Foot-driven computing: our first glimpse of location privacy issues (Invited paper)

Frank Stajano

University of Cambridge Computer Laboratory

Ubiquitous computing has been a fashionable research theme for the past twenty years, so much so that many research groups have felt the urge to give it a different name (pervasive computing, calm computing, ambient intelligence etc etc) in order to claim that they were doing something new or at least slightly different from everyone else. One of these many alternate names has been "context-aware computing", to suggest systems and devices that would sense the "context" of a situation and behave accordingly: for example, a mobile phone might sense that its owner is "in a meeting" and automatically switch from ringtone to vibration mode.

I have been professionally involved in ubiquitous computing research since 1992 [14, Chapter 2], when I joined the ORL (Olivetti Research Ltd) laboratory in Cambridge, UK, and I have long been somewhat sceptical of the vague semantics generally attributed to the term "context" in the above usage. When we cut out the fog of more or less useful abstractions introduced by the middleware layer, we find that in a majority of practical cases "context" essentially boils down to *location*; therefore, "Location-driven computing" or "Location-based services" are names I find more descriptive and concrete. At ORL we also jokingly used to say "foot-driven computing", not referring to any hypothetical pedals but to the practice of influencing the behaviour of computer systems just by walking around as opposed to doing so by typing at a keyboard.

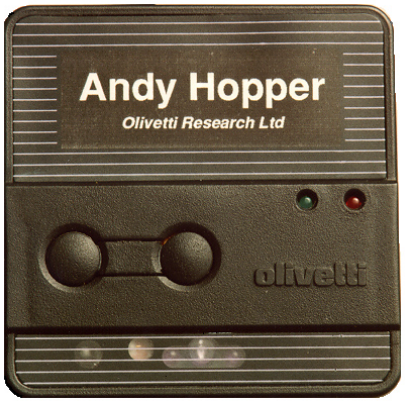

One of the ORL inventions that had the greatest impact on the worldwide research scene was the Active Badge [16], the first indoor location system: there is an image of an Active Badge in Weiser's [19] classic *Scientific American* article. A small infrared-emitting name tag worn by personnel, the Active Badge told the system the position of its wearer and enabled mobility features such as rerouting of phone and video calls, "teleporting" (moving one's desktop to the nearest workstation, without disrupting the running applications) and, last but not least, simply allowing researchers to find their colleagues within the three floors of our building. The Active Badge (1989), adopted and deployed by such pioneering research institutions as Xerox PARC (which soon hired the badge's inventor Roy Want) and MIT Media Lab, not to mention our own University of Cambridge Computer Laboratory, gave us a glimpse of the possibilities opened by the foot-driven-computing paradigm. More

ACM SIGSPATIAL Special Issue, July 2009.

importantly from a research perspective, it also raised plenty of questions, particularly on privacy, to which our daily interaction with the system allowed us to respond with experience-driven practical answers, such as enforcing reciprocity ("you are allowed to see my location only if I am allowed to see yours, and I get notified that you are monitoring me whenever you do"). Social conventions naturally developed around the use of the badge, including the nuance between leaving it on one's desk face down (which turned it off and meant "I have left or I don't want to be tracked") and leaving it on the desk face up (which to the system looked indistinguishable from when the wearer actually was in the office, and meant "I don't want to be tracked but I don't want the system to be aware of that").

With only slight technological adjustments, similar location privacy issues now directly affect hundreds of millions of users worldwide. All the press and television reporters who visited our laboratory throughout the Nineties and were quick to throw up their hands in horror at the privacy invasion that the Active Badge represented are now carrying a mobile phone in their pocket, which allows their location to be tracked not just within one building but across the city, the country and the globe. The merit of the Active Badge was to grant us a window of a few years in which to think seriously about those issues before they became truly pervasive and universal.

Did we make good use of that head start? We certainly did—witness Jackson's [9] early work on user-definable access control for Active Badge sightings and Beresford and Stajano's [3] work (actually conducted on the Active Bat [17], the higher-resolution successor of the Badge) that introduced the *mix zone* concept and a quantitative criterion for measuring location privacy. A mix zone is an area in which no spatial monitoring takes place. (We don't just want to disallow spatial monitoring completely or we'd have to give up all benefits of location-based services. But designating specific zones as non-monitorable is an acceptable compromise.) Each tag is known to the system through a pseudonym and, whenever a tag enters a mix zone, it changes to a different pseudonym. If the mix zone was originally empty, then changing to a new pseudonym offers no protection because a hostile observer can clearly deduce that the new pseudonym coming out of the mix zone belongs to the lone tag that had previously entered it. If, however, the mix zone was not empty, then, when a tag comes out of the mix zone with a new pseudonym, the hostile observer doesn't know for sure *which* of the tags that previously entered it has now changed into this new pseudonym. The mix zone technique thus offers unlinkability between the observable segments of the location trace of the tag. A quantitative measure of the amount of location privacy thus gained is obtained by computing the entropy of the population of the mix zone. The logarithm of the size of the population would be a simpler first-order estimate, but we use the entropy to take into account also what the observer knows about the movements of the tags that have entered the zone.

Elsewhere, Gruteser and Grunwald [8] introduced the techniques of spatial and temporal cloaking. Later, Buttyán et al. [5] applied the mix zone technique to protect location privacy in vehicular networks.

Was all that enough? Perhaps not, judging from the lack of location privacy safeguards in, say, today's mobile phone systems. But we eventually also learnt that, despite what good-spirited researchers might think and despite what people might say in your face if you ask them, the general public doesn't actually put a very high value on privacy in general [1] and on location privacy in particular [6], at least until something really bad happens to them personally.

To some of us it is absolutely evident that protecting location privacy is a desirable goal, and one that we have a moral duty to pursue as responsible architects of the technologies that will affect billions of citizens of our world whether they like it or not. The ability to track individuals

wherever they go, and even more so the ability to data-mine such location history retrospectively on a global scale, can be misused as an Orwellian tool of blackmail, surveillance and political oppression. In the inspiring words of Phil Zimmermann [21], whom I often quote on this subject,

> When making public policy decisions about new technologies for the government, I think one should ask oneself which technologies would best strengthen the hand of a police state. Then, do not allow the government to deploy those technologies. This is simply a matter of good civic hygiene.

But I am well aware that these privacy-oriented values are not universally shared and that a full debate on motivations would exceed the scope of this brief note. Still, I feel that researchers who concentrate only on the technical aspects and completely dodge the debate on values are myopic and irresponsible.

Mobile phones are only one of many ways through which the location of an individual can be tracked: anyone wishing to protect location privacy must look at a much wider picture. In the modern electronic society we all leave behind what Alan Westin [20] presciently defined as "data shadow" way back in 1967. Most of our purchasing and travel transactions are recorded in back-end databases [7]. The owners of such databases often have an economic incentive to take active measures to link individual transactions back to the same person, for example by offering loyalty cards, in order to be able to engage in price discrimination [13]. Governments, who also deploy other pervasive location-monitoring tools such as CCTV cameras that recognize car numberplates (or even faces, in a not-too-distant future), are keen to centralize and cross-link their own databases, often under the excuse of the fight against terrorism [2]. Whenever our laptop establishes a wi-fi connection with an access point, it leaves some traces, at many levels in the protocol stack, of having visited that location. The same happens with Bluetooth connections [10] and of course with every kind of wireless technology, of which mobile phones are just a special case. There has been widespread debate on privacy issues raised by RFID tags [11] and location privacy is among them. As first pointed out by Weis et al. [18], the "constellation" of tags of objects carried or worn by a person is likely to have enough of an invariant "core" that a person can be re-identified from one day to the next (e.g. you might be wearing the same glasses and wristwatch and overcoat as you did yesterday, even if you changed your shirt, socks and so forth).

In summary, location privacy is a hard unsolved research problem and one that applies to a variety of modern systems. What matters most is not so much the specific technology used to acquire the location information (mobile phones, loyalty cards, CCTV cameras, wi-fi laptops, bluetooth gadgets, RFID tags or whatever) as the back-end databases that store all the sightings. The underlying problem is "denied oblivion" [15], the fact that storage is so cheap that there is no incentive ever to delete any data. Technological safeguards on their own will be insufficient to protect individuals from abuse and will have to be complemented by regulatory and societal protections. From the technical viewpoint, however, since the potential for abuse is already so great, anyone offering a new location-based service or technology would do well to think about its undesirable side effects and how to minimize them at the design stage. It's the moral equivalent of "when you design this new vehicle, don't just go ahead blindly but please think about how much it will pollute". Unfortunately the location privacy problem is made harder by the misalignment of incentives of the players involved: those who could do the most to solve it are those who are least affected by the problem and the least concerned about it.

Location privacy, though pervasive, multi-faceted and unsolved, is certainly not the only security concern in location-based computing, though. We can only mention them in passing, but *secure positioning* and *secure position attestation* are two other significant classes of location security problems. The former consists of "I want to determine where I am, despite the presence of active attackers who might send me fake signals instead of the ones I expect from my references" [12], as might be of interest to the navigation system of a ship in pirate-infested waters or, in a totally different context, to a region-coded video player that does not trust its owner. The latter problem can instead be described as "I want you to prove to me that you really are where you say you are", a subcase of which is "I want you to prove to me that you are within x metres of this point" [4]. Both have a variety of practical applications and, in the grand scheme of things, they may be easier to tackle than location privacy, given that they don't suffer from the same problem of misalignment of incentives.

References

[1] A. Acquisti and J. Grossklags. "Privacy and Rationality in Individual Decision Making". *IEEE Security & Privacy*, **3**(1):26–33, 2005.

[2] R. Anderson, I. Brown, T. Dowty, W. Heath, P. Inglesant and A. Sasse. "Database State". Technical Report, The Joseph Rowntree Reform Trust, 2009.

[3] A. Beresford and F. Stajano. "Location Privacy in Pervasive Computing". *IEEE Pervasive Computing*, **2**(1):46–55, January 2003.

[4] S. Brands and D. Chaum. "Distance-Bounding Protocols". In *Proc. EUROCRYPT 93*, LNCS 765, p 344–359.

[5] L. Buttyán, T. Holczer and I. Vajda. "On the Effectiveness of Changing Pseudonyms to Provide Location Privacy in VANETs". In *Proc. ESAS 2007*, LNCS 4572, p 129–141.

[6] G. Danezis, S. Lewis and R. Anderson. "How much is location privacy worth?" In *Proc. WEIS 2005*.

[7] S. Garfinkel. *Database Nation*. O'Reilly, 2000.

[8] M. Gruteser and D. Grunwald. "Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking". In *Proc. MobiSys 2003*, p 31–42.

[9] I. W. Jackson. *Who goes here? Confidentiality of location through anonymity*. Ph.D. thesis, University of Cambridge, February 1998.

[10] M. Jakobsson and S. Wetzel. "Security Weaknesses in Bluetooth". In *CT-RSA*, LNCS 2020, p 176–191.

[11] A. Juels. "RFID Security and Privacy: A Research Survey". *IEEE Journal on Selected Areas in Communication*, **24**(2), February 2006.

[12] M. G. Kuhn. "An Asymmetric Security Mechanism for Navigation Signals". In *Proc. IH 2004*, LNCS 3200, p 239–252.

[13] A. M. Odlyzko. "Privacy, economics, and price discrimination on the Internet". In *Proc. ICEC 2003*, p 355–366.

[14] F. Stajano. *Security for Ubiquitous Computing*. Wiley, 2002.

[15] F. Stajano. "Will Your Digital Butlers Betray You?" In *Proc. WPES 2004*, p 37–38.

[16] R. Want, A. Hopper, V. Falcao and J. Gibbons. "The Active Badge Location System". *ACM Transactions on Information Systems*, **10**(1):91–102, January 1992.

[17] A. Ward, A. Jones and A. Hopper. "A New Location Technique for the Active Office". *IEEE Personal Communications*, **4**(5):42–47, October 1997.

[18] S. A. Weis, S. E. Sarma, R. L. Rivest and D. W. Engels. "Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems". In *Proc. Security in Pervasive Computing 2003*, LNCS 2802, p 201–212.

[19] M. Weiser. "The Computer for the Twenty-First Century". *Scientific American*, **265**(3):94–104, September 1991.

[20] A. Westin. *Privacy and Freedom*. Atheneum, 1967.

[21] P. R. Zimmermann. "Testimony of Philip R. Zimmermann to the Subcommittee on Science, Technology, and Space of the US Senate Committee on Commerce, Science, and Transportation", 1996.