

Total Recall: Are Privacy Changes Inevitable?*

[appeared in CARPE 2004]

William C. Cheng
Computer Science Dept. &
IMSC
University of
Southern California
<bill.cheng@acm.org>

Leana Golubchik
Computer Science Dept.,
EE-Systems Dept., IMSC, & ISI
University of
Southern California
<leana@cs.usc.edu>

David G. Kay[†]
Donald Bren School of
Information and Computer
Sciences
University of California, Irvine
<kay@uci.edu>

ABSTRACT

Total Recall is a system that records an individual perspective of the world using personal sensors such as a microphone in a pair of glasses or a camera in a necklace. There are many applications of Total Recall – patients accurately recording what they’ve recently eaten, students replaying any part of a class, and so on—that can significantly improve people’s quality of life. However, data recorded by such a system may be also used by the judicial system without the consent of the user or of those being recorded. Pervasive use of systems like Total Recall will likely change our social structure as memory becomes vastly more reliable and complete.

It is natural then that privacy advocates might consider such technology dangerous because such data can be used in unanticipated ways by government agencies or third-party civil litigants. In this paper, we discuss privacy concerns in the context of systems like Total Recall and propose a solution that may alleviate some of these concerns. We discuss the ramifications of this solution and its possible implementations.

Categories and Subject Descriptors

K.4.1 [COMPUTERS AND SOCIETY]: Public Policy Issues—*Privacy*; K.6.5 [MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS]: Security and Protection—*Authentication*

General Terms

Legal Aspects, Design, Security

*This research was funded in part by the NSF EIA-0091474 and ANI-0070016 grants as well as by an Okawa Foundation Research Award. It was also funded in part by the Integrated Media Systems Center, a National Science Foundation Engineering Research Center, Cooperative Agreement No. EEC-9529152.

[†]Member, State Bar of California

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CARPE’04, October 15, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-932-2/04/0010 ...\$5.00.

Keywords

Privacy, Personal sensors, Record and playback

1. INTRODUCTION

Technology’s ultimate purpose is to improve people’s quality of life. One aspect of improving quality of life is to provide or enhance abilities that are missing, diminishing, or otherwise in need of improvement. Memory is one such ability.

This paper focuses on legal/social as well as technical issues in the context of a project called Total Recall [1]. The idea of Total Recall is to be able to remember when an event happened, where it happened, who was there, why it happened, and how we felt. Total Recall aims to amass memories, experiences, and ultimately knowledge from an *individual* perspective and for a multitude of individuals.

It starts with the use of *personal* sensors, like a microphone in a pair of glasses or a camera in a necklace; it would include other sensors, all of which would record an *individual* perspective of the world. (This recording is intended to be continuous and under user control.)

But, Total Recall is not simply an individual memory enhancer. It could have many other applications, for example in health care, education, and support of elderly and people with disabilities:

- Placing a microphone array on a hearing impaired person’s glasses can allow collection of audio that gets converted to text and displayed on a PDA in near real time.
- Being able to recall a patient’s food intake and recent environments can help discovery of allergies.
- Monitoring food intake of diabetics can provide automatic warning signals when appropriate.
- Being able to review a patient’s state before and after a serious health problem, like a heart attack, can help doctors arrive at a more accurate diagnosis in an emergency situation.

Some people’s first reaction, when they hear about a system that records everything, at every moment, and everywhere you go, is fear. After all, who knows who else might get their hands on this information? But the reality is that this is already starting to happen around us. For instance, there are cameras (webcams) everywhere—on traffic lights, on highways, in buildings [6]. We expect that a world that is constantly recording will come sooner or later.

There are many benefits to such technology, as well as drawbacks, and keeping them in balance requires both technical and legal/social solutions. From the technological point of view, we need to design and build systems that provide proper security, privacy, and integrity mechanisms. Such mechanisms should enable a wide variety of policies so that legal/social policy development is not hampered by a paucity of technical alternatives. Without technical flexibility, the inevitable development of technology may result in poor policy by default.

There are always scary uses of technology, but we believe this technology can result in much good, if done right. We have enhanced our eyesight with glasses and our timekeeping ability with watches, so why not enhance our memories as well?

Although there are many significant technical challenges that need to be addressed in the context of Total Recall and its applications, in this paper, we focus on *privacy* and *security* issues. Such issues are of great concern to many, as indicated, for instance, by the July 2004 issue of IEEE Spectrum, which includes a number of articles on sensors and related privacy concerns [5, 6, 8, 10]. Although these works have a different focus (mostly sensors embedded in the environment), they do indicate that privacy is a significant issue in general. In this paper we focus on systems of personal sensors that are under users' control to some degree, unlike the sensors in the environment.

Without properly addressing privacy and security concerns, technologies such as Total Recall might have grave consequences (or their wide acceptance might be hampered). To provide proper technological solutions, we must first understand the privacy concerns in both technical and legal/social settings, and that is our focus here.

Specifically, we first explore the potential privacy concerns and consequences within a legal or social setting. We argue that important concerns do exist in the context of systems such as Total Recall and that from a technical point of view, hooks and mechanisms are needed that can support future legal and social changes. We then present one possible technical solution that could provide such a mechanism. We do not suggest that this is a complete or a definitive solution to all privacy concerns. However, we hope that this work can serve as a good start for fruitful discussions.

2. PRIVACY CONCERNS

As described above, Total Recall will record a user's experience (for example, in audio and video) continuously. Continuous recording means that the user need not decide consciously in advance what interactions or experiences merit recording, much like real memory. (Even continuous recording may have its exceptions. A user would likely turn it off in private moments—assuming the user remembers, since a continuous service is easy to set and forget.) We might also expect Total Recall to provide a complete and accurate picture of the events it records, but even that has its limits. Surely it is more comprehensive than unaided memory, but audio and video recordings portray just one point of view, limited by the user's position and environment, and are subject to technical constraints such as bandwidth and resolution.

Total Recall presents plenty of interesting technical issues: how to handle the volume of data, retrieval of particular "memories," annotation, alteration, and so on.

But the pervasive nature of Total Recall also gives rise to a range of legal and social questions. Since Total Recall's high-level goal is to improve quality of life, we must consider its broad social and legal effects as well as the social and legal issues that might affect the technical design of Total Recall. It is useful to consider these questions *before* the technology becomes pervasive; once a technology is widely deployed, as a practical matter it is generally too

late or too hard to make significant changes.

One of these issues is *privacy*. We first explore some legal and social issues in the context of privacy and then focus on possible technological features that could address these privacy concerns.

A key characteristic of Total Recall records is that third parties can gain access to them more easily than they can to human memory. Short of truth serums, hypnosis, or interrogation techniques, even the existence of a particular human memory can usually be concealed. Encryption and obscure access protocols can hamper third-party access to computer-based "memories," but the existence of a Total Recall system implies that certain records exist. The judicial system, moreover, can compel production of these records, which of course gives rise to privacy concerns.

2.1 Can we record everything we see and hear?

A threshold question is whether using Total Recall, recording everything the user sees and hears, is even legal under current law. As with every legal question, the answer is that it depends.

US wiretapping laws are the first line of legal control on recording, and they vary from state to state [4]. Some states require that all parties to a conversation consent to it being recorded; others require just one party. Some states have different rules with respect to video recording (which may in turn vary according to whether the recording is unattended or whether it's nudity that's being recorded).

Obtaining consent of persons being recorded poses logistical problems for a system like Total Recall, once it is in pervasive use. Mechanisms might be adopted that provide for implied consent, by analogy with a recurring beep during recorded telephone conversations, but a proliferation of perceptual cues (beeps or flashing lights) might degrade the quality of the recorded information and of the real-world experience. Subliminal system-mediated protocols would ameliorate the perceptual issues but would exclude people without system access (as the audio or video cues would exclude people with certain disabilities).

Consent to being recorded also implies an understanding of the use to which the recording may be put. A longstanding principle of fair information practices holds that *information gathered for one purpose not be used for another without the subject's consent* [3]. The person being recorded cannot know the extent to which Total Recall recordings may be used; indeed, neither can the user at the time of recording. Memory is such a central part of the human experience, and its universal availability (subject to the usual human vagaries) is so fundamental, that any advance limitation on its use would alter its nature completely.

Apart from statutory consent requirements, the fundamental principle is that *people are entitled to privacy in situations where privacy is their "reasonable expectation."* It is reasonable to expect privacy when alone at home; it is unreasonable to expect privacy when walking on a public street. If a tourist with a video camera can record a street scene for private use, there is little difference legally in the use of Total Recall. However, a significant practical difference, as yet unrecognized legally, would arise if Total Recall became widely used. The camera-bearing tourist is relatively rare; the chances of one's image being captured are low. But if Total Recall were as common as cellphones, any passer-by could be almost certain of being recorded many times. This overlapping web of recorded "memories" would be a qualitative change in the heretofore ephemeral nature of quotidian activity.

2.2 Once we have it, what can we do with it?

The legal and social issues do not end if we determine that Total Recall recordings are legally permissible. We must also consider the permissible uses of the recorded material.

The individual Total Recall user would have primary access to the recordings; appropriate security measures could largely prevent unauthorized access by others. The user's private use of legally obtained recordings is largely unrestricted, although publishing without permission the likenesses or other personal information of the recorded subjects could give rise to liability.

A user's Total Recall recordings would, however, be available to the judicial system.

In a criminal proceeding against the user, the protection against self-incrimination provided by the Fifth Amendment to the US Constitution would likely not protect Total Recall data. The Fifth Amendment protects a person from giving testimony that would relate to his or her commission of a crime. But criminal defendants are routinely required to produce records, documents, and even DNA samples, so disclosure of Total Recall data could likewise be compelled, even if they would incriminate the user.

In civil lawsuits, even where the user is an uninvolved third party who merely observed some relevant event, a court could compel production of Total Recall records, just as a court today can compel production of electronic mail records. Moreover, once the user knows that the records are requested by the court, destruction or alteration of those records would also give rise to legal liability.

One could imagine the combination of Total Recall systems with radio-frequency identification so that information captured by Total Recall would include the radio-frequency ID (RFID) information of other Total Recall users in the vicinity. The comprehensive web of recorded activity surrounding the incident in question would be feasible to identify and obtain.

Indeed, in the current US environment of terrorist threats, the political climate supports access to information by law enforcement, even without judicial intervention, if that information is perceived to have national security implications.

2.3 Will we see legal support for Total Recall privacy?

There is some question whether the legal system will develop enhanced privacy protection for Total Recall records.

The law does evolve to accommodate new circumstances, including new technology. Rules of evidence exist to ensure that courts consider only trustworthy information. A centuries-old rule of evidence states that if the original of some written document is available, the original must be introduced; a copy won't be allowed.

This rule dates back to when a copy was a handwritten copy, which of course could contain transcription errors. With the advent of carbon copies and photocopies, the rule has evolved to allow "duplicate originals" produced by mechanical means. But this change was reactive and evolutionary, occurring after the technologies had been deployed and, most significantly, without significant controversy or opposition. In theory, new rules of evidence could be adopted to exclude Total Recall recordings or limit their use, but there is reason to be skeptical that such rules would in fact be created because of the tension between legitimate privacy concerns and legitimate needs for the data.

Pro-active protection is harder to achieve. Since US courts decide actual cases based on existing situations, it is the legislature's role to consider policies for situations that have yet to occur. But legislative interest in privacy issues is hard to create and sustain in general. Legislation against unsolicited commercial e-mail (spam) was only briefly of interest and to date has been only partially effective, and that issue currently affects a broad constituency. For potential abuse of an as-yet-undeveloped technology, the likelihood of protective legislation in advance is low, not only because of limited public concern but also because of a general reluctance to inhibit

the development of rapidly evolving technologies. On the other hand, by the time any technology has even the smallest commercial foothold, its commercial supporters are likely to oppose any restrictions as an interference with the value of their investment and its economic consequences (such as employment).

The law, in general, changes more slowly than technology developments. This is generally desirable, since we would not want the rules by which society operates to fluctuate as rapidly as we see new system releases. We should expect, therefore, that systems like Total Recall will be deployed before a comprehensive policy on the privacy of its recordings is in place and that as a result, changes in the nature of privacy we experience are all but inevitable.

However, a vital role still exists for technologists: designing highly configurable systems with enough technical "hooks" to enable whatever privacy policy decisions are eventually arrived at. Below we explore one such hook, and in the following section we explore the technical issues involved.

2.4 Could technology help?

Given privacy concerns about government or third-party attempts to obtain Total Recall records and about third parties' unauthorized (if initially unintentional) recordings, it is reasonable to explore potential technical measures to address these concerns.

Suppose, for example, that Total Recall were in universal use and that it provided each user with the ability to become "invisible" to other Total Recall users by setting a preference ("Don't record me now") that other Total Recall systems would recognize; those other systems would record everything else but smoothly and seamlessly remove the "invisible" user from the record.

If every user on the street may have invisibility settings on or off at any time, a given user really doesn't know whether his or her recording is an accurate reflection of reality because no user can keep conscious track of which passers-by were invisible to Total Recall. Watching an event recorded previously, the user would have no way of evaluating the accuracy of the recording.

In the scenario above, the recording is no longer "authentic." One way to reduce or eliminate the use of Total Recall recordings in legal proceedings would be for this inauthenticity to become broadly understood. However, such comprehensive inauthenticity would also diminish Total Recall's utility for its intended applications. It is desirable to find an appropriate tradeoff between inauthenticity and intended utility.

One solution we explore in this paper is to mark each piece of data (at a granularity to be determined later) with an "authenticity bit." Briefly, this bit would be on for (portions of) recordings that were unmodified from the original data capture. The bit would be off where Total Recall made modifications (based on either automatic system modifications or those performed directly by the user). The authenticity transition can only go in one direction, from unmodified (authentic/original) to modified; the modification status would not be reversible. (Section 3 explores this approach in detail.)

The ability to track whether the recorded data has changed since the original capture has many advantages. For instance, a recording guaranteed to be unmodified (authentic) would be some protection against other forms of evidence. Conversely, if the authenticity bit were off by default, one might have some protection against the non-consensual use of recordings in legal proceedings; if this were the convention, more users might consent to be "visible" more often.

If all Total Recall system builders and all builders of player software could be required to follow these conventions of respecting users' invisibility settings and maintaining the one-way nature of

authenticity transitions, would Total Recall records avoid the embrace of the judicial system?

Probably not, because the legal system does not require provable certainty. It hardly even recognizes absolute certainty as a concept. The legal system just provides different levels of required proof—by a preponderance of the evidence (A is more believable than B), clear and convincing (A is a lot more believable than B), beyond a reasonable doubt (nobody could reasonably believe B). We cannot tell the legal system to ignore information; even if a Total Recall record’s authenticity bit is off, the legal system will make up its own mind.

An imagined exchange like the following illustrates how a Total Recall record, with or without an authenticity bit, might be *shown to be* credible and introduced as evidence.

Counsel: Now, Mr. Smith, you were wearing your Total Recall system on the day in question?

Witness: Yes, I was.

Counsel: And was it functioning properly?

Witness: As far as I know.

Counsel: Did you have occasion to look back at some of the recorded data on that date?

Witness: Yes, I’m sure I did; I use it all the time.

Counsel: And did you recognize any flaws in those recordings, or any missing conversations?

Witness: No.

Counsel: How long have you been using Total Recall?

Witness: About six months now.

Counsel: And in that time, have you ever noticed any flaws in the recordings, or any missing conversations?

Witness: No, I haven’t. Of course, I only have my own memory to compare against, and it’s been getting hazy lately ...

Counsel: Your honor, move to strike the last part as non-responsive.
The Court: Sustained.

Counsel: Mr. Smith, do you know about the “authenticity bit” that Total Recall attaches to its recordings?

Witness: It’s some preference setting. I haven’t paid much attention to it.

Counsel: You just use Total Recall to record your various activities during the day.

Witness: That’s right.

Counsel: And you refer back to it pretty frequently?

Witness: Yes.

Counsel: And it helps you out a lot, is that correct?

Witness: It’s pretty useful.

Counsel: And you have come to rely on it as you go about your work?

Witness: Sure.

Counsel: Your honor, I move to admit Mr. Smith’s Total Recall records as Exhibit No. 1.

Opposing counsel: Objection, your honor. Lack of foundation. Without the authenticity bit on the recording, we have no guarantee of its accuracy.

Counsel: The witness has testified that he finds the records accurate and reliable. It’s a question of fact, and without specific evidence of tampering or other unreliability, the recording should be admitted.

Probably the court would rule to admit the evidence under current law.

As we indicated earlier, the rules of evidence could change; Total Recall records without an authenticity bit could be made inadmissible explicitly. While there is reason to be skeptical on practical,

political grounds that such a change would occur, technical functionality (such as the authenticity bit) could provide the hooks on which policymakers could hang a legal protection scheme. In Section 3 we discuss the technical issues surrounding the authenticity bit.

3. OUR APPROACH

In this section, we present an approach to implementing the *authenticity bit* discussed above. We also present the rationale behind our approach. The term *authenticity bit* was used earlier for ease of exposition. The actual implementation of this concept requires more than a single bit per data block. (Below we assume that the continuous data stream is divided into blocks; determining their granularity is not essential for the purpose of this discussion.) Note that many of the security concepts mentioned in this section can be found in standard systems security textbooks, such as [7] and [9].

3.1 Authenticity

We assume that a user (Alice) of the Total Recall system will carry a wearable device that has a reasonably large amount of storage capacity so that data collected from her personal sensors can be stored on the device for a significant period of time. (Voice-only recording at 8 kilo-samples per second, 16 bits per sample, mono-channel, and assuming a five-to-one compression ratio, would take under 300 MB of storage per day, well within the 2 GB capacity of currently available compact flash cards.) We also assume that strong encryption (such as triple-DES) is used to encrypt the data blocks. We further assume that the storage device (such as a compact flash card) can be removed by Alice so that she can edit the data easily if she desires. Data may be uploaded to a server when Alice is connected to the network. Since data can reside on the wearable device for a long period of time, we need to provide a mechanism to verify whether or not the data on the device is original and authentic. One way to achieve this is to have the device digitally sign every block it produces. (For the purpose of this discussion, we assume that public-key cryptography is used for digital signatures. It is understood that other cryptographic schemes can be used with reduced security benefits but with increased performance benefits.)

In public-key cryptography, a private-key can be used to produce a digital signature and its corresponding public-key can be used to verify the digital signature. If the private-key can be kept secret, a digital signature can be used to provide proof that the data has not been modified since it was created. A common way of keeping the private-key secret is to embed it in a *cryptographic smartcard*. Such a smartcard is temper-resistant and can produce digital signatures without ever exposing the private-key. (Temper-resistant means that one cannot break into the smartcard without being detected.) If a bit in the data block is modified, the verification of the digital signature will fail.

Although each block of data has an attached digital signature, the time the data was produced may be in question. Even if a block contains a timestamp issued by the device, the clock on the device may be inaccurate since the battery on the device can be drained and the clock can be reset by Alice. One solution to this problem is to require the device to synchronize its clock with a clock server when it is connected to the network. However, reliable and verifiable clock synchronization may be difficult to achieve.

An alternate solution is to use third-party authentication. Instead of having a server send its clock value to Alice’s device, Alice’s device sends a cryptographic hash of a block to a public *notary* server and asks the notary server to produce a timestamp and digitally sign the timestamp and the hash. (A cryptographic hash has




Figure 1: A digital signature for every block.




Figure 2: Only blocks B_a and B_b are notarized.

the *bit-commitment* property which implies that the data block cannot be modified without detection.) Alice’s device then attaches the signed timestamp and hash to the data block. This is similar to the timestamp step in the *Bistro System* [2]. In our case, by digitally signing the timestamp and the hash, the notary server provides proof that the data block was received at the time indicated in the timestamp. (In this section we assume that the notary server is trusted to perform its functionality accurately.)

Without loss of generality, let’s assume that a new data block is produced every second. Figure 1 depicts a set of data blocks with their digital signatures. $B_{j,i}$ denotes the i^{th} data block of day j . For confidentiality, $B_{j,i}$ has been encrypted using K_j , which is the day key for day j . Here i ranges from 0 to $I - 1$, where I is the number of blocks produced per day ($I = 86400$ in this example). Block $B_{j,-1}$ is a special block (and therefore shaded differently) that contains the encrypted day key K_j ; encryption is performed using the public-key of Alice’s device. (Note that once something is encrypted with the public-key of Alice’s device, the only way it can be decrypted is with the physical presence of her smartcard.) Here, $h(X)$ denotes the cryptographic hash of X , $\tau_{j,i}$ is the timestamp issued by the notary server for $h(B_{j,i})$, a *plus* symbol denotes concatenation, and $DS[X]$ denotes the digital signature for X .

In the remainder of this section, we will omit the day index j when it is clear from the context.

We further assume that each block is timestamped by Alice’s device. Although the clock value of this timestamp cannot be trusted to be genuine, it can be used as a form of numbering. After a new block is produced, it is important to get it notarized as soon as possible. Otherwise, Alice may have enough time to modify the data. This can happen if network connectivity is intermittent or unavailable for most of the day, which is currently common for many people. (It is a requirement for Total Recall to allow Alice to modify her data. What we do not want is for Alice or anyone else to be able to claim that the data is original after it has been modified.) Also, producing digital signatures is computationally very expensive, even for servers. So, we would like to avoid the need for signing every data block.

If we can sign a block occasionally and create dependencies between the blocks, we may still be able to provide provable authenticity but at a lower cost and without requiring continuous or frequent connectivity. One way to produce dependencies between consecutive blocks is to use *chaining*. We can embed the cryptographic hash of block i in block $i + 1$. Figure 2 depicts the case where only blocks B_a and B_b are notarized, where $a < b$ (further

assume that there are no other notarized blocks between B_a and B_b). The left-pointing arrows depict dependencies.

Although the blocks are chained together, the encrypted day key is available to Alice, and Alice owns the device that contains the smartcard which can decrypt the day key. This makes it possible for Alice to modify all blocks between the time the day key is released and time index a . (For example, Figure 1 implies that the day key is released at the beginning of the day.) This modification can be performed by Alice any time she desires (even after time index b). Also, Alice may be able to modify all blocks between time index $a + 1$ and $b - 1$, inclusive, as they are being generated, as long as these modifications are performed before time index b .

Therefore, the day key should only be released when it is no longer being used, and there should be no gaps between the time the day key is released and the notarized block is generated. Since there can be many days before Alice has network connectivity, the encryption key should not be associated with the calendar, i.e., the day key should be replaced by a session key (whose use can span multiple days or a fraction of a day).

Figure 3 depicts the case where blocks between B_{a+1} and B_b , inclusive, are encrypted using the same session key, $K_{a+1,b}$. This session key, i.e., $K_{a+1,b}$, is encrypted using the public-key of Alice’s device and placed in the special block $B_{b/s}$. Instead of just sending the cryptographic hash of B_b for notarization, Alice’s device now sends the hash of the concatenation of B_b and $B_{b/s}$ for notarization. Now, Alice cannot modify *any* block without it being detected.

In the above example, blocks at time indices a and b can be thought of as *authentication anchors*.

There are additional advantages to this approach. (1) The device can decide when to get a block notarized. In the above example, time indices a and b can be any time when there is network connectivity. The frequency of getting a block notarized is also flexible. Even when a device has frequent network connectivity, it can choose to contact a notary server infrequently to reduce network traffic and notary server’s workload. (2) A notary server’s digital signature does not have to be verified on the fly. Therefore, the public keys of notary servers do not have to be stored on Alice’s device. Even if Alice can set up her own network and spoof the IP address of a notary server to trick her device into releasing the session key $K_{a+1,b}$, the fact that a new session key is used will result in eventual detection of data not being authentic. (3) How long a particular session key is used can also be controlled. Alice’s device can be offline for a long period of time. Thus, if there is




Figure 3: Proper release of session keys.




Figure 4: Data modification.

concern about encrypting with the same session key for too long, multiple session keys can be used (and released at the same time). *Lamport's hash chains* can also be employed so that every block is encrypted with a different key.

One drawback of the approach described in this section is that a session key only resides in the working memory of Alice's device. If that device dies, all data between the time of the crash and the last notarized block cannot be authenticated. If the device is not reliable, the notarization frequency should be increased.

3.2 Modifications

In order to turn the *authentication bit* off (or the *modified bit* on) for some data blocks, as suggested in Section 2, Alice's device can remove digital signature blocks within appropriate data blocks. Continuing with the example in Figure 3, removing the digital signature block attached to B_b would make it impossible to authenticate all blocks between B_{a+1} and B_b , inclusive, since they *may* have been modified. In order to ensure that there is no way to claim that the data blocks have original data, Alice's device must *modify* all blocks between B_{a+1} and B_b , inclusive. To do this, Alice's device must decrypt all blocks between B_{a+1} and B_b , generate a new session key ($K_{a+1,b}^*$), and re-encrypt all these blocks. The new session key is encrypted with the public-key of Alice's device and stored in $B_{b/s}$. Figure 4 depicts modifications where the modified blocks are marked with asterisks. Note that block chaining no longer exists between modified blocks as the hashes have been zeroed out.

If Alice only wants to modify a few minutes of her data, the above simple approach will not work because the granularity of modifications may be too coarse, i.e., determined by the density of authentication anchors. One possible solution is to upload all the related data blocks to a third-party server, verify the authenticity of all the data blocks, digitally sign additional data blocks, then remove appropriate original digital signature blocks. One drawback of this approach is that the third-party server may make a copy of the data before modifying it, without informing Alice. Therefore, this server must be a server trusted by Alice.

4. CONCLUSIONS

In this paper we focused on privacy and security concerns in the context of Total Recall, a personal information system intended to record and aid in remembering when an event happened, where it happened, who was there, why it happened, and how we felt. The paper focused on exploration of privacy concerns in a legal/social setting. It also offered a potential technical mechanism which, in combination with appropriate legal/social policy, could address at least some of the privacy concerns.

We also note that there are other broader, perhaps more speculative social implications of Total Recall deployment. For instance, eventually there might be an expectation that everyone would use a Total Recall system, just as we pretty much expect everyone to have a telephone today. Then one might imagine this courtroom query: "So, Mr. Jones, you turned your Total Recall off when you met Mr. Smith. What were you trying to hide?"

Moreover, as technology evolves, the skills and knowledge we find valuable change. Horsemanship is no longer a survival skill; knowing Morse code is not necessary for wireless communications; knowing how to get anywhere is becoming obsolete with GPS; arithmetic skills are less necessary with the ubiquity of calculators. So, will human memorization become less important a skill?

We also note that this paper is not intended as a definitive solution, but rather as a starting point for future discussions. Much is left to consider, technically, socially, legally, philosophically, and so on. However, the potential of such technology for improving quality of life is great, and hence, worth pursuing.

In conclusion, we believe that systems like Total Recall will get built, they will have valuable uses, and they will radically change our notions of privacy. Even though there is reason to be skeptical that there will be any meaningful legal protection for the privacy status quo, we believe that useful technologies are largely inevitable, that they often bring social changes with them, and that we will inevitably both suffer and benefit from their consequences.

We have air pollution. We have exploding airplanes. We have red-light cameras and sidewalk cameras. We have cellphone records and credit card records that say where we are and when. There is

not much to stop someone from collecting all that data now. As responsible technology builders and researchers, we should do our best to consider the possible long-term consequences of the systems we develop so that we can design into them, as much as possible, the flexibility necessary to address those consequences in whatever ways society chooses.

We hope this paper will be a starting place for pointing out the potential ramifications as well offering an initial technical mechanism to help enable future legal or social policies.

5. ACKNOWLEDGMENTS

The authors would like to thank Ching-Hua Chuan and the anonymous referees for their helpful comments on earlier versions of this paper.

6. REFERENCES

- [1] Internet Multimedia Lab at the University of Southern California. *Total Recall: a Personal Information Management System*. <http://bourbon.usc.edu/iml/recall/>, 2004.
- [2] W. C. Cheng, C.-F. Chou, L. Golubchik, and S. Khuller. A secure and scalable wide-area upload service. In *Proceedings of the 2nd International Conference on Internet Computing, Volume 2*, pages 733–739, June 2001.
- [3] Center for Democracy and Technology. *Generic Principles of Fair Information Practices*. <http://www.cdt.org/privacy/guide/basic/generic.html>, 2000.
- [4] Reporters Committee for Freedom of the Press. *A Practical Guide to Taping Phone Calls and In-Person Conversations in the 50 States and D.C.* <http://www.rcfp.org/taping/>, 2003.
- [5] H. Goldstein. Mike Villas’s world. *IEEE Spectrum*, pages 45–48, July 2004.
- [6] H. Goldstein. We like to watch. *IEEE Spectrum*, pages 30–34, July 2004.
- [7] C. Kaufman, R. Perlman, and M. Speciner. *Network Security: Private Communication in a Public World, 2nd Edition*. Prentice Hall, 2002.
- [8] J. Kumagai and S. Cherry. Sensors and sensibility. *IEEE Spectrum*, pages 22–28, July 2004.
- [9] W. Stallings. *Cryptography and Network Security: Principles and Practice, 2nd Edition*. Prentice Hall, 1999.
- [10] V. Vinge. Synthetic serendipity. *IEEE Spectrum*, pages 35–44, July 2004.