The Smart Tachograph – Individual Accounting of Traffic Costs and its Implications

Vlad Coroama

ETH Zurich, Institute for Pervasive Computing
8092 Zurich, Switzerland
`coroama@inf.ethz.ch`

Abstract. Today, several costs caused by road traffic may either be only roughly approximated, or cannot be clearly assigned to the drivers causing them, or both. They are typically distributed evenly among a large fraction of drivers, which is both unfair and economically inefficient. We have built a prototypical platform, called the "Smart Tachograph", that allows us to measure traffic-related costs on an individual basis, thus supporting a more fine-granular charging of the responsible parties. Sensors observe the manner and circumstances in which a vehicle is driven, while several accounting authorities can evaluate this information and charge motorists on a pay-per-use basis. The Smart Tachograph offers valuable insights for the deployment of future ubiquitous computing services in general: its implementation has obvious requirements in terms of security and privacy; its deployment model is realistic through the strong economic incentives it offers; and its usage directly affects core societal values such as fairness and trust. This paper summarizes our design considerations and discusses the feasibility and wider economic and societal implications of fielding such a system.

1 Introduction

The *Smart Tachograph* is a system that allows an individual, fine-granular analysis and accounting of traffic costs. Nowadays, many of the traffic-related costs are not accounted to their originators, but rather spread across a larger group, mainly due to the impossibility of exact measurements. Ecology-oriented vehicle taxes, for example, typically depend on the vehicle's type, more polluting vehicles having to pay a higher tax. Such taxes fall short of fulfilling their ecological aim, however, since they do not take into account the annual mileage of the vehicle, nor the conditions in which the car is being driven, like ozone levels.

Likewise, today's car insurance schemes typically divide drivers into about two dozen different risk categories, using only a few criteria such as the driver's age, gender, driving experience, place of residence, or car model. While all these parameters are being determined before the insurance goes into effect, the actual behavior of the driver after signing the policy (e.g., a safe driving style) will reflect only slowly on his or her insurance rate, typically over many years. Young people will pay high insurance fees only because, on average, young drivers tend to drive more aggressive and accident-prone. The individual young motorist often has no means of proving himself or herself to be a safe driver, other than several years of accident-free driving.

Through the use of ubiquitous computing technology, however, much of the data that has been previously unavailable might now easily be measured. According to the place, time, and manner someone is driving, the economical and ecological costs, as well as the risk of being involved in a traffic accident can be estimated with a high degree of accuracy. The Smart Tachograph is a prototypical system designed to allow a determination of these momentarily costs and risks and subsequently bill drivers in a pay-per-use/pay-per-risk manner. Its aim is to offer valuable insights for the development of future ubiquitous computing services in general by providing a realistic model for analyzing the technical, economical, and societal challenges such applications will pose.

The remainder of the paper is organized as follows: Section 2 presents in larger detail the problems that arise due to the impossibility to measure many traffic costs. Section 3 gives a detailed view of the system. Section 4 presents related work. Section 5 addresses several issues raised by the deployment of a system such as the Smart Tachograph: its economic feasibility and practicability, its privacy implications, its broader societal implications, as well as the influence of system design decisions on these dimensions.

2 Motivation

2.1 Information Asymmetry in Insurance Markets

Various authors argue that the nowadays practiced classification of automobilists into a few classes – typically based on their driving experience, accident history and type of driven car – is not optimal. According to [10], within such a class (of presumably similarly skilled drivers), there is still a large spread of risks, depending on such factors as: the annually driven mileage; the time of day and the season predominantly driven at; weather conditions; the type of route (a certain distance in a crowded city being more accident-prone than the same distance on a highway) or the neighborhood where the car is usually parked. Litman [9] also argues that today's rigid insurance premiums are both economically and ecologically obsolete: "What would be the consequences if gasoline were sold like vehicle insurance? With gasoline sold by the car-year, vehicle owners would make one annual advance payment which allows them to draw gasoline unrestricted at a company's fuel stations. Prices would be based on the average cost of supplying gasoline to similar motorists. Unmetered fuel would cause a spiral of increased fuel consumption, mileage, and overall vehicle costs, including externalities such as accident risk, congestion and pollution." Instead, the insurance should be related to the mileage driven because, all other parameters equal, there is a strong correlation between driven distance and accident risk [9]. Connecting the insurance rate to the annual mileage would be fairer and economically more sensible. Moreover, since a larger fraction of vehicle costs would depend on the driven distance, it would also have a positive environmental side-effect. Oberholzer goes further and builds up a detailed matrix of how much the insurance kilometer should cost depending on two factors: type of road (highway vs country road vs city) and the hour of driving [10].

All these distinctions not being done today, two phenomenons occur. First, inside one of the risk classes of presumably equally-skilled drivers, the ones with a higher annual accident risk (due for example to a higher annual mileage) are being cross-financed by lower risk motorists. Second, as mentioned above, all the parameters determining the insurance rate are measured before the insurance goes into effect. A safe driving style or other safe behavior (like parking only in secure areas) will not be reflected immediately on the insurance rate, but rather slowly. Neither will a low-mileage driver be rewarded with an insurance bonus. The other way around, exhibiting dangerous driving behavior will not influence the insurance rate in any way except when the driver gets involved in an accident. The aggressive driver lacks a direct feedback on how his or her driving style increases the risk of being involved in a traffic accident.

Both these problems are well-known in insurance markets. The cross-financing from low to high risks is called *adverse selection*, while *moral hazard* denotes the tendency to handle an insured good more carelessly after it has been insured. They both ultimately root in the same phenomenon of *information asymmetry*, first described in an influential article by George Akerlof [1]. Information asymmetry denotes the state in which one market side has more information than the other side. In the context of insurances, it describes the insurers' lack of information about the actual behavior of their customers and thus the exact dimension of the risks they insure. Because of this lack of information, the insurer can not reward customers that have a low-risk behavior. Instead, he has to insure an average risk through a larger customer group. This is not only unfair towards the ones having less risk, it is also economically inefficient, since it hinders a market for "high-valued goods" (low risk drivers) to emerge [1].

2.2 Road Pricing

Road Pricing is a tool for regulating the traffic flow through selectively penalizing the driving on specific roads at particular times or under specific conditions. Deploying a road pricing scheme may have several political or societal aims. For overcrowded city centers, it may be deployed for replacing the regulation of traffic through queuing (the "communist" solution) by a free-market mechanism. It may also be used to steer the traffic away from some streets to others (by penalizing the former more) or to other means of transport. Road pricing may further pursue environmental aims, like reducing emissions or noise levels. Finally, it may simply be used to raise money for the maintenance of the road infrastructure. However, as [6] argues, whatever the main reason – financing, improving the environment, or managing traffic and improving accessibility – a road pricing system will have all these effects to a certain extent.

Road pricing has become increasingly popular over the last years, since the two traditional tools for charging drivers, fuel and vehicle tax, are rather coarse and cannot fulfill all of the above mentioned aims. The annual vehicle tax penalizes people for owning dirty cars, but this says nothing about the actual pollution caused by those cars. Since the car's overall pollution is the product of its emissions per distance unit and the car's usage, this flat tax fails short of fulfilling its environmental aim. Fuel tax is better at penalizing people for the consumption of gasoline, but does not look at the other side – how dirty the emissions resulted from that consumption really are or under which circumstances the gasoline has been burned (e.g., ozone levels). Moreover, neither tax can have a traffic management effect [14].

Hence, many places worldwide have started to deploy road pricing systems as a complementary tool to the existing taxes. There is a wide spread in the level of detail that existing road pricing systems take into account. At one end are rather coarse systems, such as the London Congestion Charge. There, motorists are charged once every 24 hours for the permission to drive in the city center [6], regardless of the actual usage during these 24 hours. Even so, the introduction of the Congestion Charge has reduced traffic by 15% and increased speed by 22% in central London [14]. At the other extreme, having a much more detailed usage model, stands the ERPS (Electronic Road Pricing Scheme) from Singapore. There are different taxes for the usage of distinct roads, and they also vary with the hour of driving. Furthermore, every three months, the whole price structure is analyzed and readjusted [6].

Many other cities play with the idea of road pricing systems: Viennese officials think of a city-wide street usage charge of 2 to 8 Euro-Cents per kilometer [12]. A large study has been carried out by the Swiss Center for Technology Assessment to investigate the public's acceptance of a generalized road pricing scheme, envisioned by parts of the government [11]. Britain also thinks of a nationwide, satellite-based road pricing system, no earlier than 2014 though [13].

3 The System

To analyze the implementation requirements and the subsequent economic and societal consequences of such traffic-related cost allocation issues, we have proposed the Smart Tachograph generic platform [4]. The Smart Tachograph uses off-the-shelf ubiquitous computing technologies and a newly developed prototypical software infrastructure to allow for measurement of driving parameters, the transformation of those parameters into costs, and billing these costs to their originators (Fig. 1). The software infrastructure of the prototype runs on a laptop computer that can be placed anywhere in the car. Any number of sensors can be attached to the system. They are depicted above the car in Figure 1. The sensors gather data about the way and the circumstances in which the vehicle is being driven and send this information to the computer. Several accounting authorities (connected from below in the figure) may evaluate this information and charge motorists on a pay-per-use basis. The software platform serves not only as a sink for sensor data and as back-end connection for the accounting entities, but also (not shown in Figure 1) as a front-end interface to the vehicle's driver.

3.1 Sensors

A small plastic box (Fig. 2) has been fitted for our prototype with a collection of sensors. It contains a GPS unit and a sensor board carrying two accelerometers (for longitudinal and cross acceleration), a temperature sensor, and a light sensor. Raw GPS coordinates are not the only information that can be obtained from the GPS unit. The current time is also encoded in the satellite signal and the current speed can be inferred as distance traveled over time. The data gathered by all these sensors is sent via Bluetooth to the computer running the Smart Tachograph software infrastructure. We used a Bluetooth-enabled GPS sensor, while the sensor board sends its data through

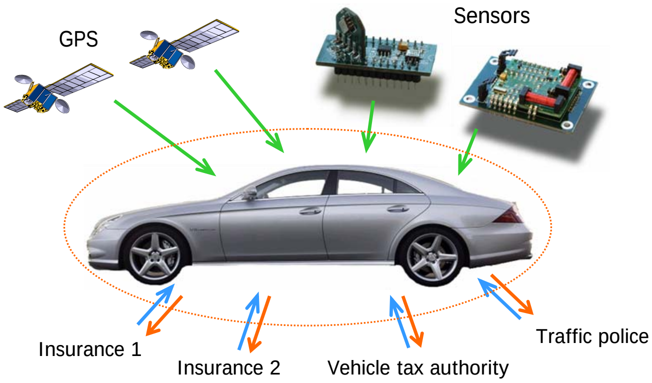

Fig. 1. Top-level view of the Smart Tachograph with the main involved instances.

a BTnode. The BTnode[1] is a small computing device for sensor network applications equipped with Bluetooth communication capabilities.

"Installing" the system is pretty straightforward. The box has to be placed in a spot where the GPS sensor can easily receive the satellite signals, for example underneath the car's windshield. The only other point to be ensured is that the sensor box is placed on an even surface and that it faces in the correct direction. Both conditions are needed for a correct functioning of the accelerometers. The controlling computer can be placed anywhere in the car, since it wirelessly communicates with the sensors.

We have chosen to deploy the mentioned type of sensors for two reasons. First, they all measure data that is potentially relevant to one of the envisioned accounting authorities. Having the raw GPS coordinates, the system can always determine on which street the vehicle is on and the speed limit for that street, using a commercial geospatial database installed on the computer. This information is obviously relevant for a road pricing scheme. But it could also influence the current accident risk, for example if the driven speed significantly exceeds the speed limit. Excessive longitudinal acceleration and especially excessive cross acceleration seem to be equally important indicators for a high accident risk. Finally, the light intensity sensor gives information about light conditions, which could also influence the accident risk. The second reason to include these sensors in the Smart Tachograph prototype was that all of them are already available in a typical medium to high end state-of-the-art car. They have been placed there for other reasons, but could be reused in a real deployment for a system such as the Smart Tachograph. A modern car is equipped with acceleration sensors for the electronic stability programs, with a temperature sensor for signaling the driver a possible slippery road, and with light sensors for automatic headlight activation. Most cars today also come with a GPS system and navigation maps.

Apart from the sensors used in our example, many modern cars come with a variety of other sensors that could be used to determine the insurance rate or road pricing tax

---

[1] See www.btnode.ethz.ch.

even more accurately. A distance sensor used in many cars as parking aid could be reused to measure the distance to the car in front. This information, correlated with the type of street and driven speed, is a major determinant for the current risks taken while driving (see section 5.4). If the car is connected to the Internet (e.g., via UMTS), it could also download environmental data that possibly determine its road pricing tax, like ozone levels or the concentration of carbon dioxide in a specific city.

Why haven't we used the car's sensors if they are already there and probably more precise than ours? The practical reason was that we did not have access to the vehicle data bus, since the work has been academic research so far. The aim of the work lies not in the highest possible sensor precision, or in the most realistic approximation, but in creating a proof of concept for a generic traffic accounting platform and to analyze the various implications of deploying such a system. The fact that similar sensors already exist in vehicles only underlines the feasibility of the presented concept.

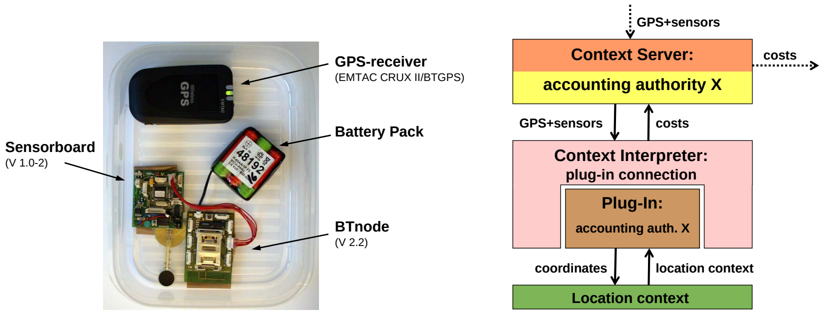

Fig. 2. The plastic box containing the Smart Tachograph's sensors (left) and the plug-in-architecture for the accounting authorities (right).

3.2 Software Infrastructure

The main role of the software infrastructure is to query data from the sensors, and to mediate communication with the accounting places. Due to the flexible software design, adding accounting entities is as easy a task as adding new sensors. At the time being, three different kind of accounting entities have been included in the system (see Figure 1): insurance companies, a vehicle tax authority, and the police. The traffic police has been included in order to show how powerful the paradigm of a smart tachograph is and what far-reaching social consequences it could have. These consequences are discussed in Section 5.

To connect with the accounting authorities, the Smart Tachograph uses a plug-in-architecture as depicted in Figure 2. The system is built on top of Anind Dey's Context Toolkit [5], with every accounting authority represented by a Context Toolkit *server*. This server is registered through the Context Toolkit's publish/subscribe mechanism to receive all events from the GPS unit and the other sensors. The server sends this information to a Context Toolkit *interpreter* that generates the corresponding costs according to the rules defined by a plug-in that has to be loaded when starting the system. The costs are then “consumed” by the accounting authorities, registered in the system as *context handlers*. The plug-ins define how the telemetry data are transformed into costs: they infer the road fee or calculate the accident risk and transform this risk to an insurance rate, for example. To be able to transform the raw GPS coordinates into meaningful location context, the plug-ins also have direct access to the commercial geospatial database installed on the computer. Every predefined period of time (a week or a month would probably be meaningful), the interpreter returns an aggregated sum to the context server, which in turn sends this sum to the accounting authority (over the vehicle’s UMTS connectivity or the home WiFi network that can be received from the garage, for example). The fact that the accounting authority does only receive the accumulated sum but not the raw sensor data will be relevant for the privacy and security discussion in section 5.2.

By using the Context Toolkit, adding new sensors or new accounting instances become easy tasks. To add a new sensor, a new *widget* has to be written that encapsulates the proprietary communication with that sensor. Similarly, a new accounting authority is being added by registering it as new *handler*.

From a functional perspective, the Smart Tachograph knows three kind of predefined accounting entities: `compulsory`, `required-select`, and `optional`. In a system configuration file, three corresponding lists have to be filled out. In a realistic setting, these could be for example:

```
compulsory: vehicle-tax-authority; traffic-police
required-select: liability-insurance
optional:own-damage-claim
```

Every item in the “compulsory” list has to be active before the system can be started. An accounting entity is active when its server has been registered in the system as a subscriber for sensor values, and the plug-in has been downloaded from the corresponding authority server. Likewise, for every item in the “required-select” list, one plug-in has to be present. The difference is that the user may choose here between different plug-ins (e.g., from different insurance companies), while for compulsory plug-ins there is no choice. Any number of “optional” plug-ins can be loaded before starting the system, but none is required. In the prototype presented here, the Smart Tachograph software does not start until all mandatory servers have been started. In a real deployment, it is conceivable that the system would be connected to the electronic anti-theft device, so that the car would not start until all the legally required plug-ins are active. Also, in a real deployment, the tasks accomplished for the prototype by the laptop computer would probably be taken over a by a computing on-board-unit.

3.3 Driver Interface

The Smart Tachograph’s software further includes a front-end interface to the driver (Fig. 3). The main system window, in the lower right corner, is needed to setup and start the prototype system. Several parameters can be set here. Among them, the driver may choose for all required but selectable items a specific instance. For example, he or she may choose a liability insurance from the existing offers. The driver may further choose any number of optional insurances. To have an up-to-date view of the existing offers, when the user starts the system (read: “enters the car” for a real deployment), the Smart Tachograph accomplishes the following two steps. First, it connects via the car’s wide area communication system to a root-server, retrieving a list of available vehicle insurance companies. Then it connects to the server of each insurance company to retrieve the available insurance schemes. After choosing everything needed, the user starts the Smart Tachograph by clicking the “start” button.

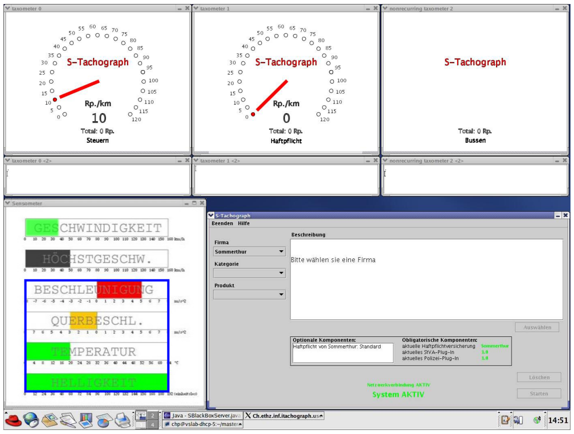

Firma: Sommerthur

Bitte wählen sie eine Firma

Haftpflicht von Sommerthur: Standard

aktuelle Haftpflichtversicherung Sommerthur
aktuelles StVA-Plug-In 1.0
aktuelles Polizei-Plug-In 1.0

Netzwerkverbindung AKTIV

System AKTIV

Fig. 3. A typical screenshot of the system’s interface.

The second window (in the lower left corner) presents the sensor data as a collection of bars. The raw GPS coordinates are translated into the actual street that’s been driven on using the geospatial database. Knowing the speed limit for all streets in the administrative region of Zurich, the system displays this information on the second bar from top. It uses this information, together with the actual speed (displayed on the topmost bar), for risk approximation. The lower bars show the data from the other sensors – longitudinal and cross acceleration, temperature, and light intensity. The data from all these sensors, as well as the GPS coordinates, are ascertained and transferred to the computer once every second.

The third type of interface windows (upper part of Figure 3) would probably be the only ones shown in a real deployment while driving. They show the current costs (insurance rate, road tax), which are continuously calculated from the received sensor data. The indicators presenting these aggregations should be perceived by the driver similar to the momentary gas consumption indication built into some cars. It allows the driver to receive instant feedback on how his or her driving habits influence the traffic costs. Traffic fines (window in the upper right corner) can also be issued automatically. They are not expressed as money per kilometer, but as one-time events (i.e., when the speed limit has been exceeded for more than ten seconds in a row).

4 Related Work

Within the ubiquitous and pervasive computing community, there has been little research regarding road pricing or pay-per-risk insurances. There have been, of course, numerous location-awareness projects in the broader domain of traffic, most prominent CoolTown's WebBus [8], but none that we are aware of regarding specifically road pricing or dynamic insurances. Our work has largely benefitted from context awareness research. The Smart Tachograph platform is built on top of Dey's Context Toolkit [5]. Some economists and business analysts [7, 10], aware of the potential of ubiquitous computing technology, have examined the impact of ubiquitous and pervasive computing on insurance markets from an analytical point of view, specifically highlighting the vehicle insurance market. Other related work can be found in publications on road pricing or vehicle insurances.

A good overview on road pricing systems in use today can be found in [6]: the London Congestion Charge, Singapore's ERPS, the systems in Oslo and Trondheim, and California's expressway SR91. All systems but the one in Singapore have a very coarse area model. They either penalize the usage of one road only (the California expressway), or the entrance to a specific area, typically the city center (London, Trondheim, Oslo). In all these examples, the fee is flat for a multi-hour usage permit, typically for 24 hours. Obviously, this only allows a coarse traffic management, keeping some traffic out of the surcharge area, but having much of that traffic redistributed on the borders of the restricted area. Moreover, such a flat system does not acquire a high degree of fairness, since everybody pays the same independent of the actual usage. In Singapore, the model is more complex, every usage being charged and the fees varying from street to street and with the hour of driving. The more jam-prone the street and the less fluent the traffic rolls, the higher the tax gets. Singapore's road pricing has a very fine-granular model, however, it can not include other parameters, such as ozone or carbon dioxide levels. Another drawback of all these systems, as compared to the Smart Tachograph architecture, is that they rely on a heavy infrastructure of active RFID tags, ubiquitous gates with powerful antennas to read those tags, numerous cameras to identify the cheaters, and partially also payment machines and manned stations.

Early discussions in the field of pay-per-risk insurances revolved around considering driven mileage only. [9] made the case for including the driven mileage into the calculation of the insurance rate, also pointing out the positive environmental and traffic safety side-effects, but ignored other criteria. Progressive, a US-insurer, initiated a pilot project (called Autograph) between 1998 and 2000. It took into account the driven distance, and in addition the time of day and the geographic location.[2] More recently, Norwich Union, a UK-based insurer, offers a black box for what they call "pay-as-you-drive" insurance,[3] but disclose only that they take into account the hours of driving, not whether their black box also considers other attributes. Privacy does not seem to be a central issue: "The black box device measures vehicle usage and sends data directly to Norwich Union using similar technology to that used by mobile phones." Progressive

---

[2] See www.epa.gov/projectxl/progressive/index.htm.

[3] See www.norwichunion.com/pay-as-you-drive/.

also started last year to offer a more sophisticated insurance product, called TripSense.[4] TripSense is based on a black box that has to be installed in the car as well, but their webpage is more detailed about what it will record: "[...] which measures your actual driving habits and allows you to earn discounts on your insurance by showing us how much, how fast and what times of day you drive." The driver may analyze the data recorded over several months at his or her PC at home (and see, for example, the per-day number of "aggressive brakes" and of times driving over 75mph) and decide for himself or herself on sending the data to the insurance company or not. If the data is not sent, a no-punishment policy is advertised. This seems to be a more privacy- and customer-friendly approach than Norwich Union's, although in order to gain the price advantage, the customer has to send all data and thus give up privacy here, too. The more important point, however, is that the responsibility lies with the customer. He or she has to decide on sending the data to the insurer, without possibly realizing what the longer-term consequences of such action will be. What will happen at the next contract renewal if the driver has not sent any data to the insurer in the previous period of time? What if the data was sent and points to a risky driving style? Who else will gain access to the data and will "my" data ever be used against me? Progressive states that: "We may retain the information that you send to us indefinitely" and further that "If you are in an accident, you may have a legal obligation to preserve the information on the TripSensor. This information may be sought by opposing parties in a civil lawsuit or by police when investigating the cause of an accident. We may be legally obligated to provide such information in response to a subpoena or as otherwise required by law."

To sum up, neither of the above-mentioned systems have been built with customer privacy at its core. All presented road pricing infrastructures record the places where the vehicle has been (or at least where it entered and exited the fee area) through transponders, while the existing insurance models continuously gather data about the driver's habits and whereabouts and send them to the insurance company. With the Smart Tachograph, we try to provide evidence that highly personalized insurance rates are also feasible without such a massive loss of privacy and control.

Also, we know of no approaches so far that tried to develop an open platform that could be used to calculate and charge a great variety of traffic-related costs. Proprietary black boxes have been the standard solution so far for pay-per-risk insurance prototypes; transponders and a heavy gate infrastructure have been the standard for road pricing schemes; but the two concepts have not been combined before. Summarizing, we believe that the power of our approach lies in its simpleness, flexibility, and open character. Such power, however, can also be misused, as the next section will discuss.

5 Evaluation and Discussion

A technical evaluation of the Smart Tachograph platform would encompass answering questions such as how well the system estimates traffic costs, how robust these estimations are with respect to different driving styles, vehicle types, and environmental conditions, or if the chosen algorithms for calculating risks and costs (that will be presented later in this section) are appropriate and how they could be improved. We have

---

[4] See https://tripsense.progressive.com/home.aspx.

not conducted thorough research in any of these directions. Many issues would require a close collaboration with representatives from the insurance sector or traffic policy makers. We hope to be able to further investigate these directions in the future. The main scope of the system so far has been to build a reasonably realistic prototype in a relevant setting and then to vertically analyze privacy-related, societal, and economical implications of the individual design decisions and of fielding such a system. These conclusions will be presented throughout this section.

5.1 Deployment experiences

During the development of our prototype, we have been out on the streets for several weeks, testing and tuning the system, gaining experiences about difficult driving situations that challenge the system, and partially solving those problems. We have also been an entire day on a closed circuit, where a professional driver tested the Smart Tachograph under different conditions – from "normal" driving to a driving style that would qualify as very aggressive and highly risky for an average driver.

Any developer of vehicle navigation systems will probably have their own tale of solving the inherent imprecision of GPS measurements in urban areas, which was one of our first practical challenges. As probably many others before, we use a circle around the reported position to search for streets, and a sliding window to minimize the erroneous reporting of another street that is in the vicinity for a short period of time (e.g., at crossroads or when driving on a highway under a bridge). After ascertaining the street the vehicle is being driven on, the next point is to determine the speed limit for that street. This should be trivial in theory – a lookup in the geospatial database that contains, among other attributes, the speed limits for individual streets. In practice, however, this data is not easily available. After extensive research, we found only two providers of street topology data worldwide – Navteq and TeleAtlas.[5] All providers of GPS-based car navigation platforms seem to buy the raw geographical data from one of these two producers. Navteq has no speed limits recorded for Switzerland and the ones from TeleAtlas database were often incorrect. Although this is in part a rather pragmatic problem, that presumably occurs only for some geographical regions, there is also a conceptual issue behind it. Speed limits on individual streets change with quite a rapid pace, so that a CD containing them will be partly outdated from the first day of usage. Speed limits may change due to changes in the street architecture (e.g., a street enlargement may come with an increase in the speed limit), or traffic-policy reasons, but also because of short-termed construction sites. To have an outdated database in a system that could in consequence charge a sum one or two orders of magnitude larger than the true one, is obviously unacceptable. A prerequisite for such a system to work is thus to have an efficient way of propagating speed limit changes to vehicles. Many solutions are conceivable, but they all require quite a massive infrastructural support, which is unavailable today. A centralistic solution could be for example to have a publicly-accessible database where all changes are published. Vehicles would lookup that database on a regular basis, updating their local copy. A distributed way would be to have electronic tags on all speed limiting signs, which could be read by the vehicles.

---

[5] See www.navteq.com, and www.teletlas.com, respectively.

We also had positive experiences with the used technology. The measurements from the cheap, off-the-shelf acceleration sensors were always exact and we did not experience any failures. Having both the BTnode and the Bluetooth-enabled GPS-sensor wirelessly transmitting data to the laptop computer has also proven to be robust – the communication never failed.

5.2 Privacy and Security Considerations

The architecture of a system that continuously analyzes the driving parameters to ascertain momentary costs on a pay-per-use basis can be realized in three different ways. The first, chosen by some insurance companies and which is also the standard for road pricing schemes (see related work in Section 4), is to send all sensed data to the accounting entity, be it tax authority or insurance company. The data could be sent online (via the car's wide area communication system) or offline, on a regular basis. This solution is the most simple yet most privacy invasive.

A second possible implementation – at least for the insurance part – would be not to disclose the data by default, but store it and reveal it in order to get a retroactive reduction for a safe driving style. The data could be stored either locally in the car's blackbox, or it could be encrypted with the motorist's private key and transferred to the insurance company. To qualify for the reduction, the customer has to reveal the data to the insurer. This is the model chosen by Progressive for its TripSense product. Aside from being a different insurance model than the risk-dependent momentary insurance rate presented here, this solution has several drawbacks. As pointed out in section 4, in order to get the price advantage, the customer has to reveal all data and thus give up privacy in this model, too. Furthermore, the responsibility lies with the customer, and it is unclear what the long-term consequences of both revealing or not revealing the data would be. Finally, it is unclear how this approach could work for a road pricing scheme, being thus rather unsuited for a generic solution for all types of traffic costs.

In the third and chosen model, all data is processed locally and only the total rate is transferred to the accounting authority, as described in section 3.2. The data is processed by a plug-in that has been downloaded from the respective authority (technically a Java-class). Such a client-side personalization insurance scheme [3] guarantees a high-level of privacy, since the accounting authority receives only a monthly sum. Should it be high, there can still be a multitude of reasons – a large amount of safe driving, or a small amount of risky driving under bad weather conditions, to name just two. Past whereabouts and behavior of users are protected, yet they pay their fair share. Moreover, since the sensor data needs not to be cached (the continuously incremented overall sum will suffice), the driver retains full control over his or her data – at least in theory. In practice, however, the sensor data might be stored, e.g., for a possible later legal dispute on the charged amounts. In such a scenario, the data could be used in a lawsuit against the driver, after all.

In terms of mutual safety between insurer and customer, all models have to face several challenges. As more thoroughly analyzed in [3], one issue induced by our client-side personalization model is the driver tampering with the software module downloaded from the accounting authority. However, through the use of a trusted computing platform on the vehicle, this problem could be overcome. The trusted platform would verify the software signed by the accounting authority, deciding upon its authenticity. In all implementation versions, on the other hand, the user could try a man-in-the-middle attack, modifying the message transmitted by the black box to the accounting authority. By having the black box signing the messages, such an attack can be prevented. By using timestamps in the messages, replay-attacks (i.e., replacing the message with an original but older message) can also be easily avoided. Finally, the vehicle owner could try to tamper with the vehicle sensors. For example, he could cap the acceleration sensors, or cover the rain sensor. This security issue is also common to all implementation models. However, this would not only potentially endanger the driver by disabling some important security functions such as the electronic stability programs, it can also be further avoided by including most of the sensors in a tamper-proof hardware environment (this would hardly work for a rain sensor, though).

The most obvious way an accounting authority could try to cheat would be to charge a different (i.e., higher) sum than the real one. This problem is common to all architectures as well, but in the client-side personalization paradigm the user has the most effective means to verify the claims coming from different accounting authorities. One way would be to run the same software on a client-trusted platform with the same sensorial input. Another possibility would be to have the black box issued by a third party trusted by both sides (e.g., a governmental agency). The black box would receive the billing contract (digitally signed by both parties), verify the signatures, extract the billing formula, and compute the sum in a secured environment.

Although we have not implemented in our first demonstrator any of these security features, it seems that the client-side personalization paradigm does not add crucial security-related issues when compared to a more privacy-invasive approach.

5.3 Economic Feasibility of Pay-per-Risk Insurances

Among the first issues to suggest themselves when thinking of pay-per-risk insurances is the question of economical feasibility. Is there a market attractive enough for both sellers to offer such insurances and customers to buy them? The mere technological possibility does not imply economical feasibility. There are many examples of products and services that could easily be offered per-use, but due to economical reasons are charged on a flat basis. Ski resorts, for example, could charge skiers for every ride, yet all over the world ski passes are offered almost exclusively as per-day flat charges. Likewise, breakfasts in hotels could also be charged “per-use”, yet more often than not they are “all you can eat” (a flat fee). Looking at insurances, the short answer would be that there is a market for pay-per-risk insurances, since at least two companies have started offering these (see section 4 on related work). A more detailed answer would be that it is in the insurer’s best interest to identify the good risks, and offer them a more advantageous rate that reflects their actual risk. The insurer would thus be able to gain new market shares in the attractive market of low-risk drivers, while at the same time filtering out the “bad risks” (due to the then increased rate they would have to pay to stay with the insurer) [2, 10]. On customer side, there would probably also be enough interest for such a model. Lower risk customers, who nowadays partially subsidize the higher risk drivers, would presumably welcome a model that could bring them important savings.

Even if individual pay-per-risk insurance schemes still lie a few years in the future, there are some areas more likely than others to be early adopters of such technology on a large scale. All domains where people drive cars that do not belong to them, but are merely borrowed, rented, or co-owned, fall into this category. In such areas, the driver of the car has an explicit or implicit obligation to handle the confided vehicle with care. And the lending or owning part can more easily enforce a black box analyzing the way customers handle the assets. For example, a car rental company has started as early as 2001 to protect its vehicles from overspeeding by charging customers exceeding 70mph with a high fine.[6] Car sharing models, such as the popular Mobility-network in Switzerland,[7] have to pay relatively high insurance rates because of the few accident-prone risky drivers. Eventually (after two or three accidents in a short period of time) these drivers are sorted out, but the damage is done, and through the continuous flow of new members the insurance rates stay high. Car sharing networks would presumably be happy to detect such risky drivers before they cause accidents, by having their driving style analyzed from the very first ride. And the "How am I driving?" sign could soon disappear as well, if companies would start equipping their car fleets with black boxes analyzing (and reporting) the way their drivers behave. Furthermore, in such examples it seems less probable that drivers would try to tamper with the system since the vehicle returns periodically to its owner, who may detect the fraud.

5.4 Measuring the Risk

As a proof of concept of the Smart Tachograph paradigm, we designed two plug-ins, one for road pricing, the other one for a liability insurance. The road pricing plug-in is a simple one, it only differentiates between streets in the city and roads outside the city, and between high-peak and low-peak hours. It charges, depending on these two parameters, a fee between 2 and 8 cents per driven kilometer. Since we have access to comprehensive geospatial data, refining the plug-in to differentiate between individual streets and taking into account more fine-granular time slots would be a trivial (yet laborious) task.

The insurance plug-in is more complex. It takes into account five parameters: type of street, difference between driven speed and speed limit, the two acceleration types (longitudinal and cross acceleration), and the time of day. Converting sensor data into an accident probability and expressing this risk in a monetary way are obviously no trivial tasks. We made the following assumptions: There is a basic per-kilometer risk, that depends on the type of street and on the time of day, as suggested in [10]. The per-kilometer accident risk is lowest on highways, followed by country roads, and is highest in cities. It varies between 2 and 10 cents per kilometer. In a real deployment, this minimum would presumably also depend on the "classical" risk factors such as driver age or experience, which are used today to classify customers into driver categories.

With respect to acceleration, we acknowledge that some thresholds have to exist. At every traffic light stop, every departure, and every curve taken by the driver, there are accelerations involved. Such accelerations within normal limits pose no special danger

---

[6] See http://archives.cnn.com/2001/TECH/ptech/06/22/gps.airiq/.

[7] See www.mobility.ch.

and have to be allowed without punishment. According to our subjective danger sensation correlated to the measured accelerations, the thresholds for the prototypical plug-in have been set to $2m/s^2$ ($1/5g$) for the cross accelerations as well as the positive longitudinal acceleration, and to $3.5m/s^2$ for negative longitudinal acceleration (braking). After that, we assume that the risk increases exponentially with a low base of 1.5. We further assume that exceeding the speed limit with 10% does not notably increase the risk, and only after that it increases exponentially as well, but with a lower base than in the case of acceleration. This base further depends on the street type, varying between 1.05 for highways and 1.2 within cities. The overall formula for the momentary insurance rate for the experimental system (expressed in cents per kilometer) thus results in

$$R = B_{st,t} * 1.5^{Max(0,\frac{A_c-2}{2})} * 1.5^{Max(0,\frac{Abs(A_{ln})-3.5}{3.5})} * VC_{st}^{Max(0,\frac{V_{driven}}{1.1*V_{limit}}-1)} \tag{1}$$

where R is the resulting momentary insurance rate, B the basic rate depending on street type and time, A_c the cross acceleration, A_{ln} the negative longitudinal acceleration (actually, this is a slightly simplified version of the formula, ignoring the positive acceleration), VC the base for the speeding coefficient, V_{driven} and V_{limit} the driven speed and the speed limit, respectively.

From subsequent discussions with representatives from the vehicle insurance industry, we learned that they have a pretty clear picture of when a driving style becomes dangerous. Vehicle insurance companies have a strong accident research tradition, where they analyzed such data for many years. Until now, they have not been able to transfer this know-how to insurance schemes due to the bulky equipment and high costs involved. It appears that some of our assumptions have been quite exact, while other were rather erroneous. The experts confirmed that a defensive driving style usually remains under $0.2g$ and that the cross acceleration becomes dangerous around $1g$. At $1g$ cross acceleration, our formula results in an insurance fee almost 8 times higher than the basic fee (and then continues to grow exponentially). Choosing an exponential function seems to have been the right decision, details aside. Our formula seems to have over-estimated the velocity component though. Overspeeding is one determinant for accidents, but this depends much more on the actual context than, for example, with accelerations. Often, speed limits are set very conservative or with respect to other criteria (such as noise reduction), thus overspeeding may often have no influence whatsoever on the risk of causing an accident. Another highly relevant parameter to be considered, according to the industry representatives, would be the distance to the car in front, especially when it is correlated to the driven speed.

5.5 Practicability

Requesting customers to sign a complex mathematic formula as the one presented in the previous section as part of an insurance contract seems to be unacceptable, no matter how much a realistic risk estimation it encompasses. Insurance companies know that a good contract (one easily accepted by customers) must comprise two features: simplicity and upper bounds. To keep complexity low, two or no more than three new attributes should be considered. Because of their outstanding importance, these could be cross acceleration and the distance to the car in front. To further reduce the complexity, discrete intervals are preferable over a continuous function. There could be for example three different classes of driving: safe, normal, and dangerous driving. After a journey, the motorist would get different per-kilometer prices for the times he or she exhibited any of these styles.

Not having upper bounds in the insurance tariff could lead in extreme circumstances (when an accident will most likely happen) to a chronic situation in which the insurance rate would get as high as the expected damage costs. This would, of course, undermine the idea of an insurance, making such a practice unacceptable. But even in less extreme situations, customers seem to feel much more at ease if they know what to expect in terms of maximum possible costs, no matter of their behavior.

A common misconception, on the other hand, is that such highly personalized insurance rates, that finally lead to “risk communities of one” (Andreas Schraft, Head Risk Engineer of SwissRe, a reinsurance company, in a recent talk) would undermine the idea of insurances and would thus not be a realistic concept. This often-heard interpretation originates in the (wrong) assumption that an insurance company has to insure the same risk for a large number of people in order to work. As a matter of fact, insurance companies insure unlikely, but high-cost events of individuals with a sum that represents the probability of that event’s occurrence times the costs it will produce (plus the insurance’s security margin and its profit). Since many individuals are insured with one company, some of these events will occur, most will not, and due to the law of large numbers the company will pay exactly the expected cumulated costs (if it estimated correctly the individual risks). While the company does need a large number of individuals for the system to work, it does not need to insure the same risk for everyone – highly personalized risks will do just fine.

5.6 Societal Issues

“Having everyone paying for his or her individual risk and usage pattern is much fairer than today’s pricing scheme!” Is it? There are several examples of costs that could easily be allocated to their originators, yet they are burdened by the society as a whole. One such example are health insurances. Instead of evaluating the individual risk of illness based on age, gender, and health history, many countries have decided to spread those costs throughout the society, willingly cross financing the elderly or the ones with chronical diseases from the young and healthy. However, personal mobility does not seem to have the same societal value as health and does not seem to be in need of redistribution. Today’s cross-financing of traffic costs is probably more the consequence of the imperfections stressed out in this paper, than the result of a socio-political master plan. Hence, charging the responsible parties for the caused traffic costs would presumably be seen by many as fairer indeed. Economists have also argued that in the extreme situation where all information asymmetry (and thus adverse selection and moral hazard) could be eliminated, a Pareto-type welfare improvement could appear, thus having all drivers being better off [7].

The ubiquitous presence of a system such as the Smart Tachograph could have other consequences as well. If many vehicles would have such devices installed, then the technical prerequisites would be given that also authorities such as the traffic police can require access to that data. From a technical perspective, this would be an easy game: generating a new plug-in and enforcing it to all vehicles. To illustrate this scenario, we have included a traffic police plug-in into our first prototype. When driving above the speed limit, it issues an audio warning: "You are above the speed limit. Please slow down." If ten seconds later the car is still above the speed limit, an automatic traffic fine is issued. Since there have been some similar examples lately (e.g., airlines forced to give away data on passengers and their flying habits, data that they had initially collected for their frequent flyer programs), this scenario seems not to be that far from reality. We thus think that the Smart Tachograph is a good example to illustrate some of the Pandora boxes that could be opened by ubiquitous computing technology in the near future.

6 Conclusion and Outlook

We have presented the Smart Tachograph, a prototypical platform that facilitates an individual and accurate accounting of generic traffic costs. We have built a prototypical black box containing the sensors used by the system and a software infrastructure supporting a set of basic features. Even if our prototype does not yet address many practical deployment issues (such as guaranteeing mutual security between accounting authority and customer), it already supports various core features of a deployment-ready system. Being a generic platform, it allows different kinds of traffic costs to be measured and billed through it – we included road pricing and insurance rates for illustration. The prototype is easily extensible to include new sensors that may be relevant for measuring some of these costs. It can be extended to include new accounting authorities, too. Where applicable, the infrastructure allows several accounting authorities to compete for the driver's favor. The prototype already includes several insurers. The system automatically downloads and presents their different offers to the driver, who may choose among them.

We have subsequently analyzed several technical, societal, and economic issues that could arise from the deployment of such a system. Such consequences are often outside the focus of typical ubiquitous computing prototypes. Thus, the Smart Tachograph seems to be a good ubiquitous computing case study. It involves a collection of technologies and concepts typical for ubiquitous computing applications, such as sensors and sensor nodes, wireless communication, location and context awareness, and machine-to-machine communication. There is a realistic business model behind it. And, as mentioned, it allows further exploration of many highly-relevant questions regarding economic models, welfare, security, or fairness and trust throughout the society, as well as the specific tradeoffs among these. As a consequence, we have further made the point for a privacy-friendly solution that could at least alleviate some of the potential societal drawbacks, and included this solution in our prototype.

There are several interesting future directions of research. For example, evaluating the system – how well does it ascertain different kind of costs? A more intense dialogue with stakeholders from the insurance industry and government agencies could be enlightening. How robust are these calculations with respect to different driving styles and environmental conditions? A more structured testbed seems necessary in order to answer this question. A study with other categories of stakeholders, mainly with potential customers, also seems imperative. By presenting them different possible implementations of the Smart Tachograph with their specific advantages and drawbacks, a substantial overview of user opinions could result. Finally, including the different mutual security features in the system would imply a large leap towards a realistic system and would more clearly show the related challenges.

Acknowledgements. Christoph Plüss has programmed most of the Smart Tachograph during his master thesis and did an outstanding work. Dr. Jochen Jagob pointed out information asymmetry, adverse selection, and moral hazard as relevant economic issues. Prof. Friedemann Mattern, Prof. Hans Gellersen, and Dr. Marc Langheinrich have provided many helpful comments on earlier drafts of this paper. The Gottlieb Daimler- and Karl Benz-Foundation, Germany, has generously supported a large part of this work as part of the project "Living in a Smart Environment – Implications of Ubiquitous Computing".

References

1. George Akerlof. The Market for Lemons: Qualitative Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
2. Vlad Coroama and Norbert Höckl. Pervasive Insurance Markets and their Consequences. *First Int. Workshop on Sustainable Pervasive Computing at Pervasive 2004*, April 2004.
3. Vlad Coroama and Marc Langheinrich. Personalized Vehicle Insurance Rates – A Case for Client-Side Personalization in Ubiquitous Computing. *Workshop on Privacy-Enhanced Personalization. CHI 2006*, April 2005.
4. Vlad Coroama and Marc Langheinrich. The Smart Tachograph. *Video submission abstract. Adjunct Proceedings of UbiComp 2005*, September 2005.
5. Anind Dey. *Providing Architectural Support for Building Context-Aware Applications*. PhD thesis, College of Computing, Georgia Tech, December 2000.
6. Jonas Eliasson and Mattias Lundberg. Road Pricing in Urban Areas. www.transport-pricing.net/download/swedishreport.pdf, January 2003.
7. Lilia Filipova and Peter Welzel. Reducing Asymmetric Information in Insurance Markets: Cars with Black Boxes. In *Proceedings of the 32nd Conference of the European Association for Research in Industrial Economics (EARIE)*, September 2005.
8. Tim Kindberg, John Barton, Jeff Morgan, Gene Becker, Debbie Caswell, Philippe Debaty, Gita Gopal, Marcos Frid, Venky Krishnan, Howard Morris, John Schettino, Bill Serra, and Mirjana Spasojevic. People, Places, Things: Web Presence for the Real World. *Mobile Networks and Applications*, 7:365–376, 2002.
9. Todd Litman. Distance-based vehicle insurance. Victoria Transport Policy Institute, 2003.
10. Matthias Oberholzer. *Strategische Implikationen des Ubiquitous Computing für das Nichtleben-Geschäft im Privatkundensegment der Assekuranz*. PhD thesis, Basel University, Switzerland, 2003.
11. Lucienne Rey. publifocus – Road Pricing. Technical report, TA-Swiss, July 2004.
12. "Der Standard" staff. Wien droht doppelt so viel Autoverkehr. *Der Standard – Austrian daily newspaper*, April 24th-25th, 2004.
13. The Economist staff. Driven to radicalism. *The Economist*, June 11th:33–34, 2005.
14. The Economist staff. Jam yesterday. *The Economist*, June 11th:14, 2005.