Pervasive and Mobile Computing 11 (2014) 168–187

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Fast track article

Inference management, trust and obfuscation principles for quality of information in emerging pervasive environments☆

Chatschik Bisdikian [a], Christopher Gibson [b], Supriyo Chakraborty [c], Mani B. Srivastava [c], Murat Sensoy [d,e,*], Timothy J. Norman [e]

[a] IBM Research, T. J. Watson Research Center, Yorktown Heights, NY, USA
[b] IBM United Kingdom Ltd, Hursley Park, Winchester, UK
[c] University of California, Los Angeles, CA, USA
[d] Computer Science, Ozyegin University, Istanbul, Turkey
[e] University of Aberdeen, Aberdeen, UK

ARTICLE INFO

Article history:
Available online 25 September 2013

Keywords:
Quality of information
Value of information
Risk of information
QoI
VoI
RoI
Obfuscation
Inference management

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, information transformation, such as obfuscation, is used to manage the inferences that could be made 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 inference management and the role that trust and obfuscation plays in it 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 the tightly-coupled to loosely-coupled sensory-inference systems and describe how quality, value and risk of information relate in collaborative and adversarial systems. Next, we discuss quality distortion illustrated through a human activity recognition sensory system. We then present a system architecture to support an inference firewall capability in a publish/subscribe system for sensory information and conclude with a discussion and closing remarks.

1. 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 transacting parties share with each other and impacts the utility of their mutual activities.

☆ This research was sponsored by the US Army Research Laboratory and the UK 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 US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Dr. Şensoy thanks to the US Army Research Laboratory for its support under grant W911NF-13-1-0243 and The Scientific and Technological Research Council of Turkey (TUBITAK) for its support under grant 113E238.

* Corresponding author at: Computer Science, Ozyegin University, Istanbul, Turkey.
E-mail addresses: bisdik@us.ibm.com (C. Bisdikian), christopher.gibson@uk.ibm.com (C. Gibson), supriyo@ee.ucla.edu (S. Chakraborty), mbs@ee.ucla.edu (M.B. Srivastava), murat.sensoy@ozyegin.edu.tr (M. Sensoy), t.j.norman@abdn.ac.uk (T.J. Norman).

1574-1192/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.pmcj.2013.08.003

Such transactional activities will typically involve the exchange of information between the parties, such as the collection of patient medical and healthcare records and their sharing between health-care providers; the collection of census, polling, or other populace information and their sharing between governmental agencies and non-governmental organizations, e.g., while participating in humanitarian aid activity or supporting emergency response and disaster relief efforts; the collection and sharing of "where I am" information to (seemingly) friends via social networking media; the gathering and sharing of intelligence and surveillance information, e.g., reporting people and vehicle movement patterns at cross-roads, boarders, public venues, etc., and their sharing between city, state, and federal law enforcement agencies.

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 the stated or implied purposes and not for other unspecified, possibly illicit, purposes. To this end, the information masking, or *obfuscation*, serves as the mechanism for protecting against the inappropriate use of shared information. It is one of the processes by which information providers deliberately conceal and, in general, alter aspects of the information they provide to protect sensitive information. Such information alterations affect the ability of making certain inferences while allowing information consumers to still hopefully derive value from the information they receive [1]. These techniques can range from manipulating the information content directly, e.g., removing a name entry from an information record; translating, generalizing, or adding noise to a location entry; altering features or dimensionality of the information, e.g., truncating Fourier coefficients, concealing variance information, etc. They aim to influence, i.e., manage, the set of inferences that can be made with [the parts of] the information shared by altering the quality of the shared information and hence altering its value as well. In this paper, we will summarily refer to such inference managing techniques as obfuscation, even though, strictly speaking, obfuscation (e.g., as in [1]) typically refers to a specific subset of these techniques.

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 reflect 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 (e.g.,. patient) 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 heterogeneous 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 extended version of our contribution to IQ2S 2012 [7] and represents a foray in this area. It aims in 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 and inference management in dynamically created, loosely-coupled sensory systems. To this end, we have drawn great inspiration from [8], which considers privacy issues in social network based applications, and, hence, serves as a bridge between past work and ours. Related and background work is provided as needed throughout the paper; we note though that we could find none that was aligned along the particular scope of this work.

The contributions of this work from [7] are:

- defining trust for consumers and providers and obfuscation within the context of our broader application space;
- highlighting their role while migrating from tightly-coupled to loosely-coupled sensory-inference (M2M) systems; and
- introducing the notion of various quality distortions to capture the various transformations made to information prior to sharing it.

To the above contributions, this paper's additional contributions are:

- providing an extended interpretation of the quality and value of information (QoI and VoI) in the presence of risk in collaborative and adversarial settings;
- giving a concrete example of inference management using accelerometer data for human activity inference; and
- providing a system architecture and highlight an implementation of an inference firewall capability for managing inferences on a publish/subscribe system for accessing and delivering sensory information.

The organization of the paper is as follows. Section 2, presents the extended definitions of trust and obfuscation along with the interpretation value, risk, and quality of information in collaborative and adversarial settings. Section 3 introduces representations of tightly- and loosely-coupled inference systems along with a discussion of quality distortions. Section 4 gives an example of inference management for a human activity detection use case and Section 5 presents a system architecture for an obfuscation and inference management capability in a sensory information delivery system. Section 6 presents related work on reputation and trust assessment; other related and background work is provided as need throughout the paper. Finally, we conclude in Section 7 with a discussion and concluding remarks.