

ON TRUST AND PRIVACY IN CONTEXT-AWARE SYSTEMS

Waleed Wagealla, Sotirios Terzis, Colin English and Paddy Nixon

The Smartlab research group
Department of Computer and Information Sciences
University of Strathclyde, Glasgow, Scotland.
waleed.wagealla@cis.strath.ac.uk

Recent advances in networking, handheld computing and sensors technologies have led to the emergence of context-aware systems. The vast amounts of personal information collected by such systems has led to growing concerns about the privacy of their users. Users concerned about their private information are likely to refuse participation in such systems. Therefore, it is quite clear that for any context-aware system to be acceptable by the users, mechanisms for controlling access to personal information are a necessity.

According to Alan Westin “privacy is the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information is communicated to others”¹. Within this context we can classify users as either *information owners* or *information receivers*. It is also acknowledged that *information owners* are willing to disclose personal information if this disclosure is potentially beneficial. So, the acceptance of any context-aware system depends on the provision of mechanisms for fine-grained control of the disclosure of personal information incorporating an explicit notion of benefit.

In the SECURE project², we envisage that trust could be exploited to protect users’ privacy, in the sense that reasoning about the trustworthiness of *information receivers* allows us to decide the amount of information that can be disclosed to them. Our approach also incorporates an explicit notion of risk. Reasoning about the risk involved in interactions allows us to adjust the amount of disclosed information according to their expected benefit. For example, information could only be revealed to trustworthy users, i.e. users that are expected to provide significant benefits to the *information owner*. Moreover, each *information owner* can specify his/her own privacy policy, in which he/she can articulate his/her preferences by adjusting his/her risk aversion (acceptable costs) for the various outcomes of an interaction. Allowing information owners to specify their own privacy policy is very important because users have significantly different attitudes towards privacy. Furthermore, our model of trust and risk supports learning from past interactions. We observe the outcomes of each interaction and we change the *information receiver*’s trustworthiness to reflect our observations.

We have applied the approach outlined above in a smart space scenario. The scenario looks at a university department equipped with a context information server, which tracks the location of users and can provide location information on demand. The access to the location information is controlled by the tracked user’s privacy policy (*information owner*), which is expressed in terms of the trustworthiness of the requesting users (*information receivers*). The application of our approach to this scenario has provided some useful insight on the engineering of trust-based privacy solutions. We are currently evaluating the performance of our approach in the context of this scenario.

¹ Alan F. Westin. Privacy and Freedom. Publisher: Bodley Head.

² <http://secure.dsg.cs.tcd.ie/>