Will Your Digital Butlers Betray You?

Frank Stajano
University of Cambridge

ABSTRACT

The cost of data storage is now so low that there is little necessity ever to delete anything. The consequence is *denied oblivion*—digital systems that remember forever and can be data-mined retroactively, years after the event, ignoring any privacy promise under which the original data may have been acquired.

Even for systems under your own control, though, the situation is alarming. As your capacious digital butlers faithfully collect as much data as possible about you, your private information is increasingly likely to become compromised.

New solutions are needed. But technical countermeasures alone are not the whole story.

Categories and Subject Descriptors: K.4.1 [Computers And Society]: Public Policy Issues—*Privacy*.

General Terms: Security.

Keywords: Denied Oblivion, Data Mining.

1. DENIED OBLIVION

Processor speeds have increased by three orders of magnitude ($1,000\times$) over the past 20 years. A more significant though less frequently glorified improvement has occurred in mass storage capacity: the size of today's hard disks has increased by about 4.5 orders of magnitude ($30,000\times$) over the same time span. It is now possible to collect data to an extent that was previously unthinkable.

Thirty years ago, no three-letter agency would have had the budget to monitor *all* the international telephone traffic of a country with the population of Canada. Yet a rough estimate, the numerical details of which have been omitted in this concise position paper, shows that a 10 M$ server farm could now transcribe all that speech into text in real time. Much more significant, though, is the fact that a whole month's worth of searchable full-text transcripts would fit into a single 300 GB hard disk. The running costs of storage have become practically nil.

There is no economic requirement ever to delete anything. Whatever was once digitized is now stored forever. This property, which I shall call **denied oblivion**, is the source of many new privacy problems. The privacy violation may not occur right now; but, since everything is logged, there is always the sword of Damocles that some intrusive data mining may occur retroactively at a later date.

The time shift inherent in denied oblivion is responsible for another major threat, namely that data acquired for one purpose by one agent will later be accessed and searched for another purpose by another agent. The regulations under which data was acquired may have changed; the agent that originally acquired the data may have gone out of business; but the data itself is still there, ready to be mined by its new owner who, in practice and despite theoretical claims to the contrary, is no longer bound by the original rules. Captured data tends to have a much longer lifetime than the privacy policy under which it was captured[1].

The interception-of-communications scenario, while not unrealistic, is just one of many possible examples. The commercial world offers many more, from supermarket loyalty cards to adware that spies on your web browsing habits. Garfinkel's *Database Nation* will raise your awareness on the extent to which data collection and dossier building already takes place in today's society.

Perhaps one of the most worrying aspects of the problem is the apathy of the general public towards it: most people can be bribed out of their shopping privacy by the 1% discount offered by the loyalty card. Except for small vocal minorities, the complacency of the public extends to much more intrusive developments such as the ubiquitous *Minority Report*-style CCTV surveillance to which people are subjected in the UK, or the US-VISIT program under which all the non-American attendees of this "Workshop on Privacy in the Electronic Society" were fingerprinted and photographed at their port of arrival in the US[2]. Both practices are dismissed by many honest members of the public as a small price to pay in order to ensure "the safety of the country".

"Denied oblivion" simply means that those fingerprints taken from you at the airport will stay on file forever. Policies may come and go, but this acquisition will never be undone.

2. A CONTROVERSIAL DIGITAL BUTLER

Intrusive data acquisition about you by other parties is not necessarily the most worrying development. Consider also legitimate data acquisition about you by devices under your control, later misused against you by adversaries who take over these devices against your wish.

[1]Assuming, optimistically, that one was in place at the time of acquisition.

[2]This practice, introduced in January 2004 for citizens of certain nations, was widened to most of the remaining ones in October 2004. It would not be entirely illogical to predict that something similar will eventually extend to domestic citizens too.

WPES'04, October 28, 2004, Washington, DC, USA.

Let me introduce a hypothetical recording device—call it "Omnirec"—that would forever store everything you hear. This is feasible today: modern digital dictaphones already compress voice-grade audio to 10 MB/h, so the 24×7 storage requirement amounts to less than 100 GB/year. A few years from now, a portable Omnirec would also be capable of transcribing and indexing all speech, thereby making it searchable; and it might be capable of recording video as well as audio. With an only slightly more daring stretch of the imagination, the Omnirec might eventually sample your auditory and visual neurons rather than external microphones and cameras. This invention would be a wonderful memory prosthesis: you would be able to recall, instantly and accurately, any event at which you had been present.

This invention would also, however, raise a number of serious privacy concerns, both towards others and towards you. The most significant difference between the prosthetic and the wetware memory has to do with *transferability* of content from the original viewer to other persons. With standard human memory you can tell a third party what you saw and heard; but this indirect and imperfect report is quite different from a video. There is also the significant difference that the recipient has no reliable way to distinguish between the objective facts and your own (intentional or unconscious) additions and omissions.

Imagine you visit my home, where all walls are covered in books. Later, by reviewing past Omnirec footage at your leisure, you could compile a list of all the thousands of books in all the rooms that you visited, learning the titles of many more of my books than you could possibly have noticed on your own. Imagine further that someone asked you whether I have an embarrassing or controversial title—say *Mein Kampf* or *How to build an atom bomb*. You didn't notice it when you visited; but with the Omnirec you could search all the books that entered your field of vision, instead of just saying "I don't know". If you found it, you could then show this third party the video frames of the book on my shelf—a much more convincing and damning report than just "I saw it". (And let's not get into forgeries...)

It is easy to see that showing Omnirec media to a third party would therefore make both of you guilty of eavesdropping. As a matter of fact, under many jurisdictions, if you had been using your Omnirec without my knowledge, you would be deemed guilty of some form of spying *regardless* of whether you showed the recordings to anyone else. Some may consider this an exaggeration—a consequence of the accidental fact that, with existing or foreseeable technology, there is no way to prevent you from showing your Omnirec recordings to others. If it were possible to build a memory prosthesis with the same guaranteed non-transferability property as the human memory, then there might be grounds to consider single-user operation of the Omnirec as quite distinct from spying, just as we have no objections to people having a good memory, or to people writing down accurate debriefing notes after having witnessed something of which they want to keep a record.

So long as enforcing this limitation is technologically impossible, though, the Omnirec remains a sinister spying device that others might not like you using. Both reactions to and justifications for the Omnirec may be similar to those that apply to Mann's Wearcam. The device is also, however, a dangerous double-edged sword, as illustrated by the nightmare scenario in which your Omnirec is actually owned and operated by the secret police, who will also beat you up if you attempt to turn it off.

Encrypting the content before storing it, with a key only known to the wearer, is an obvious first attempt towards a solution. Unfortunately it still leaves the wearer at liberty to reveal the content to third parties, therefore not altering the status of the Omnirec as a potential spying device. It also doesn't prevent a determined adversary from obtaining the content from the wearer—it just forces this adversary to resort to more convincing attacks, of the kind in which the locution "brute force" reverts to its literal meaning as opposed to the one usually attributed to it by cryptographers.

Various steganographic solutions have been proposed to prevent the extraction of encrypted data under duress. None is entirely satisfactory when the threat model includes taking you to Orwell's "Room 101".

3. SUMMING UP: THE RIGHT TO REMAIN SILENT

The threats to privacy are changing.

From a technological standpoint, the spectacular improvements in capacity and affordability of mass storage are responsible for *denied oblivion*: data, once acquired, is never forgotten. Since data will outlive the privacy policy under which it was acquired, retroactive data mining is a first, obvious privacy problem.

Our growing reliance on digital butlers, from cellphones to car navigation systems, means that more and more data nuggets about us are being logged on an ongoing basis, with ever-increasing temporal, spatial and semantic resolution. This in itself would not be a privacy threat if not for the fact that your digital butlers can be "forced to speak" with much less effort and risk than that needed to force *you* to say something you wouldn't. The more our butlers become privy to intimate information about our us and our lives, the more serious this threat becomes. Transferring information from brain to butler enhances availability but threatens confidentiality.

Sometimes people tolerate surveillance and intrusion in exchange for the promise of greater security. Wouldn't the job of the police be a lot easier if they knew everything about every citizen? Of course. They might even be able to *prevent* crimes, rather than stop them. Yet not many would like the idea of having to live in a society in which the secret police kept detailed files on everyone.

As technologists we enjoy devising access control cryptotricks that might make the butlers a little safer. Technical countermeasures, however, can always be overcome. At a higher level, therefore, it is important to discuss *principles*.

To the extent that your digital butlers are increasingly knowledgeable about your thoughts, as the purposefully exaggerated Omnirec example illustrates, their lack of "rights" is a problem. There is practically nothing—in our cultural perception of right and wrong, much less in law—to protect them when information is being techno-tortured out of them. Your digital butlers do not enjoy the "right to remain silent". Yet, given what might be in them, to dismiss this issue as ridiculous just because they are not sentient beings is to leave the back door open to the thought police.