

# Architecture is Politics: Security and Privacy Issues in Transport and Beyond\*

Bart Jacobs

Digital Security group, Radboud University Nijmegen, the Netherlands  
[www.cs.ru.nl/B.Jacobs](http://www.cs.ru.nl/B.Jacobs)

**Abstract.** This paper discusses the political relevance of ICT-architecture through a review of recent developments in the Netherlands, involving the bumpy introduction of a national smart card for public transport and the plans for electronic traffic pricing based on actual road usage of individual cars. One of the underlying themes is the centralised or decentralised storage of privacy-sensitive data, where centralised informational control supports centralised societal control.

## 1 Architectural issues

Mitchell Kapor is one of the founders and the first chairman of the Electronic Freedom Foundation (EFF), an international non-profit organisation that aims to defend civil rights in the digital age. Kapor has coined the phrase *Architecture is politics*, see also [7]. It does not refer to architecture as the design and structure of buildings or bridges, but of ICT systems. In a broad sense ICT-architecture involves the structural organisation of hardware, applications, processes and information flow, and also its organisational consequences, for instance within enterprises or within society at large.

Why is this form of architecture so much a political issue—and not chiefly a technical one? An (expanded) illustration may help to explain the matter.

Many countries are nowadays in the process of replacing traditional analog electricity meters in people's homes by digital (electronic) meters, also known as smart meters. The traditional meters are typically hardware based, and have a physical counter that moves forward under the influence of a metallic disc that turns with a speed proportional to the electricity consumption. This process takes place within a sealed container, so that tampering is not so easy and can be detected. It is in general fairly reliable. Such domestic meters may last for decades.

From the (information) architectural perspective the main point is that the usage data are stored decentrally, in one's own home. The reading of the meters

---

\* Based on a keynote talk at the *Computers, Privacy, and Data Protection* (CPDP) conference, Brussels, 16-17 jan. 2009. This paper is intended for the CPDP proceedings. Version of May 15, 2009.

must also be done locally, either by the customer or by a representative of the electricity provider. This is of course cumbersome.

Hence, one of the main reasons for the introduction of digital meters is to optimise this process of meter reading. These digital meters can be read from a distance, by the electricity provider. What happens<sup>1</sup> is that the digital meters report the electricity consumption every 15 minutes to the provider, for instance via a dedicated GSM connection or via a special signal on the power line. The electricity provider thus stores very detailed usage information about domestic electricity consumption via such remote meter reading. We are moving from a decentralised (information) architecture to a centralised one!

Within this centralised set-up electricity providers offer their clients the option to view their own electricity consumption patterns via a personalised webpage. An important argument in favour of the introduction of such smart meters is an ecological one: energy conservation. The idea is that if customers are more aware of their actual consumption, they will become more careful and buy for instance more energy-efficient appliances.

Remarkably, these two arguments—remote meter reading and energy conservation—have nothing to do with the architectural change, from decentralised to centralised storage of usage data. Remote reading can also happen within a decentralised architecture: consumption data are stored locally within the smart reader, and only passed once every three months (say) to the electricity provider. Also, if it is beneficial for energy preservation that I get direct access to my electricity consumption, why should my usage data make such a detour, and first go to the electricity company, after which I can see them again via the web? One can also put USB connectors on smart meters so that people can read their data locally, via some open standard. Then I can decide for instance to connect my meter to a big screen in my hallway that shows consumption levels and starts flashing if my consumption is above a certain threshold.

Moreover, a detailed picture of my electricity consumption gives information about my way of life: what time I get up or go to bed (per day), whether or not I have visitors, when I'm on vacation *etc.* Especially this last point shows once again that privacy protection is important for one's own personal security. If people at the electricity company can see when I'm on holidays, burglars may become interested to get access to this data, for instance via bribery or blackmail, or even infiltration. Why am I exposed to such security risks and privacy violation with the introduction of smart meters?

Information is power. This is a basic law in politics, in marketing, and in many other sectors (like the military). The more electricity companies know about their customers, the better they can influence and steer their behaviour.

Is this far-fetched? Based on one's long-term detailed electricity consumption, the electricity provider can observe certain (statistical) patterns, like when the fridge switches itself on and how much electricity it uses. The provider can even observe if such a fridge becomes old (less efficient) and could be replaced. This

---

<sup>1</sup> at least in the Netherlands; the situation in other countries may vary, but the differences are not essential at this stage.

information can be used for targeted (fridge) advertisements. Consumers may view this positively as a form of service or negatively as a form of intrusion. Also, one can of course be skeptical how much of a service this really is: the electricity company will probably have a contract with a specific fridge supplier and provide information selectively. These scenario's are not imaginary but exist within the electricity supply community.

The (information) architecture determines how information flows within a particular IT-system: who can see what about whom. Since knowledge about others gives a stronger position and thus more power, IT-architecture is a highly political matter. After all, centralised informational control supports centralised societal control. Therefore, such architectural power issues are best discussed and decided upon within political fora, such as parliaments. Unfortunately the level of awareness of the sensitivity of these matters is fairly limited. In practice these issues are decided in different places, such as enterprise board rooms, or within ministeries, via civil servants, consultants and lobbyists. Those in control tend not to opt for architectures that reduce their power. But who defends the interests of individual citizens, and tries to protect their autonomy?

An architectural issue that comes up time and again is centralisation versus decentralisation (see also later in this paper). The architecture that often first comes to mind is a centralised one. It is the obvious one. Traditional privacy-enhancing technologies (PETs) focus on this approach. They use relatively low-level cryptography and rely to a large extent on organisational measures. Over the last years we are seeing that data-leakage and privacy incidents are becoming structural<sup>2</sup>. The centralised approach concentrates power and its protection model does not work well.

The decentralised approach, with data storage under direct control and influence of individuals, is typically not so well-developed. It often requires more high-tech crypto, such as zero-knowledge proofs or secure multi-party computation. Also, identities are less prominent, and can often be replaced by (anonymous) attributes, so that the possibilities for tracing of individuals and identity theft are reduced.

Whether our ICT infrastructures will be organised in a centralised or in a decentralised manner will have a deep impact on the organisation of our society and in particular on the division of (political) power. It is a sign of the level of appreciation of individual autonomy.

## 2 What went wrong: smart cards in public transport

This section will sketch the story of the bumpy introduction of a national smart card—called *OV-chipkaart*, or *OV-chip*, for short—in public transport in the Netherlands. It will focus on the architecture and privacy issues involved.

Almost a decade ago planning started in the Netherlands for the nation-wide introduction of a smart card based payment system for all forms of public transport: bus, train, metro, *etc.* Organisationally this is a non-trivial matter, because

<sup>2</sup> See for instance <http://datalossdb.org/> for an overview.

there are so many public transport companies, public authorities (at various levels, local, regional, national) and stakeholders involved. Various motivations for such a card exist: fair division of revenues and/or subsidies and improvement of service via detailed travel logs, public safety via restricted access (via electronic gates), fraud reduction, cost reduction (fewer inspectors needed), convenience for travelers, behavioural targeting and direct marketing via personal travel profiles, and simply the desire to look modern and high-tech.

Technically the matter became a nightmare. The RFID chip that was selected for the *OV-chip* was the Mifare Classic, because it was cheap and “field-proven” technology. It is indeed the most widely used RFID chip, of which around 1 billion copies have been sold worldwide, with around 200 million still in use, largely for access control (buildings and premises) and public transport (about 150 cities worldwide, including London, Boston and Beijing). The design of the Mifare Classic is from the early nineties. Its security depends on a proprietary authentication protocol and stream cipher using keys of only 48 bits long. In the beginning of 2008 it became clear that the card was broken [10, 9, 6, 3, 4] and that its content (esp. its balance) can be accessed and changed within a matter of seconds. Hence the chip is not a data-vault, as it should be.

After a phase of denying, dismissing and trivialising these findings the main players started accepting them and began working on a replacement plan towards a new card. In the meantime the actual roll-out went ahead. The expectation (and hope) of the public transport companies is that fraud can be detected quickly in the back-office, leading to black-listing of fraudulent cards. These fraud detection mechanisms run overnight on the transactions of the previous day. Hence they are not perfect. For instance, buying a €5 prepaid card and manipulating its balance to let’s say €100 enables you to travel anywhere (first class!) for a day. The card will be black-listed the next day, but by then you are ready to manipulate another card. The intention is to start migrating to a new card as soon as fraud levels reach a certain (unknown) threshold.

Apart from these technical difficulties, the *OV-chip* is a privacy disaster, at various levels.

- (1). Each Mifare Classic chip has a unique identity in the form of a UID. It is an identification number of 32 bits that is used in the so-called anti-collision phase when a card and a reader start to set up a conversation. The numbers’ main purpose is to keep different cards apart.

Hence the card tells this unique number immediately, in unencrypted form, every time it enters the electromagnetic field of an arbitrary card reader. It can be used to recognise and trace people. Soon, almost everyone in the Netherlands will be carrying such an *OV-chipkaart* and will thus be electronically recognisable via the UID of the card. Parties not affiliated with the *OV-chipkaart* may use these UIDs for whatever they like. Shopkeepers may start using this UID: “hey, hello, long-time-no-see; maybe it is time you buy new underwear!?” Since card readers can be bought for around €10 everyone can set up its own card detection and tracing, for benevolent or malevolent (*e.g.* terrorist) purposes.

More modern smart cards, such as used in the Dutch e-passport, use a random UID. This is good enough for anti-collision, but prevents tracing.

- (2). As said, the cryptographic protection mechanism of the Mifare Classic is broken and the contents of the card can be accessed with relatively easy means—certainly when the technique becomes available for script kiddies. Hence someone sitting next to you in the bus or in a cafe can in principle read the contents of your card. It contains the logs of your last ten trips so that it is publicly visible (in principle) where you have been (ah, . . . , the red-light district!). If the card is personalised, it also contains your date of birth, but not your name or address.

In principle someone who is reading your card in this manner can also surreptitiously change the data on the card. In this way your balance or validity date or date of birth can be changed, without you noticing. Of course you can also do this yourself. The card is an open wallet, that does not protect the money you store on it.

European data protection acts generally put the obligation on data processors to appropriately protect privacy-sensitive data in their custody. It can be argued in this case that Mifare Classic based travel cards are not in accordance with the law. They have not been challenged in court so far.

- (3). Each entry or exit into the public transport system, in a bus or at a train station, generates an entry in the back-office of the travel companies<sup>3</sup>, involving among others the identity of the entry/exit point (often connected to a fixed location), time-of-day, and identity and balance of the card. This yields a huge database of travel transactions that can often be linked to specific clients, for personalised cards. Hence individual travel patterns can be determined easily. The travel companies have left no doubt that they are quite eager to do so and to use these data for behavioural targeting and direct marketing: if I often travel to Rotterdam I will get advertisement for certain shops or hotels in Rotterdam; and if I travel first class I will get advertisement for different shops and hotels. Travel companies have been fighting for years with the Dutch data protection officer over this matter, and have finally managed to reach a half-baked compromise with an opt-in for such direct marketing and an opt-out for more aggregated profiling for optimised travel advice—in order to stimulate the use of public transport<sup>4</sup>. The former East-German *Stasi* would have been jealous of such a database. It will not take long before police and intelligence services start exploiting its potential. One may expect “public transport taps”, analogous to phone taps, data retention obligations<sup>5</sup> and datamining. Reportedly, this is already happening with London’s Oyster card.

- (4). Traditionally, when you enter a bus or a train you only have to show, when asked, that you possess an appropriate “attribute” (ticket) that permits you

---

<sup>3</sup> which, for simplicity, are treated here as one organisation.

<sup>4</sup> It seems to be a new trend to use ecology, instead of terrorism, as argument to undermine privacy.

<sup>5</sup> which are quite unnecessary actually, because the travel companies show no inclination whatsoever to delete any transaction data.

to travel in this way. Such an attribute is typically not personalised. With chip card based systems access to the public transport has suddenly changed from attribute-based to identity-based. The chip card tells its identity to the card reader upon entry (and exit). Is it really necessary to tell who you are when you enter a bus? Do we want such a society?

There are modern, more advanced cryptographic protocols (see *e.g.* [1, 11]) that allow anonymous, attribute-based access control. However, they require non-trivial computational resources and advanced, very fast smart cards. Getting them to work in practice is a topic of ongoing (applied) research.

- (5). Anonymous, prepaid cards do exist, but they are a sad joke. As soon as you charge them with your bank card, your anonymity is gone. They are even sold via the internet! They clearly show that privacy is the last thing the designers of the *OV-chip* system cared about—in sharp contrast with the principle of privacy by design [8].

With a migration to a new card the points (1) and (2) will hopefully be tackled, via a random UID and proper cryptographic protection of the card contents. The centralised, card-identity based architecture will then remain. Changing this architecture as well involves moving to attribute-based access control, which is not foreseen for the near future.

As described, the *OV-chip* smart card based ticketing system involves a centralised architecture giving the travel companies unprecedented access to individual travel behaviour. This aspect is not emphasised, to use an understatement, in the advertisements of these companies (or of the government). There is little transparency. Politicians and the general public are only slowly (and partly) becoming aware of these architectural matters. Reactions vary, from complete indifference to strong indignation. Independent regulators, such as the Dutch Data Protection Authority (CBP) or the Dutch Central Bank (DNB), are reluctant to take decisive action. The invested economical and political interest and prestige are high.

What is the rôle of academics when such technically and societally controversial infrastructure is introduced? The Digital Security Group of the Radboud University Nijmegen had to deal with this question because of its research into Mifare Classic chipcards [6, 3, 4]. It has decided to speak openly about the security vulnerabilities, which in the end involved standing up to legal intimidation and defending its right to publish freely in court, see [2] for a brief account. The group decided to draw a line and stay away from political activism, but not to stay away from academic *engagement*, in the sense of speaking openly about the issues and giving a number of demonstrations, for instance during an expert session in national parliament where an *OV-chipkaart* of a volunteer member of parliament was cracked on the spot and read-out in public, and in raising the balance of a card of a journalist with one cent in order to see if this would be detected in the back-office<sup>6</sup>. In the end it is of course up to others to decide how

---

<sup>6</sup> The responsible company *Translink* claimed (afterwards) it had noticed the event as a mismatch of balances but had decided not to respond because one cent did not

to evaluate the warnings and to make their risk analysis. The underlying motivation was to make clear that the *OV-chip* system should not set the information security standard for large scale future (public) ICT-projects. It is probably too early to evaluate the developments but it seems that this approach may have had indirect influence on other projects but has had little direct impact on the actual roll-out of the *OV-chipkaart*<sup>7</sup>. In The Netherlands politicians and industrials have become aware of the fact that large ICT-projects can be made or broken by security issues. This applies for instance to electronic patient dossiers, electricity meters (see the previous chapter), or road pricing—which is the next topic.

### 3 What can still go right: road pricing

There is another major ongoing ICT-project in the transport sector in the Netherlands, namely road pricing (also called electronic traffic pricing). It is however still in the design phase, with prototypes being developed. The idea is to equip every car registered in the Netherlands with a special box containing GPS and GSM functionality for determining its detailed whereabouts and for communication with the central back-office of the traffic pricing authority. The aim is to replace the existing flat road tax by a fair, usage-dependent road charge that allows congestion reduction—by making busy road segments expensive—and environmental impact reduction—by making charges depend on vehicle characteristics. The choice for time, location and vehicle category dependent kilometre tariffs makes this approach ambitious and new in the world. For each individual vehicle detailed time and location information must be collected and processed without endangering privacy. The correct amounts have to be calculated reliably with the help of a digital tariff and/or road map.

The main architectural choice in road pricing is, as before between a centralised and a decentralised system.

- In a centralised architecture the on board box is simple and all intelligence resides with the pricing authority. The box frequently sends, say every minute, its location data to the authority. At the end of each period, say each quarter year, the authority calculates the total fee due.

This architecture is simple and rather naive. Privacy violation is the main disadvantage: it means that the authorities have detailed travel information about every vehicle. This is unacceptable to many. An additional disadvantage is the vulnerability of this approach. The central database of location data will be an attractive target for many individuals or organisations with

---

look like serious fraud. Hence the journalist could travel for a couple of days with the adapted card, until the manipulation was made public. Not everyone was amused. A similarly manipulated (Oyster) card was used earlier in London to demonstrate that free rides are possible.

<sup>7</sup> It did result in some additional delays, and in the development of a migration plan.

unfriendly intentions, like terrorists or blackmailers. The system administrators who control this database may not always behave according to the rules, voluntarily or unvoluntarily.

In this approach one needs to have confidence that the box in each car registers and transfers all actual road use correctly. This may be enforced by “spot checks” along the road, where vehicles are photographed on specific (but random) locations and times. These observations can then be compared with the transferred registrations. A fine can be imposed in case of discrepancy.

- In the decentralised architecture the on board box contains enough intelligence to calculate the fee itself, requiring additional local computing resources. The box must securely store data, contain tariff & road maps, and carry out tariff calculations. Additionally, in such a decentralised approach the road-side checks must involve two-way communication in order to be able to check that the last few registrations and associated fee calculations have been performed correctly. Such checks require requests and responses by the vehicle, and thus an additional communication channel.

This decentralised approach apparently also has disadvantages, mainly related to the required local complexity, making it more vulnerable. It does however leave privacy-sensitive location information in the hands of individuals, who may access it via direct (USB or Wifi) connection with their box—and plot it for instance on their own PC (for fun) or use it for reimbursement.

A privacy-friendly and fraud-resistant alternative is sketched in [5] where the on board box regularly—for instance, once per day—sends hashes of its time-stamped location data to the pricing authority, in order to commit itself to certain trajectories without revealing any actual information. After a random photo spot-check the driver (or actually, the box) is asked to send in the pre-image of the submitted hash in order to show correctness of the submission. Fee submission can also be organised in such a way that checks can be performed “locally”, so that the privacy violation associated with checks is limited. Further details can be found in [5].

The underlying point is that privacy-friendly architectures are often available, but may require a bit more thinking (and work). It is a political issue whether or not we, as a society, wish to use them (see also [8]).

## 4 Privacy and trust for business

This section addresses the question: is there any honest business in privacy protection? An actual proof of commercial viability can only be given in practice, but a scenario will be sketched, taking economic incentives (instead of regulation) as starting point.

The first observation is that (commercial) services are quickly becoming both (a) electronic, and (b) personalised. It is especially in this personalisation that future business opportunities are expected, with targeted, on demand delivery, based on personal preferences. Hence it is important for businesses to become



*trusted*, so that customers allow these businesses to know them better and to deliver specifically what they really want.

So how to obtain trust? Keeping users ignorant about business processes and information flows seems, certainly in the long run, not to be the right strategy. Hence transparency is needed. Standardly one uses *a priori* mechanisms such as certificates, or *a posteriori* mechanisms based on reputation to obtain trust. Here we propose an approach based on three points:

- build your infrastructure (architecture) so that a privacy freak would be happy;
- allow customers to switch various transparent options on or off, for additional services;
- do not ever try to make any money out the privacy freaks, who typically have every possible additional option switched off.

Such mode of operation may be covered by some sort of privacy seal (see *e.g.* EuroPriSe).

The rôle of the privacy freak in this model is to generate trust. Given enough trust, you, as a business, may hope that a large majority of your customers will actually switch on additional services, enabling you to do (more) business. If you ever dare to act against what you promised towards a privacy freak, this will have detrimental effect on the level of trust and thus on your reputation as a business. Hence these freaks form a necessary foundation for your business, which will not directly generate revenue, but only indirectly via the customers with enough trust and willingness to obtain services via your infrastructure.

Here is a more concrete version. In many countries postal services want to become “trusted parties”. One way of starting such a service is starting a web portal offering “throw-away connectivity”. This can be done for (mobile) phones or email addresses<sup>8</sup>, but here we shall describe ordinary, physical addresses.

Many people shop online, but online shops have a bad reputation in handling customer data. They frequently leak credit card details or other sensitive customer data. Hence my trusted (postal) web portal could offer to act as a (transparent) intermediary when I purchase an item online. After making my selection on a (redirected) webpage, the portal could handle the payment on my behalf—so that my details remain hidden and cannot be abused by the shop owner—and provide the shop with a throw-away address in the form of one-time barcode that is printed and put on the box containing the item that I purchased. The box is put in the mail by the shop owner and ends up in one of the distribution centers of the postal company. There, the barcode can be read and be replaced by my actual physical (home) address, so that the box can be delivered to me<sup>9</sup>.

Personally, I would be strongly interested in such a service, even if it involves registration with one particular (trusted) company that can see everything that

---

<sup>8</sup> like at [spamgourmet.com](http://spamgourmet.com).

<sup>9</sup> Such an address-blinding service is not something only postal services can do; anyone can do it by acting as distribution center oneself. However, this may involve double postal charges.

I do. Even more if such a company has a privacy seal and offers various options with respect to the level of logging that they carry out, either at a central level or at a decentralised level with logs (and personal preferences) stored under my own control, but signed by the portal for authenticity/integrity. Maybe I do want some limited level of logging for dispute resolution (for certain purchases only). Maybe the portal even has switches for (temporary!) advertisements for specific goods, whenever I need them (like a new car, or fridge). Long term trust requires the portal to really stick to its policy of not sending me anything else or bothering me with things I didn't ask for.

Such a portal can develop into a more general identity provider (or broker), using mechanisms such as OpenId. Who makes the first move, and wants to transparently use modern technology for personalised use without hidden control and manipulation agendas?

**Acknowledgements** Thanks to Wouter Teepe for his critical reading of an earlier version.

## References

1. S. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, 2000. Freely available via [www.credentica.com](http://www.credentica.com).
2. A. Cho. University hackers test the right to expose security concerns. *Nature*, 322(5906):1322–1323, 2008. Available from <http://www.sciencemag.org/cgi/reprint/322/5906/1322.pdf>.
3. F. Garcia, G. de Koning Gans, R. Muijers, P. van Rossum, R. Verdult, R. Wichers Schreur, and B. Jacobs. Dismantling MIFARE Classic. In S. Jajodia and J. Lopez, editors, *Computer Security – ESORICS 2008*, number 5283 in Lect. Notes Comp. Sci., pages 97–114. Springer, Berlin, 2008.
4. F. Garcia, P. van Rossum, R. Verdult, and R. Wichers Schreur. Wirelessly pick-pocketing a mifare classic card. In *30th IEEE Symposium on Security and Privacy (S&P 2009)*. IEEE, 2009, to appear.
5. W. de Jonge and B. Jacobs. Privacy-friendly electronic traffic pricing via commits. In P. Degano, J. Guttman, and F. Martinelli, editors, *Formal Aspects in Security and Trust*, number 5491 in Lect. Notes Comp. Sci., pages 143–161. Springer, Berlin, 2009.
6. G. de Koning Gans, J.-H. Hoepman, and F. Garcia. A practical attack on the MIFARE classic. In G. Grimaud and F.-X. Standaert, editors, *8th Smart Card Research and Advanced Application Conference (CARDIS 2008)*, number 5189 in Lect. Notes Comp. Sci., pages 267–282. Springer, Berlin, 2008.
7. L. Lessig. *The Future of Ideas*. Vintage, 2001.
8. D. Le Métayer. Privacy by design: a matter of choice, 2009. This volume.
9. K. Nohl, D. Evans, Starbug, and H. Plötz. Reverse-engineering a cryptographic RFID tag. In *17th USENIX Security Symposium*, pages 185–194, San Jose, CA, USA, 2008.
10. K. Nohl and H. Plötz. Mifare, little security, despite obscurity, dec. 2007. Presentation on the 24th Congress of the Chaos Computer Club in Berlin, see <http://events.ccc.de/congress/2007>.
11. E. Verheul. Self-blindable credential certificates from the Weil pairing. In C. Boyd, editor, *Advances in Cryptology - ASIACRYPT 2001*, number 2248 in Lect. Notes Comp. Sci., pages 533–550. Springer, Berlin, 2001.