

Trust and Obfuscation Principles for Quality of Information in Emerging Pervasive Environments

Chatschik Bisdikian
IBM Research
Hawthorne, NY, USA
bisdik@us.ibm.com

Murat Sensoy and Timothy J. Norman
University of Aberdeen
Aberdeen, UK
{m.sensoy,t.j.norman}@abdn.ac.uk

Mani B. Srivastava
Univ. of California
Los Angeles, CA, USA
mbs@ee.ucla.edu

Abstract—The emergence of large scale, distributed, sensor-enabled, machine-to-machine pervasive applications necessitates engaging with providers of information on demand to collect the information, of varying quality levels, to be used to infer about the state of the world and decide actions in response. In these highly fluid operational environments, involving information providers and consumers of various degrees of trust and intentions, obfuscation of information is used to protect providers from misuses of the information they share, while still providing benefits to their information consumers. In this paper, we develop the initial principles for relating to trust and obfuscation within the context of this emerging breed of applications. We start by extending the definitions of trust and obfuscation into this emerging application space. We, then, highlight their role as we move from tightly-coupled to loosely-coupled sensory-inference systems. Finally, we present the interplay between trust and obfuscation as well as the implications for reasoning under obfuscation.

Index Terms—trust; obfuscation; Quality of Information; Value of Information; QoI; VoI; reasoning

I. INTRODUCTION

Even though not always at the forefront, trust is a key underlying element of any transactional activity. It characterizes the “bond” and “comfort” that the transacting parties share amongst themselves and impacts the utility of their mutual activities. In pervasive applications, transactional activities will typically involve the exchange of information between the parties. For example, the (electronic) sharing of: medical and health-care records from patients to health-care providers and between health-care institutions; information between city, state, and federal law enforcement agencies; information about population status between governmental and non-governmental organizations (NGOs) supporting emergency response and disaster relief efforts; “where I am” information to (seemingly) friends via social networking media; intelligence information, e.g., reporting people and vehicle movement patterns at cross-roads, borders, public venues, etc.

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

In all of the above, trust has primary and secondary implications. Naturally, transacting parties are *primarily* concerned with whether the information and its quality are as desired to satisfy the needs for which information was exchanged. However, there is a *secondary* concern as to whether the exchanged information, as a whole or in part, will be utilized only for stated or implied purposes and not for unspecified, possibly illicit, purposes. To this end, obfuscation serves as the mechanism for protecting against inappropriate use of shared information. It is the process by which information providers deliberately alter the content of the information they provide to protect sensitive information while allowing information consumers to still derive (hopefully) value from the information they receive [1].

With particular interest in protecting people’s private information, past studies in pervasive computing had focused on privacy-preserving obfuscation mechanisms. These include anonymization by removing or abstracting a person’s identifiers, and hiding, generalizing, or perturbing personal context, such as location [2], [3]. Trust, on the other hand, had been manifested through access policies to pertinent information [4].

However, while humans are an essential part for a significant portion of pervasive computing applications, we see the need to broaden the scope of trust and obfuscation to commensurate with the broader spectrum of sensor- and actuator-based *machine-to-machine* (M2M) and *Internet of Things* (IoT) [5] applications that emerge. These emerging areas are part of the so called *smarter planet* solutions [6] and include areas such as traffic and utility grid management, supply chain monitoring, infrastructure (and habitat) monitoring, environmental control, inter-city agency coordination, and so forth, that rely on fast-paced manipulation and analysis (of large amounts) of streaming data gathered from heterogenous collections of sensory sources and possibly from across administrative domains. In these settings, one may question, for example, the privacy implications of air-temperature measurements.

This paper is an early work in the area, setting the stage by identifying key components and establishing terminology for the principles upon which systems in these areas could be designed to deal with trust, obfuscation, and their interplay. With applications such as remote patient monitoring and social-networking-based (participatory) sensing lying at the

intersection of human-oriented and M2M pervasive applications, revisiting trust and obfuscation does not seek to obviate past work but rather augment it to cover the emerging smarter applications. To this end, we have drawn great inspiration from [7], which considers privacy issues in social network based applications, and, hence, serves as a bridge between past work and ours. Due to space limitations, prior art is provided as needed throughout the paper; we note though that there was none found aligned to the particular scope of our work.

The contributions of this paper are: (a) defining trust for consumers and providers and obfuscation within the context of our broader application space; (b) highlighting their role while migrating from tightly-coupled to loosely-coupled sensory-inference (M2M) systems; and (c) presenting the interplay between trust and obfuscation as well as the implications for reasoning under obfuscation. These contributions are covered in sections II, III, and IV, respectively.

II. TRUST AND OBFUSCATION DEFINITIONS

Sensor-enabled applications collect sensory information of their “world” (i.e., surroundings) to support reasoning and inferences about the world’s state and evolution alternatives. The significance and effectiveness of the inferences made and of the ensuing actions taken depend on the *quality of information* received and the value it brings to the sensing tasks at hand. These in turn are influenced by the relationships that are developed between information providers and consumers. With these relationships described by levels of *trust* and instantiated (at least, in part) by *obfuscation*, in this section we define these concepts for our purpose. We start, though, with the definitions of quality and value of information from [8].

A. Quality and value of information (QoI and VoI)

To accommodate the many uses of information (and information products), we have adopted a layered view of quality and value. Quality captures usage-independent “facts” about the information, and any usage dependencies are captured by its value. Specifically:

QoI: Quality of information represents the *body of tangible evidences* available (i.e., the innate information properties) that can be used to make judgments about the fitness-of-use and utility of information products.

VoI: Value of information represents an *outcome of such judgement*, which is an assessment of the utility of an information product when used in a *specific usage context*.

Both quality and value are described by collections of attributes (e.g., information metadata) such as accuracy, latency, and provenance for QoI, or relevance, timeliness, presentation/usability for VoI [8]. Trust relates to provenance while obfuscation impacts accuracy.

B. The consumer’s view of trust

Trust can be broadly defined as the willingness of one party (*trustor*) to rely on the actions of another party (*trustee*) [9]. Trust is critical in the large-scale, open distributed pervasive systems considered here, enabling interactions between

parties in uncertain and constantly changing environments. In our case, these parties are the information *consumers* and *providers*. For a consumer, we define:

Trust (consumer’s view): Represents the information consumer’s degree of belief that she can rely on the information that a provider has provided her with.

Later we will define a provider’s view of trust as well. Note that the term “belief” in the definition will allow us in the future to exploit both probabilistic and logic-based techniques for further investigating the pertinent systems [10].

If a consumer has a low level of trust in the information it receives, ensuing inferences may be considered unreliable. Trust in a body of available information, and hence the value that it may bring to bear on a process, is influenced by many factors. These include how this information has been collected, i.e., its provenance, or even how different pieces of information within this body relate and corroborate with each other, i.e., consistency in the body of information. For example, if information derived from a third-party provider conflicts with established facts and knowledge, the level of trust, and subsequently the value ascribed to such information, would be low.

Naturally, trust in information providers affects the trust in the information they produce. If a provider historically behaved in an untrustworthy manner, e.g., due to the consistent use of inappropriate, faulty, ill-calibrated, etc., sensory devices to gather information, any subsequent information provided may also be considered untrustworthy. Hence, the trust in providers and the information they produce are both highly correlated and build upon each other. However, this may not always be true, as in the case where information is provided by a third-party information aggregator and provenance linking it to its true source is incomplete or even missing.

We also consider a refinement of the above trust definition that focuses on the facts (or, meta-information) about the provided information. Under this refinement, trust relates to the *confidence in the meta-information* that is known about the information. In this case, a provider and the information they provide are trusted as long as its provider is forthcomingly truthful in disclosing this meta-information. For example, a provider may state that the accuracy of the information provided about the location of an object has an error variance σ_e^2 . Should this error be (statistically and verifiably) true, the provider may be trusted in the sense that whatever he provides is what he says he does. Whether the value of σ_e^2 is satisfactory to the consumer should be secondary with regard to trust.

The above refined definition may lend to a seeming paradox that all knowingly *incompetent* providers (e.g., these who claim very high σ_e^2) can be fully trusted. Of course, in this case, as long as such providers are indeed truthful, under any reasonably well-designed provider selection mechanism, the possibility that the consumer will engage with them, and hence the impact the information they provide will have in any inferences the consumers made, will diminish to 0. The latter is similar to the reaction expected by a consumer when dealing with highly untrusted providers as well. Hence,

one may draw parallels between trust and competence and not necessarily distinguish the two. We do, however, see benefits in this refined definition of trust, because it semantically separates trust and competence, with the former implying *intention* by the provider, and the latter capturing his information gathering *abilities*.

C. Obfuscation and the provider’s view of trust

The dictionary definition of *obfuscation* is “to make so confused or opaque as to be difficult to perceive or understand” (<http://www.thefreedictionary.com>). Common examples of obfuscation include GPS where, by adding noise and not sharing certain data, only lower accuracy location and time information is made available to civilian users, and satellite imagery where resolution is reduced when sharing with specific users.

In computing, obfuscation refers to the process of hiding certain data, while maintaining the usefulness of the data for an intended purpose, i.e., allowing authorized inferences [1]. It has been typically used to prevent leakage of personal information that would allow to, for example, identifying a patient in medical records, or identifying a person and/or creating permanent records of a person’s exact location presences when a mobile or online services providers sell location-based information to third parties, e.g., for marketing purposes. Techniques like anonymization and location abstraction are typically used for this purpose [1], [3].

However, when opportunities abound for collecting and fusing information derived from multiple providers (including physical sensors, knowledge bases, human observers, experts, etc.), sufficient knowledge may be gained and inferences be drawn that could go beyond the (seeming) intentions that caused the gathering of information in the first place, such as, when GPS tracking information or credit card purchase patterns are correlated. Thus, for a provider, we define:

Trust (provider’s view): Represents the information provider’s degree of belief that a consumer will use her information only for expressed purposes.

Note that restricting the scope of this definition to personal privacy leads to the inference violation problem in [7].

The above definition leads to a broader purpose for obfuscation to protect not only the primary piece of information provided, e.g., an object’s id or location, but derivative meta-information such as the location and capabilities of one’s sensor resources. For example, consider the implication of the latter in military coalition operations involving coalition members that develop ad hoc and transient alliance relationships, where one partner, owning high-quality sensory resources, shares localization information with another partner that has a need for but it is reluctant to share the location and nature of its sensors. This use scenario illustrates situations of collaboration across administrative domains involving parties of varying and evolving trust levels. Therefore, we define obfuscation more broadly as follows:

Obfuscation: Is the process that an information provider uses to influence *consciously* the set of inferences that could be made involving the information she provides.

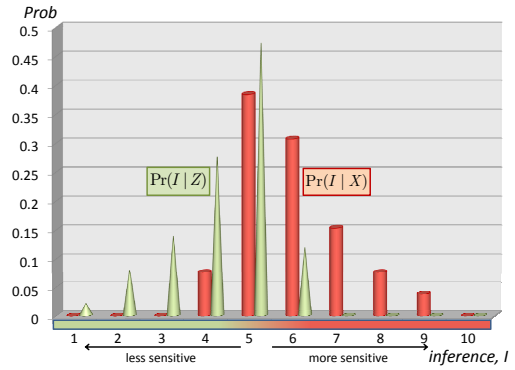


Fig. 1. An illustrative example of influencing inferences that can be made (red: original; green: after obfuscation).

This definition centers on the ultimate intention of a provider to affect deliberately the range of uses that the information it provides may have. Fig. 1 provides a *conceptual* visual for this case. It depicts two probability mass function assignments over a collection of (all) possible inferences; the collection is shown to be ordered for convenience. The (reddish) cylindrical-bar assignment represents the probability $\Pr(I|X)$ that inference I could have been made if access to the original information X were possible. On the other hand, the (green) conical-bar assignment corresponds to the probability $\Pr(I|Z)$ that inference I could be made when using the obfuscated information Z . In this example, both assignments exhibit a single peak at inference #5, which we assume also coincides with the “permissible” inference for the information provided based on the needs expressed by the consumer. However the available information may give credence to additional inferences as well. For the obfuscated information Z in Fig. 1, these inferences oppose others that X would have given credence to, which may represent an inference area that relates to sensitive information that needs protection.

Obfuscation may occur *in-situ* where, for example, a provider deliberately calibrates his sensor resources to collect data in a particular way, introducing uncertainties, e.g., bias and error, in addition to normal sensing errors that are beyond the provider’s control. Obfuscation may also occur after-the-fact by *post-processing* a piece of information prior to delivering it to consumers. The boundaries between *in-situ* and *post-processing* are not crisply defined. For example, GPS systems for civil uses can be viewed either way depending on where within the GPS hardware one may consider the boundary between the military and civilian portions of it lie.

In the next section we consider sensor-enabled inference systems that experience operational conditions where trust and obfuscation plays a role.

III. TIGHTLY- AND LOOSELY-COUPLED SENSORY-INFERENCE SYSTEMS

Sensing systems are deployed to enhance one’s capability to observe the world and estimate its state in portions of it of interest. Fusing the sensory observations collected produces

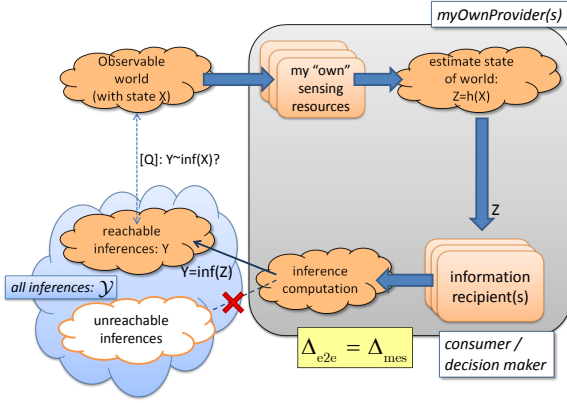


Fig. 2. A traditional (tightly-coupled) sensor-enabled inference system.

information about the presence (detection), type (classification), location, amplitude, motion (tracking), etc., of objects, conditions, and events in the physical world. Processing and correlating pertinent information builds situation awareness about the world that aids effective decision making and action taking. These simple steps of inference making are applicable from the rudimentary on-off operation of an energy-saving home thermostat to elaborate supply chain management systems that track and schedule various inventories to attain just-in-time delivery of goods at minimum cost.

Fig. 2 shows at a high-level the structure of a typical sensor-driven “inference” system that makes inferences regarding the world under observation. Specifically, suppose that a decision maker (a software agent representing the consumer) is interested in knowing the state X of certain objects in the (observable) world, e.g., the number, type, and velocity of vehicles crossing an intersection. The information about X would allow the consumer to make inferences (and take actions) about a potential situation of concern Y that may occur, such as the possibility of congestion or accident in a particular road segment, the potential presence of persons of interests in a particular area, or estimate latent quantities.

To support the consumer’s information needs, sensing resources are deployed (*myOwnProvider(s)*) trained on X . Based on the sensing capabilities and operational conditions of the deployed resources, the consumer experiences X via the sensing transformation $Z = h(X)$, e.g., Z could be a noisy variant of X , where the noise relates to the sensing process. With the sensory evidences Z at his disposal, the consumer can make inferences regarding the likelihood of various situations Y . In particular, the sensory evidences will steer the consumer towards a set of inferences $Y = \text{inf}\{Z\} = \text{inf}\{h(X)\} \subset \mathcal{Y}$ noted in the figure as *reachable inferences*. The set \mathcal{Y} represents the set of *all possible inferences* that the entire range of information Z that can be provided could influence. The (yellow) boxes with the Δ ’s in them will be explained in the next subsection.

What the figure shows is typical of a well-planned deployment of what we refer to as *tightly-coupled* sensory system where sensors and the applications that use their observations are deployed in coordinated manner. In a sense, applications

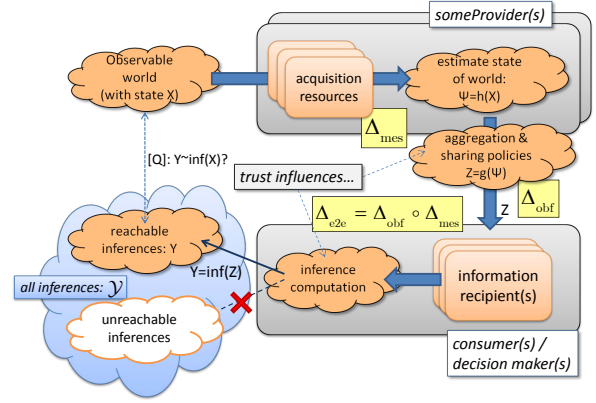


Fig. 3. A loosely-coupled, collaborative end-to-end inference system.

“own” the sensors feeding them with observations and, hence, have reasonable knowledge of (or, access to) the capabilities and deficiencies of deployed sensing resources. Thus, assessing the QoI and VoI of the sensory information they receive is, to a reasonable degree, within the applications’s control.

However the rigid structure of the tightly-coupled systems in Fig. 2 is being challenged by the new sensing opportunities and information collection paradigms that are emerging. These are exemplified by trends in multi-agency, heterogenous *Internet of Things* (IoT) deployments [5], crowd- and participatory-sensing applications [11], the aforementioned smarter planet applications [6], ad-hoc collaborative operations in multi-domain environments such as emergency response or coalition military operations, and so forth. In these cases, the tight, single-administrative-domain association between the information providers and consumers is challenged by more open *loosely-coupled* and, hence, more unpredictable, collaborative multi-administrative (and even no-administrative) domain associations. Knowledge about the capabilities of the sensing resources may be unavailable and unknown, or policy-constrained, and certainly of questionable reliance. In addition, shared data may be deliberately manipulated for various reasons. As a result, it becomes harder, in this case, for consumers to assess the QoI received and, hence, the ensuing VoI and risks associated when acting on any inferences made.

Undoubtedly, in such cases, the policies that affect the sharing of information and the underlying trust between the parties involved will play a key role in mediating effective collaboration. Fig. 3 shows an example of a collaborative, end-to-end inference system where information providers and consumers associate only as necessary. Like in Fig. 2, we use X to represent the state of the observable world, Z the way the consumer experiences it, and Y the inferences that could be made. In contrast to Fig. 2, the consumer may bind to third-party providers (*someProvider(s)*) who could obfuscate their information prior to releasing it, as noted by the transformation $Z = g(\Psi)$ in the figure.

As will be discussed further in section IV, information obfuscation policies could be affected by the level of trust that the provider has to different consumers. Likewise, the inferences made by the consumer can too be effected by

the “reverse” level of trust that consumers have towards the providers. Note that trust relationships between parties are built based on both direct and indirect, context-dependent associations between them. As a result, trust relationships may not necessarily be symmetric. Direct associations apply to “personal” experiences dealing with the provider, while indirect associations apply to knowledge gained through third-party opinions, recommendations, ratings, their direct associations, etc., typically associated with *reputation*.

A. Quality distortions

With regard to the (yellow) boxes in figures 2 and 3, measuring the QoI and the ensuing VoI that results from using the observations Z will relate to the “distance:”

$$\Delta_{\text{mes}} = \|X - Z\| = \|X - h(X)\|, \quad (1)$$

where the $\|\cdot\|$ represents an applicable operator. For example, if X and Z are n -dimensional vectors of real numbers, $\|\cdot\|$ could be the Euclidean distance between them, or, in the case of random vectors, the trace of their covariance matrix, or the information loss metric in estimating the distribution of a parameter [2], etc. The subscript “mes” stands for “measurement” to underscore that the measurement process is the key contributor to any QoI degradation experienced at this stage.

In the simple case, where Z is a noise-additive version of X , this distance may represent the error process, e.g., the error variance. However, in the general case, this distance needs to reflect not only, say, the accuracy but of the entire mismatch (i.e., *distortion*) that exists between the world state X and the collection of observations Z available about it, i.e., accommodate various QoI attributes, such as accuracy, latency, and spatiotemporal context [8].

In the absence of any additional deliberate distortion, in Fig. 2, we also equate Δ_{mes} to the *end-to-end* (e2e) distortion Δ_{e2e} , where, for simplicity, we ignore any impact that the communication networks will have on QoI, be it due to QoS degradation, e.g., increased latency, in-network information processing, e.g., data aggregation, etc. However, in Fig. 3 the information provider may further obfuscate the observations it has collected prior to sending them to consumers. In this case, the end-to-end distortion will need to account for the obfuscation as:

$$\Delta_{\text{e2e}} = \Delta_{\text{obf}} \circ \Delta_{\text{mes}} = \|X - Z\| = \|X - g(h(X))\|. \quad (2)$$

Note that for ad-hoc information-exchanging collaborations, the consumer may know neither the obfuscation $g(\cdot)$ nor the measurement $h(\cdot)$ processes, which at best can only be estimated from historical data including both direct (personal) and indirect (third party, social network) experiences.

The impact of QoI degradation to inferences, i.e., the *quality of inference* QoInf, made relates to the distance:

$$\Delta_{\text{inf}} = \|\text{inf}\{X\} - \text{inf}\{Z\}\| = \|\text{inf}\{X\} - \text{inf}\{g(h(X))\}\|, \quad (3)$$

which determines by how much the set of reachable inferences has been affected from using the observations Z instead of the original state X . We refer to this as the *interference distortion*.

Computing the QoInf is highly application specific and depends on the type of inferences desired, whether they are coarse-grained, e.g., categorical with few broad categories or classes of interest such as a four-wheeled vs. two-wheeled vehicle, or fine-grained, such as localization to within a few meters or even centimeters. Considering QoInf in the determination of human activity (referred to as context) for a tightly-coupled patient monitoring system (a categorical case), [12] introduces a QoInf metric based on the probability of error in estimating a context state.

QoInf is outside the scope of this paper, but we believe that intuitively appealing inference metrics based on probability of error [12] will not be sufficient and will need to be augmented when broader sets of applications and inference classes over loosely-coupled systems are considered. For example, [13] looks into relevance metrics for comparing pieces of desired and provided sensory information based on QoI and spatial context. Such metrics could form the conceptual basis for information distortion in the context of this research.

Next we highlight the interplay between trust and obfuscation and discuss reasoning under obfuscation.

IV. TRUST AND OBFUSCATION

A. Relationship between obfuscation and trust

Effective information sharing policies need to accommodate the dynamic nature of the relationships that consumers and providers develop in the loosely-coupled, pervasive applications and systems in consideration. They need to be responsive to the context underlying them and, in particular, obfuscation must be customizable to the various consumer, consumer classes and contexts a provider relates with. As a result, the level of obfuscation will commensurate with the level of trust that the provider has developed with each particular consumer.

Lowering the obfuscation level may allow a consumer a spectrum of inferences that include sensitive ones, such as those in the sensitive (red) area on the right-side of figure 1. This will increase the risk from sharing the information if the consumer uses it in a way not as (or, in addition to what) was originally expressed. However, if the provider has a sufficient level of trust that the consumer will not act with ill-intent, he may be willing to forgo high levels of obfuscation. Of course, the impact from breaching the provider’s trust could be severe. Therefore, the provider needs to always assess the risks from secondary uses of information prior to deciding the levels of obfuscation it should apply. Should a breach in trust be discovered by the provider, the consumer should expect that, at a minimum, the trust to the consumer will decline.

On the other hand, excessive levels of obfuscation may significantly impact the value that the consumer could derive from using the information shared. This could then negatively effect the level of trust the provider may enjoy with the consumer. Furthermore, certain types of obfuscation, such as deliberately adding bias or noise in the data to hinder certain inferences, can blur the boundary between the obfuscation and *deception*. Deception is defined as the act of misleading another through

intentionally false statements or fraudulent actions. If obfuscation done by a party is considered as a sort of deception by the other, the trust of the consumer towards the provider will be significantly damaged, leading to a cycle of decreasing trust, increasing obfuscation, and, ultimately, decreased information sharing. Avoiding this cycle is therefore an important consideration for information providers and consumers who depend on long-term interactions with others. An interesting use study about personal deceptive practices, i.e., lying, driven by a Bayesian-network model can be found in [14].

B. Reasoning under obfuscation

Obfuscation allows two parties to communicate and cooperate by sharing information while still protecting their interests. Hence, obfuscation is a key to establishing dynamic collaborative relationships between parties in loosely-coupled settings.

Obfuscation in GPS systems for civilian uses is public knowledge, as is the quality of localization attained; higher accuracies are possible with costly specialized hardware that the general public cannot (or is unwilling) to afford. In general, we see the following levels of knowledge about obfuscation: (a) *complete knowledge*, where both the type of obfuscation and its extent are known; (b) *partial knowledge*, where only the type of obfuscation is known; (c) *awareness*, where only that information is obfuscated is known; and (d) *ignorance*, where there is no knowledge regarding obfuscation.

Any knowledge about obfuscation could suggest remedial actions when dealing with the obfuscated information. Explicit information sharing “contracts” could be a source of such knowledge. Specifically, providers and consumers may establish *QoI-level agreements (QLAs)*, expressing QoI expectations for received information, as implied by the discussion on the four viewpoints of QoI in [15]. These QLAs may be further enriched by expressing the intended uses of information, i.e., the inferences to be made with the information provided. This, in turn, allows the provider to determine appropriate levels of obfuscation *a priori* based on the stated inferences and the trust the provider has toward the consumer.

At the same time, such agreements can form the basis for consumers to hold providers accountable as well. For example, if the agreement is to provide localization information, say, for moving trucks to within 20 meters of their true location, then the consumer may assess its trust to the provider based on whether he honors the agreement or not, as discussed with regard to the refined definition of trust in section II-B.

V. CONCLUSIONS

In the emerging world of multi-domain, M2M, smart pervasive applications, the trust between information producing and consuming parties with ephemeral associations will play key role. In such environments, obfuscating information will allow consumers to derive information of usable value from providers that have secrets to protect. The deployment and operational paradigms in these dynamic environments deviate significantly from the deployment and operational assumptions under which trust and obfuscation

has been traditionally considered, and, hence, need to be reconsidered. In this initial work, we have sought to lay the foundations for such a reconsideration.

We started by extending the definitions trust(s) and obfuscation to these emerging areas, highlighted applicable functional models of loosely-coupled systems, and discussed the interplay between these in reasoning with obfuscated information. Next we plan to enrich our foundations with concrete examples, take a deeper look at the interplay between trust and obfuscation and reasoning with obfuscated information, and consider dynamic obfuscation policies and the fusion of obfuscated information from various providers.

ACKNOWLEDGMENTS

The authors would like to thank C. Burnett, S. Chakraborty, A. Fokoue, L. Kaplan, F. Meneguzzi, N. Oren, and K. Sycara for enlightening discussions during the course of this project.

REFERENCES

- [1] D. E. Bakken, R. Parameswaran, D. M. Blough, A. A. Franz, and T. J. Palmer, “Data obfuscation: Anonymity and desensitization of usable data sets,” *IEEE Security and Privacy*, vol. 2, no. 6, pp. 34–41, 2004.
- [2] D. Agrawal and C. C. Aggarwal, “On the design and quantification of privacy preserving data mining algorithms,” in *20th ACM SIGMOD-SIGACT-SIGART Symp. on Principles of Database Systems (PODS’01)*, Santa Barbara, CA, USA, May 21–24, 2001.
- [3] M. Duckham and L. Kulik, “A formal model of obfuscation and negotiation for location privacy,” in *3rd Int’l Conf. on Pervasive Computing, PERSASIVE 2005*, Munich, Germany, May 8–13, 2005.
- [4] U. Hengartner and P. Steenkiste, “Access control to people location information,” *ACM Trans. Information System Security*, vol. 8, no. 4, pp. 424–456, Nov. 2005.
- [5] N. Gershenfeld, R. Krikorian, and D. Cohen, “The Internet of Things,” *Scientific American*, vol. 291, no. 4, pp. 76–81, Oct. 2004.
- [6] J. L. Martin, H. Varilly, J. Cohn, and G. R. Wightwick (eds.), “Special issue on: Technologies for a smarter planet,” *IBM Journal of Research and Development*, vol. 54, no. 4, July–Aug. 2010.
- [7] S. Chakraborty, H. Choi, and M. B. Srivastava, “Demystifying privacy in sensory data: A qoi based approach,” in *3rd IEEE Information Quality and Quality of Service Workshop (IQ2S’11, part of IEEE PerCom’11)*, Seattle, WA, USA, Mar. 21, 2011.
- [8] C. Bisdikian, L. M. Kaplan, M. B. Srivastava, D. J. Thornley, D. Verma, and R. I. Young, “Building principles for a quality of information specification for sensor information,” in *12th Int’l Conf. on Information Fusion (FUSION’09)*, Seattle, WA, USA, July 6–9, 2009.
- [9] D. Gambetta, *Trust: Making and Breaking Cooperative Relations*. Blackwell, 1990.
- [10] A. Jøsang and T. Bhuiyan, “Optimal trust network analysis with subjective logic,” in *SECURWARE’08*, Cap Esterel, France, Aug. 25–31 2008.
- [11] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, “Participatory sensing,” in *World Sensor Web Workshop (in ACM Sensys’06)*, Boulder, CO, USA, October 31, 2006.
- [12] N. Roy, A. Misra, S. K. Das, and C. Julien, “Quality-of-inference (QoINF)-aware context determination in assisted living environments,” in *1st ACM Int’l Workshop on Medical-Grade Wireless Networks (WiMD’09)*, New Orleans, LA, USA, May 18, 2009.
- [13] G. Tychogiorgos and C. Bisdikian, “Selecting relevant sensor providers for meeting ‘your’ quality information needs,” in *12th IEEE Int’l Conf. on Mobile Data Management (MDM 2011)*, Sweden, June 6–9, 2011.
- [14] X. An, D. Jutla, and N. Cercone, “Reasoning about obfuscated private information: who have lied and how to lie,” in *5th ACM Wkshp on Privacy in Electronic Society (WPES’06)*, Alexandria, Virginia, USA, Oct. 30–Nov. 3, 2006.
- [15] C. Bisdikian, L. M. Kaplan, M. B. Srivastava, D. J. Thornley, D. Verma, and R. I. Young, “Quality of sensor-originated information in coalition information networks,” in *Network Science for Military Coalition Operations: Information Exchange and Interaction*, D. Verma, Ed. IGI Global, 2010, pp. 15–41.