluate several techniques, including voting, Bayesian inference, and the Dempster-Shafer Theory of evidence. Notably, Bayesian inference takes into account prior knowledge, whereas the Dempster-Shafer Theory accounts for the *uncertainty* about data. More specifically, while trust establishment mechanisms based on popular decision logics, such as voting and Bayesian inference, consider uncertainty as refutal of evidence, our framework considers uncertainty as either supporting or refuting evidence, thus making the decision process more realistic. We show in this work that this distinction affects the flexibility and resilience to attackers in some scenarios.

In the rest of this paper, we present our framework in Sec. II. In Sec. III we mathematically develop our approach. Then, we instantiate our framework in the context of vehicular communication systems in Sec. IV. We evaluate the effectiveness of our scheme through simulations in Sec. V, and conclude with a survey of related work in Sec. VI.

II. General framework

A. Preliminaries

We consider systems with an authority responsible for assigning identities and credentials to all system entities that we denote as *nodes*. All legitimate nodes are equipped with credentials (e.g., certified public keys) that the authority can revoke. Specific to the system and applications, we define a set $\Omega = \{\alpha_1, \alpha_2, \ldots, \alpha_I\}$ of mutually exclusive *basic events*. Composite events γ are unions or intersections of multiple basic events. Examples of basic events are "ice on the road" and "traffic jam". If the ice on the road causes a traffic jam, this becomes the composite event "ice on the road and traffic jam ahead". Each α_i is a perceivable event generated by the environment, network, or an application running on vehicles. There may be multiple applications, each having its own set of relevant events. These sets are overlapping, as their events belong to the pool of basic events.

We consider V, the set of nodes v_k, classified according to a system-specific set of node types, $\Theta = \{\theta_1, \theta_2, \ldots, \theta_N\}$. We define a function $\tau : V \rightarrow \Theta$ returning the type of node v_k. *Reports* are statements by nodes on events, including related time and geographic coordinates where applicable. For simplicity, we consider reports on basic events, as reports on composite events are straightforward. We do not dwell on the exact method for report generation, as this is specific to the application.

B. Default Trustworthiness

We define the *default trustworthiness* of a node v_k of type θ_n as a real value that depends on the attributes related to the designated type of node v_k. For all node types, there exists a trustworthiness ranking $0 < \theta_1 < \theta_2 < \ldots < \theta_N < 1$. For example, some nodes are better protected from attacks, more closely monitored and frequently re-enforced, and, overall, more adequately equipped, e.g., with reliable components. As they are less likely to exhibit faulty behavior, they are considered more trustworthy.

We stress here that the data-centric trust establishment framework does not aim to replace or amend source authentication, as in reputation systems, but uses it as an input to the data trust evaluation function. In fact, if a node reputation system were in place, its output scores could also be used as input to the data trust function. Hence, data trust builds on the information provided by source authentication and reputation systems without trying to supplant them. The choice of the entity trust establishment system is orthogonal to the scope of this paper and has been prolifically addressed in the literature (Sec. VI).

C. Event- or Task-Specific Trustworthiness

Nodes in general perform multiple tasks that are system-, node- and protocol-specific actions. Let Λ be the set of all relevant system tasks. Then for some nodes v_1 and v_2 with types $\tau(v_1) = \theta_1$ and $\tau(v_2) = \theta_2$ and default trustworthiness rankings $t_{\theta_1} < t_{\theta_2}$, it is possible that v_1 is more trustworthy than v_2 with respect to a task $\lambda \in \Lambda$.

Reporting data on events is clearly one of the node tasks. For the sake of simplicity, we talk here about event-specific trustworthiness implying that it is actually task-specific trustworthiness. Nevertheless, the two can be easily distinguished, when necessary; e.g., when tasks include any other protocol-specific action such as communication.

With the above considerations in mind, we define the event-specific *trustworthiness* function $f : \Theta \times \Lambda \rightarrow [0, 1]$. f has two arguments: the type $\tau(v_k)$ of the reporting node v_k and the task λ_j. f does differentiate among any two or more nodes

of the same type, and if $\lambda_j = \emptyset$ (no specific event or task), f is the default trustworthiness $f = t_{\tau(v_k)}$.

D. Dynamic Trustworthiness Factors

The ability to dynamically update trustworthiness can be valuable, especially for capturing the intricacies of a mobile ad hoc networking environment. For example, nodes can become faulty or compromised by attackers and hence need to be revoked. In addition, the location and time of report generation change fast and are important in assigning trustworthiness values to events.

To capture this, we define a *security status* function $s : V \rightarrow [0, 1]$. $s(v_k) = 0$ implies node v_k is revoked, and $s(v_k) = 1$ implies that the node is legitimate. Intermediate values can be used by the system designer to denote different trustworthiness levels, if applicable.

Second, we define a set of *dynamic trust metric* functions $M = \{\mu_l : V \times \Lambda \rightarrow [0, 1]\}$ indexed by a selector l indicating different node attributes (e.g., location) that dynamically change. That is, for each attribute, a different metric μ_l is defined. μ_l takes node $v_k \in V$ and task $\lambda_j \in \Lambda$ as inputs and returns a real value in $[0, 1]$.

E. Location and Time

Among the possible values of l for metric μ_l, *proximity* either in *time* or *geographic location* is an attribute of particular importance. Proximity can increase the trustworthiness of a report: The closer the reporter is to the location of an event, the more likely it is to have accurate information on the event. Similarly, the more recent and the closer to the event occurrence time a report is generated, the more likely it is to reflect the system state.

Cryptographic primitives, such as digital signatures, can ensure that location and time information cannot be modified if included in a report. However, the accuracy of such information can vary, due to nodes' differing capabilities or (malicious or benign) faults. This is especially true for reports that depend on fine-grained time and location data. Hence, different types of nodes are more or less trustworthy when reporting such data. In some cases, time- or geo-stamping a report can be a distinct task.

F. Scheme Overview

We compute the trustworthiness of a report e_k^j, generated by node v_k and providing supporting evidence for event α_j, by using both (i) static or slow-evolving information on trustworthiness, captured by the default values and the event-specific trust f, and (ii) dynamically changing information captured by security status s and more so by metric μ_l. We combine these as arguments to a function

$$F(e_k^j) = G(s(v_k), f(\tau(v_k), \lambda_j), \mu_l(v_k, \lambda_j))$$

that returns values in the $[0, 1]$ interval. If v_k reports no evidence for α_j, $F(e_k^j) = 0$. These values are calculated locally for each report received from another node and are called the *weights* (or *trust levels*) of the reports. Fig. 1 illustrates our scheme.

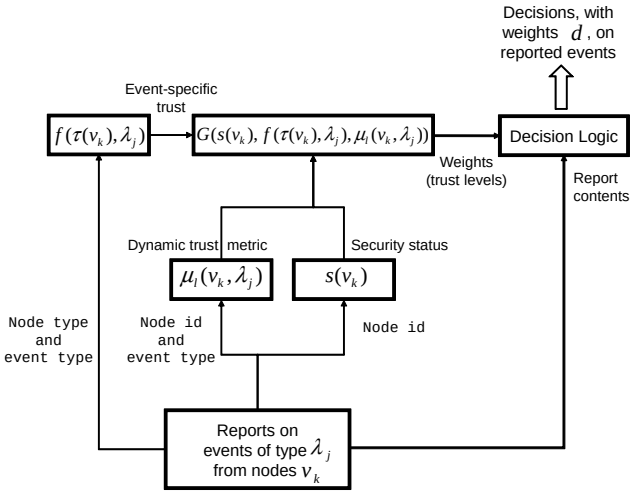

Fig. 1. Data-centric trust establishment framework.

Nonetheless, such a per message assessment may often be insufficient. It can be hard to decide whether the reported event took place based on a single message, and it is vulnerable to faults (e.g., equipment failures or compromised nodes). Instead, we propose the collection of multiple reports related to the same event and of their weights, i.e., the accompanying F values, and their combination into a robust decision scheme. Thus, the reports along with their weights are passed to a *Decision Logic* module that outputs an assessment on the event in question. The way to use such decisions and inferences is beyond the scope of this paper, as it is specific to particular systems.

The above process is tightly related to the multisensor data fusion techniques [11]. In fact, F can be computed using rule-based expert systems; the output of the *Decision Logic* module can be used by another expert system that makes decisions based on reported events. There are several algorithms to implement the *Decision Logic* module; we will compare next a selected subset of these algorithms in the context of data-centric trust establishment. It should be stressed here that data-centric trust establishment in wireless networks is a new application of data fusion techniques, to the best of our knowledge.

III. Evidence evaluation

The literature on trust in ad hoc networks proposes several approaches for trust establishment, which we survey in Sec. VI. In this work, we propose a new technique and compare it to four other existing techniques. These techniques are described below.

To mathematically model our approach, assume a node A has to decide among several basic events $\alpha_i \in \Omega$, based on K pieces of evidence e_k^j (reports from K distinct nodes).

Let d_i denote the combined trust level computed by evaluating evidence corresponding to event α_i. The *Decision Logic*

module outputs the event that has the highest combined trust level, i.e., $\max_i(d_i)$.

A. Basic techniques

The following two techniques are used for reference and serve as a basis of comparison for the remaining three techniques.

1) Majority Voting: In this technique, the majority wins (e.g., [19]). The combined trust level corresponding to event α_i is defined by:

$$d_i = \frac{1}{K}\sum_{k=1}^{K} F(e_k^i) \tag{1}$$

where $F(e_k^i) = 1$ if v_k reports α_i and it is 0 otherwise.

2) Most Trusted Report: The Most Trusted Report (MTR) decision logic outputs a trust level equal to the maximum value of trust levels assigned to reports about the event; the point of using MTR is to show the effect of isolated high trust values (in data or entities) on the system. The combined trust level corresponding to event α_i is defined by:

$$d_i = \max_k(F(e_k^i)) \tag{2}$$

B. Weighted Voting

As its names implies, Weighted Voting (WV) sums up all the votes supporting an event with each vote weighted by the corresponding trust level to output the combined trust level:

$$d_i = \frac{1}{K}\sum_{k=1}^{K} F(e_k^i) \tag{3}$$

It should be noted here that decisions on composite events are harder to do using the above three techniques since they do not provide formalisms for handling unions and intersections of events. In contrast, the next two techniques provide such formalisms.

C. Bayesian Inference

Among the data fusion techniques, Bayesian Inference (BI) [20] is the one most frequently used for trust establishment. In BI, the combined trust level corresponding to α_i is the posterior probability of α_i given new evidence $e = \{e_1^j, e_2^j, \ldots, e_K^j\}$; it is expressed in terms of the prior probability $P[\alpha_i]$ using the Bayes' theorem:

$$P[\alpha_i|e] = \frac{P[\alpha_i]\prod_{k=1}^{K} P[e_k^j|\alpha_i]}{\sum_{h=1}^{I}(P[\alpha_h]\prod_{k=1}^{K} P[e_k^j|\alpha_h])} \tag{4}$$

where we assume that reports are independent for the sake of mathematical tractability (the receiver cannot sort out the dependencies among reports from distinct vehicles since such information is not provided in the reports).

The computation of posterior probabilities for composite events γ (recall that they are unions or intersections of basic events) follow the rules of probability theory.

$P[e_k^i|\alpha_i]$ is the probability that report k confirms event α_i, given that α_i happened. Using trust levels as weights of reports, this probability is equal to the *trust level*: $P[e_k^i|\alpha_i] = F(e_k^i)$.

For $j \neq i$, $P[e_k^j|\alpha_i]$ is the probability that report k does not confirm α_i (hence, it confirms $\bar{\alpha_i}$, the complement of α_i in Ω), given that α_i happened. This is equivalent to a malfunctioning or cheating node (ideally, a node would report a real event). Hence, $P[e_k^j|\alpha_i] = 1 - P[e_k^i|\alpha_i] = 1 - F(e_k^i)$.

D. Dempster-Shafer Theory

In Dempster-Shafer Theory (DST) [22], evidence evaluation is inspired by human reasoning. More specifically, the lack of knowledge about an event is not necessarily a refutal of the event. In addition, if there are two conflicting events, uncertainty about one of them can be considered as supporting evidence for the other. The major difference between BI and DST is that the latter is more suitable for cases with uncertain or no information. More precisely, in DST a node can be uncertain about an event, unlike in BI where a node either confirms or refutes the event. For example, if a node A confirms the presence of an event with probability p, in BI it refutes the existence of the event with probability $1 - p$. In DST, probability is replaced by an uncertainty interval bounded by $belief$ and $plausibility$. Belief is the lower bound of this interval and represents supporting evidence. Plausibility is the upper bound of the interval and represents non-refuting evidence. Hence, in this example, node A has p degree of belief in the event and 0 degree of belief in its absence.

In DST, the *frame of discernment* contains all mutually exclusive possibilities related to an observation. Hence, in our context, it is the set Ω defined previously. The belief value corresponding to an event α_i and provided by report k is computed as:

$$bel_k(\alpha_i) = \sum_{q:\alpha_q \subset \alpha_i} m_k(\alpha_q)$$

which means it is the sum of all basic belief assignments $m_k(\alpha_q)$, α_q being all basic events that compose the event α_i. In this case, only $\alpha_i \subset \alpha_i$ and hence $bel_k(\alpha_i) = m_k(\alpha_i)$.

The plausibility value corresponding to event α_i represents the sum of all evidence that does not refute α_i and is computed as:

$$pls_k(\alpha_i) = \sum_{r:\alpha_r \cap \alpha_i \neq \emptyset} m_k(\alpha_r)$$

Belief and plausibility are related by $pls(\alpha_i) = 1 - bel(\bar{\alpha_i})$.

The combined trust level corresponding to event α_i is the belief corresponding to α_i:

$$d_i = bel(\alpha_i) = m(\alpha_i) = \bigoplus_{k=1}^{K} m_k(\alpha_i) \tag{5}$$

where pieces of evidence can be combined using Dempster's rule for combination:

$$m_1(\alpha_i) \bigoplus m_2(\alpha_i) = \frac{\sum_{q,r:\alpha_q \cap \alpha_r = \alpha_i} m_1(\alpha_q) m_2(\alpha_r)}{1 - \sum_{q,r:\alpha_q \cap \alpha_r = \emptyset} m_1(\alpha_q) m_2(\alpha_r)}$$

As before, using trust levels as weights of reports, the basic belief assignment that confirms α_i is equal to the *trust level*: $m_k(\alpha_i) = F(e_k^i)$.

For composite events γ, belief can be computed similarly using the above equations.

IV. Case Study

To illustrate the application and utility of the data trust framework, we present in the following a case study of a real ephemeral ad hoc network instantiation, namely vehicular networks. We first describe the system and adversary models, then explain through examples how the different components of data trust can be practically derived.

A. Secure Vehicular Communications System

Vehicular Ad hoc NETworks (VANET) and *Vehicular Communication (VC)* systems [26] are being developed to enhance the safety and efficiency of transportation systems, providing, for example, warnings on environmental hazards (e.g., ice on the pavement) and traffic and road conditions (e.g., emergency braking, congestion, or construction sites). From a networking point of view, the nodes are vehicles and road-side infrastructure units (RSUs), all equipped with on-board processing and wireless modules, thus enabling multi-hop communication in general.

Authorities are public agencies or corporations with administrative powers, e.g., city or state transportation authorities entrusted with the management of node *identities* and *credentials*. A subset of the infrastructure nodes serves as a gateway to and from the authorities.

We assume that each node v_k is equipped with a pair of private/public cryptographic keys Pr_k/Pu_k, and a certificate issued by an authority X as $Cert_X\{Pu_k\}$. Nodes are equipped with a clock and a positioning system (such as GPS or Galileo). This allows them to include their time and location information in any outgoing reports. Source authentication, required to prevent Sybil attacks, is achieved by digital signatures according to both industrial and academic proposals [2], [21]. In this example, source authentication identifies the type of the report sender and enables the assignment of default trustworthiness, as explained in Sec. II-B.

Unicast and multicast communication is possible; however, local broadcast (single hop) and geocast (flooding to a given geographic area) are predominantly used. Vehicle-specific information (e.g., velocity, coordinates) is transmitted frequently and periodically in the form of *safety messages*.[1] Reports on in-vehicle or network events are included in these messages. Safety and other messages, generated by vehicles and RSUs, can result in an abundant influx of information about events. It is important to note here that our approach, based exclusively on local processing, does not add any communication overhead and very little computation overhead to a secure VC system where the actual overhead is due to frequent broadcasting and asymmetric cryptography and is inherent in VANETs.[2]

[1]Typically, for a highway, every 300ms over a nominal range of 300m.

B. Adversary Model

Nodes either comply with the implemented protocols (i.e., they are correct) or they deviate from the protocol definition intentionally (attackers) or unintentionally (faulty nodes). Both attackers and faulty nodes can cause damage to the network and hence we consider them both as adversaries. The attacks that can be mounted by either internal (equipped with credentials and cryptographic keys) or external adversaries vary greatly. In brief, adversaries can replay any message, jam communications, and modify (yet in a detectable manner due to the digital signatures) messages. More importantly, they can inject faulty data and reports, or control the inputs to otherwise benign nodes and induce them to generate faulty reports.

We assume that at most a small fraction of the nodes are adversaries, and consequently the fraction of the network area affected by them is bounded. This bound on the presence of adversaries could be further refined by distinct values for different node types. But this assumption does not preclude that a few adversarial nodes surround a correct node at some point in time.

C. Framework Instantiation

We focus on the use of our scheme on board a vehicle. Clearly, it could be run on RSUs; nonetheless, the challenge is to design a scheme practical for nodes that are not part of the system infrastructure.

The forms of the f (event-specific trust), s (security status), μ_l (dynamic trust metric), and G (trust level) functions are determined by the secure VC system: They are either preloaded at the time the node is bootstrapped, or updated after the node joined the system. Their values are either provided by the authorities or distributed by the infrastructure.

To illustrate our instantiation, we consider an example scenario: a highway accident in which vehicle B is involved. Now, let us consider a vehicle A, several communication hops away from the accident location. A receives safety messages indicating that there is an accident on its route and has to decide whether to trust this information. In this case, we assume the event α_1: "There is an accident at location L_B". The granularity of the event location should be properly defined to avoid having reports on several different events while, actually, all these reports refer to the same event but with slightly different locations. Now assume that one or more attackers generate safety messages supporting the null event $\alpha_2 = \emptyset$: "There is no accident at location L_B". If there are several events (e.g., several distinct locations, given the defined

[2]It should be clarified that, although this overhead seems unreasonable for typical ad hoc networks, VANETs have distinct properties and requirements (making networking and security infrastructure necessary as in cellular systems) [14] and were shown to be able to support public key cryptography [21].

granularity), the data trust is computed for each of them. The resulting values can be used by the application to decide the consequent action; the specific use of these values by the application is beyond the scope of this paper.

Two important system parameters are the set of reports and the time needed to make a correct judgment. The reports considered valid for making decisions should be sent by vehicles that are on the communication path between the accident location and A. The time to correct judgment should be equal to the time needed by the *Decision Logic* module to converge to a stable output value; this time depends on the frequency of message reception. It is also constrained by the tolerable decision delay (e.g., in critical situations, decisions should be made very fast, given the available data) that depends on the event in question. The convergence in a typical VANET scenario is illustrated in Sec. V-D.

V. Performance evaluation

In this section, we examine the performance of the decision logics described in Sec. III. Recall that a vehicle computes the combined trust in an event based on the reports it receives from distinct vehicles. We compare the four decision logics: MTR, WV, BI, and DST against the basic majority voting scheme. We use the example scenario with two events α_1 and α_2 described in Sec. IV-C.

As noted earlier, the exact choice of the actual values of the f, s, μ_l, and G functions depends on the system designer and hence is out of scope of this paper. In order to provide an analysis that is independent of administrative decisions, we studied the effects of several general but representative parameters, namely the *percentage of false reports*, *prior knowledge*, *uncertainty*, and *evolution in time*; this allows us to draw conclusions that are independent from the specific choice of default trustworthiness values. We study the effect of these parameters on the *probability of attack success*, which is very important in a security context. In the case of basic majority voting, this probability is equal to 1 if the percentage of attackers is larger than 50%; basic majority voting is represented in the following figures by a vertical or horizontal dashed line corresponding to a percentage of false reports equal to 50 or a probability of attack success equal to 0.5, respectively.

We use a Beta distribution, with its mean equal to an average trust level (defined for each scenario), to assign the trust levels to the reports received by a vehicle A. We chose the Beta distribution because it approximates the Normal distribution, a common choice in statistics, but with bounds (0 and 1). We simulate scenarios with 10 or 50 valid reports (i.e., sent by vehicles on the communication path between an accident location and A, as described in Sec.IV-C).[3] This means that A includes all these reports in its decision process; each report confirms either event α_1 or α_2. Table I lists the parameters used in the following simulation scenarios.

[3]We do not simulate the wireless medium in this case since it is orthogonal to our evaluation. Sec. V-D simulates a VANET, including the wireless communication.

Scenario number	Parameter		
	$E[F(false)]$	$E[F(correct)]$	N
1	0.6	0.8	50
2	0.8	0.6	50
3	0.6	0.8	10
4	0.2	0.4	50
5,6	0.6	0.8	17

TABLE I
Simulation scenario parameters. $E[F(false)]$ denotes the average trust level of false reports and $E[F(correct)]$ is the average trust level of correct reports. N is the number of reports. In scenarios 5 and 6, the value of N is determined by the ns-2 simulations as further explained in Sec. V-D.

Simulations were performed in MATLAB (Sec. V-A to V-C) and ns-2 (Sec. V-D) with results averaged over 100 randomly seeded runs and plotted with 95% confidence intervals.

The results show that: First, trust decisions based on MTR are the most sensitive to different parameters since the MTR is not corroborated by other vehicles in this case. Second, under realistic conditions, the other three decision logics outperform both majority voting and MTR. Third, there is no clear winner among these decision logics as each performs best in certain scenarios. The details follow.

A. Effect of Data Trust

To see the effect of data trust on the resilience of the decision logic, we compare the different decision logics to majority voting. The graphs in Figs. 2(a) and 2(b) provide insight into the effect of the percentage of false reports in order to disturb the perception of the observing vehicles (Sec. IV-B). There are two different pieces of information, the false one (originating from colluding attackers or malfunctioning honest vehicles) and the correct one from functional honest vehicles, that are conflicting in their content. Collusion in this case means that all attackers report the same false information. In addition, the trust distributions of the reports generated by honest nodes and by attackers follow Beta distributions with different means. We examine two scenarios: in Scenario 1 (Fig. 2(a)), the average trust of false reports is lower than that of correct reports; Scenario 2 (Fig. 2(b)) illustrates the opposite situation (e.g., because attackers are positioned closer to the event). The mean values corresponding to both scenarios are listed in Table I.

In Figs. 2(a) and 2(b), we observe that MTR is both little resilient to small percentages of attackers and highly resilient (on average) to high percentages of attackers. This can be explained by the fact that MTR relies on the trust value of only one report, which can differ significantly from the average trust value. The other three decision logics are more resilient to attacks than majority voting when correct reports are more trustworthy than false ones (this is a realistic situation). BI is the most resilient of all three methods. When false reports are more trustworthy than correct ones, the situation is reversed and weighted voting becomes the most resilient technique. There are two curves for BI, each corresponding to a different prior probability; these plots are discussed next.

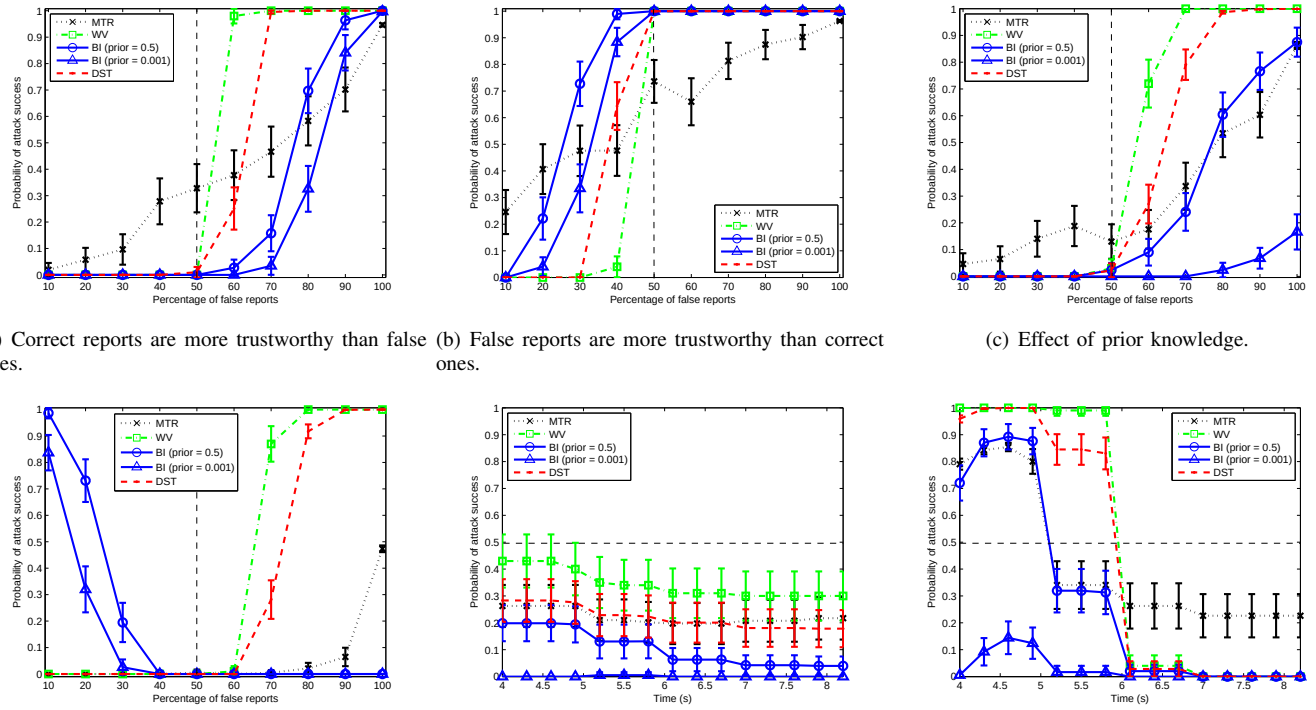

(a) Correct reports are more trustworthy than false ones.

(b) False reports are more trustworthy than correct ones.

(c) Effect of prior knowledge.

(d) Effect of uncertainty.

(e) Evolution in time; evenly distributed false reports.

(f) Evolution in time; false reports received first.

Fig. 2. Performance of the decision logics with respect to the percentage of false reports, prior knowledge, uncertainty, and time.

B. *Effect of Prior Knowledge*

One of the properties of BI is that it uses a prior probability to compute the posterior probability of an event (Sec. III-C). The prior probability represents the amount of knowledge about the event prior to the reception of new evidence; in our example, this is the probability of the presence of an accident. In this section, we study the effect of prior probabilities on the performance of BI. In VANETs, prior probability can be derived from ITS (Intelligent Transportation Systems) studies developed to estimate the probabilities of crash occurrence in different scenarios (e.g., [3]). Since these estimates depend on many parameters, notably the human factors, and cannot be computed in a generic case, we chose a rather conservative prior probability of 0.001 for our simulations. We also simulate the effect of the neutral prior probability 0.5 (i.e., lack of prior knowledge).

Figs. 2(a) and 2(b) show to some extent that the availability of prior knowledge increases the resilience of BI to false data attacks. In Fig. 2(c), there are fewer reports (only 10 compared to 50 in the previous two scenarios) and we can clearly see the benefit of prior knowledge. The reason for this increase in resilience when the number of reports decreases is that large numbers of reports damp the effect of prior probability in the calculation of the posterior probability (Sec. III-C).

C. *Effect of Uncertainty*

In a nutshell, BI does not take uncertainty into account whereas DST does (Sec. III-D). To simulate the effect of uncertainty (Fig. 2(d)) on the decision logics, we use low mean data trust levels for both false and correct reports; the exact values are listed in Table I (such values can result from low values of the security status s, e.g., due to the discovery of a virus in the network). In this case, DST is indeed the most resilient of the decision logics.

An interesting observation is related to the behavior of BI. At high trust levels (Fig. 2(a)), it exhibits behavior similar to that of WV and DST. But at low trust levels (Fig. 2(d)), it behaves opposite to the other two methods. This is so because BI deals with only the probabilities of true and false hypotheses and if a report is assigned a 0.2 trust level (i.e., 0.2 probability of being correct), it is assumed to have a 0.8 mistrust level (i.e., 0.8 probability of being false). In fact, BI requires that hypotheses be mutually exclusive and hence does not support general uncertainty that may overlap with either hypothesis. Thus, given a small percentage of supporting reports with low trust levels, there is a high percentage of refuting reports with low trust levels also. In BI, this high percentage is transformed into a high percentage of supporting reports with high trust levels (i.e., the opposite). Similar reasoning applies to high percentages of supporting reports.

It is important to note here that we applied BI in the typical way widely used in both research and industry (e.g., [15]). There are efforts to transform belief function models to probability models, thus enabling BI to handle uncertainty [6]. The use of such methods may reduce the adverse effects of uncertainty on the performance of BI shown in this section.

D. Evolution in Time

In ephemeral networks, it is important to evaluate data trust rapidly in order to permit an application logic to use the resulting values. Hence, a decision logic should be able to output the final result as fast as possible, based on the freshly received reports. This property distinguishes the mechanisms explored in this work from other approaches that rely on a longer history of available reports (e.g., reputation systems [4], [18], [27]). In this section, we are only interested in the decision delay as reports arrive. The total event detection delay by the observing vehicle depends also on how fast the reporting vehicles detect the event, which in turn depends on the particular detection sensors and hence we do not consider it in this work.

To simulate ephemeral networks, we use VANETs with mobile vehicles. Our scenario is a 2 km-long highway with 3 lanes in each direction. There are 300 vehicles moving at speeds between 90 km/h and 150 km/h; the average distance between two vehicles on the same lane is 40 m. Vehicles periodically broadcast safety messages every 300 ms within a radius of 300 m (single hop), according to the DSRC specification [1]; the broadcast start times are uniformly distributed between 0 and 2 seconds, approximately. In our simulations, we study the reception of reports at a vehicle A positioned in the middle of the scenario on the 90 km/h lane. We assume that an event (e.g., "ice on the road") is generated by honest vehicles between coordinates 1300 m and 1400 m (the icy section). The attackers report the opposite event ("no ice on the road"). As A moves towards the icy section, it receives reports from vehicles that pass inside this section. Only the last report from each vehicle is considered; this allows A to update its decision as vehicles enter the icy section and change their reports.

The parameters for this scenario are listed in Table I. We observe that each vehicle receives on average (over 100 runs) 17 reports sent by distinct vehicles from inside the icy section. The received reports are assigned corresponding trust levels in MATLAB. In Fig. 2(e), the percentage of false reports received in each timestep is drawn from a Binomial distribution with a probability 0.5 (i.e., the mean percentage of false reports is 50); this figure shows the stability of the decision logics when the percentage of false reports varies. In Fig. 2(f), the total percentage of false reports is also 50, but all false reports are received at the beginning of the simulation time to simulate the speed of convergence of the decision logics.

By examining Figs. 2(e) and 2(f), we can see that the speed of convergence of all four decision logics depends on the number of received reports and hence the scenario parameters (event generation time, vehicle density, etc.). We leave further investigation of these parameters to future work due to the lack of space. Nevertheless, we can observe the general trend of convergence to a stable output value before A reaches the icy section and despite variations in the percentage of false reports (of course, as long as the total percentage of false reports is constant). The output values roughly correspond to the results of Fig. 2(a), as expected.

E. Discussion

Based on the above results, we can see that there is no clear winner among the decision logics that fits best all scenarios. But we can elaborate several guidelines for the evaluation of data-centric trust:

- If the uncertainty in the network is low, BI is the most resilient to false reports. To avoid the case of few highly trustworthy false reports (Fig. 2(b)), the decision of BI should be positioned with respect to another logic, such as DST or WV, and the most conservative value (i.e., the one that yields the lowest probability of attack success) should be taken.
- The availability of prior knowledge can further improve the resilience of BI.
- If the uncertainty in the network is high, DST performs consistently better than other methods (MTR does not always yield better results).

VI. Related work

Work on trust has produced rich literature in conventional, P2P and ad hoc networks. In the latter, most contributions assume that there is no infrastructure and no PKI; trust is a relation among entities; trust is based on observations, with a history of interactions needed to establish trust. To the best of our knowledge, the computation of trust values in the context of ad hoc networks has been considered in only two cases: certification [7], [25] and routing [4], [27]. Otherwise, trust evaluation assumes the prior establishment of trust relations. In both certification and routing, trust values are established in very specific ways that cannot be generalized to other approaches.

Eschenauer et. al. [7] introduce the general principles of trust establishment in mobile ad hoc networks and compare them to those in the Internet. They describe examples of generic evidence generation and distribution in a node-centric authentication process.

Several papers [4], [8], [18], [27] describe the use of modified Bayesian approaches to build reputations systems with secondhand information to establish trust in routing protocols. As mentioned throughout the paper, reputation systems monitor node actions over several interactions to compute node trust values. In contrast, data trust, as defined in this work, focuses on evaluating data rather than nodes and based on only one message per node (to cope with the ephemerality of the network). In addition, all of these works relied on BI to compute reputation scores, whereas we showed that DST is more resilient to attacks when there is high uncertainty in the network (Sec. V-C). Data trust can actually complement reputation systems in non-ephemeral networks.

The main approach advanced by Jiang and Baras [12] is based on local voting that is a weighted sum of votes. Conflicting votes are mitigated by each other when summed. They also favor local interactions that we use as well.

The main idea behind the work by Sun et. al. [24] is that trust represents uncertainty that in turn can be computed using

entropy. They also introduce the notion of *confidence of belief* to differentiate between long-term and short-term trust. Trust can be established through direct observations or through a third party by recommendations.

Theodorakopoulos and Baras [25] assume the transitivity of trust to establish a relation between two entities without previous interactions. In this context, they model trust evaluation as a path problem on a directed graph. Routing protocols are the main target of this approach.

More closely related to VANETs and thus the case-study instantiation of our framework, Ostermaier et al. [19] analyze the performance of voting schemes for local danger warnings in VANETs. As mentioned earlier, voting schemes cannot properly express composite events. Another paper by Golle et. al. [10] proposes a framework for data validation in VANETs; it consists in comparing received data to a *model of the VANET* and accept their validity if both agree. Building and updating such a model in real-time may not satisfy the requirement of fast data processing in VANETs. Klein [15] describes the application of several data fusion techniques at the traffic management center.

There is little work on applying the Dempster-Shafer Theory to ad hoc networks, the most relevant to our work is the paper by Chen and Venkataramanan [5] that describes how DST can be applied to distributed intrusion detection in ad hoc networks. Siaterlis and Maglaris [23] apply DST to DoS anomaly detection. The notion of belief, disbelief, and uncertainty appears in the work of Jøsang [13]. The paper describes a certification algebra based on a framework for artificial reasoning called *Subjective Logic*.

VII. Conclusion

In this work, we developed the notion of data trust. We also addressed ephemeral networks that are very demanding in terms of processing speed. We instantiated our general framework by applying it to vehicular networks that are both highly data-centric and ephemeral. We evaluate data reports with corresponding trust levels using several decision logics, namely weighted voting, Bayesian inference, and Dempster-Shafer Theory. Simulation results show that Bayesian inference and Dempster-Shafer Theory are the most promising approaches to evidence evaluation, each one performing best in specific scenarios. More specifically, Bayesian inference performs best when prior knowledge about events is available whereas Dempster-Shafer Theory handles properly high uncertainty about events. In addition, the local processing approach based on either one of the above techniques converges to a stable correct value, which satisfies the stringent requirements of a life-critical vehicular network.

Acknowledgements

We would like to thank Maciej Kurant, Jean-Yves Le Boudec, Patrick Thiran, Slavisa Sarafijanovic, George Theodorakopoulos and the anonymous reviewers for their helpful feedback that helped shape the current version of this work.

References

[1] Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems – 5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ASTM E2213-03, 2003.

[2] IEEE P1609.2 Version 1 - Standard for Wireless Access in Vehicular Environments - Security Services for Applications and Management Messages. *In development*, 2006.

[3] M. Abdel-Aty, A. Pande, C. Lee, V. Gayah, and C. Dos Santos. Crash risk assessment using intelligent transportation systems data and real-time intervention strategies to improve safety on freeways. *Journal of Intelligent Transportation Systems*, 11(3):107–120, Jul. 2007.

[4] S. Buchegger and J.-Y. Le Boudec. A robust reputation system for P2P and mobile ad-hoc networks. In *Proceedings of P2PEcon'04*, 2004.

[5] T. M. Chen and V. Venkataramanan. Dempster-Shafer Theory for intrusion detection in ad hoc networks. *IEEE Internet Computing*, 9(6):35–41, Nov.–Dec. 2005.

[6] B. Cobb and P. Shenoy. On the plausibility transformation method for translating belief function models to probability models. *International Journal of Approximate Reasoning*, 41(3):314–330, Apr. 2006.

[7] L. Eschenauer, V. D. Gligor, and J. Baras. On trust establishment in mobile ad hoc networks. In *Proceedings of the 10th Int. Security Protocols Workshop*, 2002.

[8] S. Ganeriwal and M. Srivastava. Reputation-based framework for high integrity sensor networks. In *Proceedings of SASN'04*, 2004.

[9] V. Gligor, S. Luan, and J. Pato. On inter-realm authentication in large distributed systems. In *Proceedings of the IEEE Symposium on Security and Privacy'92*, 1992.

[10] P. Golle, D. Greene, and J. Staddon. Detecting and correcting malicious data in VANETs. In *Proceedings of VANET'04*, 2004.

[11] D.L. Hall and J. Llinas. An introduction to multisensor data fusion. *Proceedings of the IEEE*, 85(1):6–23, Jan. 1997.

[12] T. Jiang and J.S. Baras. Trust evaluation in anarchy: A case study on autonomous networks. In *Proceedings of IEEE Infocom'06*, 2006.

[13] A. Jøsang. An algebra for assessing trust in certification chains. In *Proceedings of NDSS'99*, 1999.

[14] W. Kiess, J. Rybicki, and M. Mauve. On the nature of Inter-Vehicle Communication. In *Proceedings of WMAN'07*, 2007.

[15] L.A. Klein. *Sensor Technologies and Data Requirements for ITS Applications*. Artech House Publishers, 2001.

[16] R. Kohlas and U. Maurer. Confidence valuation in a public-key infrastructure based on uncertain evidence. In *Proceedings of PKC'00*, volume 1751 of *Lecture Notes in Computer Science*, pages 93–112. Springer-Verlag, 2000.

[17] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems: theory and practice. *SIGOPS Oper. Syst. Rev.*, 25(5):165–182, 1991.

[18] J. Mundinger and J.-Y. Le Boudec. Reputation in self-organized communication systems and beyond. In *Proceedings of Inter-Perf'06 (Invited Paper)*, 2006.

[19] B. Ostermaier, F. Dotzer, and M. Strassberger. Enhancing the security of local danger warnings in VANETs - a simulative analysis of voting schemes. In *Proceedings of ARES'07*, 2007.

[20] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[21] M. Raya, P. Papadimitratos, and J.-P. Hubaux. Securing vehicular communications. *IEEE Wireless Comm. Magazine, Special Issue on Inter-Vehicular Comm.*, 13(5):8–15, Oct. 2006.

[22] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.

[23] C. Siaterlis and B. Maglaris. Towards multisensor data fusion for DoS detection. In *Proceedings of SAC'04*, 2004.

[24] Y. Sun, W. Yu, Z. Han, and K.J. Ray Liu. Information theoretic framework of trust modeling and evaluation for ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(2):305–317, 2006.

[25] G. Theodorakopoulos and J.S. Baras. On trust models and trust evaluation metrics for ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(2):318–328, Feb. 2006.

[26] Q. Xu, T. Mak, J. Ko, and R. Sengupta. Vehicle-to-vehicle safety messaging in DSRC. In *Proceedings of VANET'04*, 2004.

[27] C. Zouridaki, B. L. Mark, M. Hejmo, and R. K. Thomas. Robust cooperative trust establishment for MANETs. In *Proceedings of SASN '06*, 2006.