

# *User Interface Design as Systems Design*

**Harold Thimbleby,<sup>1</sup> Ann Blandford,<sup>1</sup> Paul Cairns,<sup>1</sup>  
Paul Curzon<sup>2</sup> & Matt Jones<sup>3</sup>**

*1 UCLIC, University College London Interaction Centre  
London, WC1E 6BT, UK.*

*2 Interaction Design Centre, Middlesex University  
London, N14 4YZ, UK.*

*3 Department of Computer Science, University of Waikato  
Private Bag 3105, Hamilton, New Zealand.*

Emails: *h.thimbleby|a.blandford|p.cairns@ucl.ac.uk, p.curzon@mdx.ac.uk, always@acm.org*

**User interfaces interface people to systems. In many cases the systems are large and complex, and entangled with organisational structures and policies. Inappropriate design can lead to many of the idiosyncrasies of the deeper system negatively influencing the user's experience. So far as possible, then, it is important to separate how things work from what they do for the user.**

**We illustrate this argument with a case study of a Graphical User Interface (GUI) as an interface to a particular complex system, the British Rail ticketing system. We show that the level of interaction imposed on the user is inappropriate to the user's task needs. We show how user interface problems are based on deeper organisational conventions that are of little interest to users. Furthermore, the GUI permits a level of complexity that is beyond the resources of the organisation to manage well: the GUI has software engineering quality problems as well. Personal technologies are recognised as a standard solution to this sort of problem, and we discuss how they would work in this context. The insight, however, is not that personal technologies are a good thing (in fact, we highlight some problems) but that, in this case, they allow the user interface to work at a level appropriate to the user's actual task: travelling. We argue that personal technologies might be a driver to simplify the interface to organisational conventions, and so further improve the user experience.**

**Keywords:** Personal technology, ubiquitous computing, interaction design, public user interfaces, systems engineering, system complexity.

## **1 Introduction**

Standard good HCI practice is to find out about user tasks as one of the first steps in design, otherwise it is tempting to start design from the functionality provided by (or planned to be provided by) the system. The system functions may interact in interesting and complex ways, but the HCI issue is to support the user task, rather than give access to the underlying system *per se*. One of the most successful user interfaces to general purpose systems is the graphical user interface (GUI), and its success in providing access to PC functionality means it is frequently a candidate for use in other applications, which are supposed to support much simpler tasks. Unfortunately a GUI does not constrain the designer much, and does not really motivate finding a clear view of the task, because it is possible to support almost any functionality and hence almost any task.

This paper argues that proper design can transform the user experience by hiding irrelevant complexity. We support our abstract discussion with a case study of rail ticket vending machines as an example of complex public technology, but the principles underpinning the discussion apply very widely across a range of technologies and contexts of use. We show how appropriate design can eliminate conventional user interface problems for the task in hand, including some problems of

implementation quality. We will see, further, that having eliminated conventional problems, latent issues become salient: user acceptance now depends on solving less technical but more social problems.

### ***1.1 The Abstract Systems Design View***

Many goals cannot be achieved by simple systems. Most interesting systems involve people interacting with each other and often with complex computer systems, and all within some more abstract system that is intended to achieve goals such as staying in business, whilst adhering to regulatory and legal requirements. It is well known that interacting subsystems can exhibit complex behaviour that no subsystem alone may be able to exhibit, such as deadlock, livelock, non-determinism and feature-interaction (multiple features combining in unexpected and usually unhelpful ways). In the specific case of human-system interaction, there are particular problems such as closure errors (including post- and pre-completion errors), which are psychological representations of the abstract system concepts like deadlock.

To avoid such undesirable problems arising in new systems it is crucial that the interfaces between subsystems are well-defined and understood and, moreover, simple. Once interfaces are well-defined, it is possible to change subsystems (for instance, to make them more efficient) without detriment to the overall system behaviour provided the interfaces are maintained. In fortuitous cases, of course, one may improve systems by changing interfaces, but in practice it has been found that carefully managing the separation of concerns of different subsystems through well-defined (and simple) interfaces is crucial to successful design. Conversely, it has been found that many system failures can be attributed to circumventions, accidental, overlooked or (in hindsight) unwisely exploited, of what were supposed to be well-defined interfaces. These systemic issues are no less true of human-system or human-computer interaction than of systems engineering.

Typically the interface between the human and the technical system user interfaces are deeply entangled with the internal details of the system. User interfaces are often designed to control as many features of the system as possible, so this entanglement is inevitable. Indeed there is a line of thought within HCI (e.g., Norman, 1986) that one central role for the interface is to communicate the underlying system model to the user. Interfaces are also deeply entangled with the social, perceptual and cognitive details of the human: hence the lively debates about design, colour, screen layout, culture, and so forth. These issues have a significant impact on the success of the combined human-system. It is a truism that there are no easy design solutions, and one often has to build prototypes and iterate design to identify feasible improvements.

In human computer systems, there are usually many parallel interfaces (in the systems sense). For example, if interaction using the primary interface deadlocks, there are may be interfaces available that support reset, cancel and timeout: and these operations may be initiated by other subsystems in the computer or in the user (or even bystanders helping the user). Users, of course, usually reflect on their behaviour and can walk away or take other action. In some contexts, avionics being a case in point, with hard real time constraints, the design of the control interfaces is very difficult if they are not to create, rather than solve, further problems of interaction. There is a great deal of interest in designing systems so that users awareness of interaction (and likely consequences) does not lead to panic or other extreme actions outside of the specification of the primary interfaces.

Technology is changing, and ideas about the best ways of designing user interfaces are also changing. For many tasks, GUIs are easier to use than command lines, and for a while they have been the best way of interacting with general purpose PCs. We have excellent visual computing skills, and GUIs exploit these skills (e.g., recognising overlapping windows: see Friedhoff & Peercy, 2000) to raise the level of complexity that is easy to interact with. However, the recognised strengths of GUIs do not mean they are best in general. Indeed, because GUIs are so good, they can give the surface impression of providing good solutions while actually forcing the interaction to an inappropriate level.

If we are to realise the significant benefits of new technologies, we must think in terms of system design, rather than in terms of conventional user interface design. Providing fancy, powerful technology for user interfaces will not necessarily circumvent the problems of complex interaction; rather, new technologies provide opportunities to rethink the design of user interfaces.

## 2 Case Study

Many cities world-wide depend on public walk up and use ticket vending machines for the smooth running of their transportation infrastructure. As a case study, we take a ticket machine installed in early 2001. The machine is a realistic representation of the state of the art in walk up and use technology and is used by the public to undertake real and important daily tasks.

The ticket machine considered here has numerous user interface problems; thirty years of relevant usability research and practice has had very little obvious impact on its design, even though incorrect operation of the machine can result in the user being unable to travel or being fined for travelling with incorrect tickets. An alternative view is that the designers are concerned with usability, but that they are only empowered to address surface-level usability concerns (such a colour any button layout) rather than the complex usability concerns that might be regarded as a systems engineering enterprise. The design well illustrates the endemic poor usability in the current every-day computing infrastructure, and the lack of impact of deep usability concerns on designers, manufacturers and those charged with installing new systems.

Although interaction problems with public technology are well known, the literature offers very little explicit help for designers to overcome the fundamental problems; researchers in recent years have been more interested in adding new interaction technologies and intelligence to such devices including personal technologies (PTs), the option that this paper examines. We show that appropriate use of PTs solves many user interface problems, and we show that the reason is because the user interface level can be set appropriately. That is, PTs are not making user interfaces better because they are exciting new solutions, but because they allow the user interface designer to define the systems interface more appropriately than might otherwise have been possible. However, as noted above, they also introduce new usability challenges.

Although some work has been done on analysing the interactive problems of public technology (e.g., Stanton & Baber, 1997), much work involves attempts at adding sophisticated interactive capabilities to ticket machines and kiosks. This work has included adding emotional avatars (Christian & Avery, 1998) and new input techniques (e.g., for sign language interaction). Our analysis suggests this approach, at least if not guided by more systems level thinking, is misguided. Our case study shows very clearly how organisational issues can cause bad user interface experiences that cannot be hidden from the user in conventional approaches.

### 2.1 Methodological Issues

The context of work studied here is enormously complex, and there is no established methodology for investigating issues such as the interaction between usability, organisational and regulatory conditions. There are also ethical challenges: if the work were commissioned by the interested parties then the independence of researchers would be compromised; conversely, those parties should have a right to answer any criticisms of their systems. We have tried to engage in discussion with the main interested parties (the manufacturers and the rail operators).

The ticket machine manufacturer has been very co-operative with us, though they work under constraints. They provide hardware to their customer s requirements. In turn, much of the hardware they use is constrained by external issues, such as financial regulation beyond their control. For example, there may be a better way of users interacting with the credit card facilities, yet this is supplied as a sealed unit for security reasons, and it cannot be changed. The rail operators have declined to comment.

New evaluation techniques for difficult-to-study contexts and applications are needed (Cheverst, *et al.*, 2000). Some have argued that the definition of usability needs to be extended for these new contexts to include notions of intimacy, beauty and social acceptance. The work reported here is based partly on the experiences of the authors and exploratory investigation of the system interface, partly on periods of observation of others using the ticket machines in question, and partly on informal discussions with various stakeholders: users, railway staff and the hardware manufacturers.

---

<sup>†</sup> There is a difference between “state of the art” relative to the potential of current technology and “state of the art” *in practice* as actually used on the British national rail system. We mean the latter here.

All examples discussed in the body of this paper have been witnessed by the authors or have been experienced during everyday use by at least one of them; the scenarios are not contrived.

A JavaScript simulation has been made of the ticket machine user interface that is adequate to explore some of the usability issues raised in this paper, and from which most of our claims can be double-checked. (The simulation obviously doesn't provide tickets or change.) It can be used in a browser over the web (e.g., for evaluation) and the complete source code is available. See <http://www.cs.mdx.ac.uk/harold/tvm/> for the simulation, which for copyright reasons may be used for non-commercial research purposes only.

## ***2.2 A Brief History of Passenger Transport Systems***

Before explaining the case study further, it is helpful to briefly review the organisational history in which the user's task is situated.

To stay in business, transport systems must collect revenue from travellers (or from other sources, such as subsidies) and, in turn, must provide an acceptable service to those travellers to ensure repeat business. Given human nature and the unreliability of transport systems, relying on trust is insufficient: some passengers do not pay, and some modes of transport fail. Thus tickets were invented as a reliable token of the contract between passengers and transport operators. In Victorian times, railway transport became very sophisticated, and the range and type of tickets expanded. Tickets were printed on card by mechanical devices; as more uses for tickets became apparent (e.g., for distinguishing between contracts with different railway companies) different colours of tickets were introduced. And so on.

Today, over a century later, many interacting layers of complexity have been accreted. Rail companies introduced many sorts of ticket fares, depending on modes of travel, and numerous loyalty schemes involving identity cards in addition to conventional tickets such as discounts for family travel at certain times of day. Passengers may use credit cards to pay for tickets: these transactions are heavily regulated. In the United Kingdom another layer of complexity was introduced by the privatisation of the railway network, resulting in many competing organisations and authorities with overlapping interests.

It is well known that the British railway system is impenetrable, particularly for passengers. In our terms, the complexity that passengers see arises because the interface between passengers, as one component of the transport system, and the railway companies, as another component of the system, has been defined so that passengers see and interact with the internal complexity of the transport infrastructure. Almost all of that is irrelevant to the needs of the passenger, namely to travel economically and in good time.

Rail companies provide interfaces between passengers and the complex ticketing systems: they provide human ticket clerks. The human clerks are supposed to provide optimal ticketing (and other advice) for passenger travel requests. Unfortunately ticket clerks are an expense and managing their costs, in turn, encourages further complexity in ticketing. If clerks cannot always be present, then new forms of tickets have to be introduced (e.g., permits to travel)!

There is a straightforward business case to automate the ticketing process, and there are substantial potential gains. Passengers will be pleased to have shorter ticketing queues, for instance, and rail companies will be pleased to better manage the demands on ticket clerks in order to save costs, and perhaps to improve ticket clerks' morale.

All this, the history and business case, is well known, though the systems complexity might not be widely appreciated as a highly relevant factor of user interface design for ticketing systems.

## ***2.3 The Conventional Solution***

A conventional user interface design solution to rail ticketing would be to automate as much as possible of the ticket clerk's role, and ticket machines were widely used from the 1980s onwards. Ticket machines are intended to replace the clerk for a given fraction (e.g., 80%) of passenger ticketing needs. As might be expected from our introductory comments, automating the existing interface between passenger and complex rail system cannot avoid the complexities apparent in the original interface itself, unless the specification of that interface is simplified. So, not only do we find the expected deadlock and other symptoms of bad design, we also find that the new automated

system is itself so complex as to introduce its own quality control problems. It is incomplete and unreliable: presumably because making it work well was beyond its designers' abilities.

We now explore the design of the current ticket machine user interface: we show how user interface problems are entwined with system complexity.

The particular ticket machine discussed here was installed on Welwyn North station (north of London's King's Cross station) in early 2001. Similar models are installed throughout the UK. The machine has a rectangular purple label saying it is 'Easy to use'.

ASCOM, the hardware manufacturer, has installed over 30,000 ticket vending machines around the world (<http://www.ascom.com>, Sept. 2001), with half still in use, handling over 40 million transactions a day. With this scale of operation, thousands of installed machines, any usability problems will have a huge aggregated social impact. In Britain, the 1000+ ASCOM machines have an annual revenue well in excess of £100 million: bad user interface design can be costly; conversely, incremental improvements to user interface design has worthwhile payoffs. If users make mistakes – even if in any one location particular types of mistake are experienced infrequently – the overall cost to rail operators (and in turn to the machine manufacturers) will be huge. On individual stations the usability of a ticket machine has an impact both on traveller satisfaction and on rail staff satisfaction. The larger the station, of course, the greater the local impact will be. There is also the knock-on loss of reputation when travellers see other users having difficulties. Ticket machines are a serious application to study: not only do they raise revenue for operators, but difficulties with them can mean users miss trains and have their lives disrupted, which has a further impact on others.

## 2.4 The Ticket Machine

Figure 1 shows the ticket machine kiosk discussed in this paper. The main buttons are on each side of the screen, and there are two hard buttons, one for information (towards the bottom left of the screen, blue with a white italic *i*) and one for cancelling (towards the bottom right of the screen, in red, and labelled explicitly as CANCEL, in capitals). The small *i* button provides help, yet there is no explanation of this (perhaps on the principle that users should recognise *i* as the international standard sign meaning information).

The interface employs a standard GUI paradigm, but is not, in fact, a GUI: the on-screen labels look like buttons, but do not act like them. Likely surface damage would make a touch screen unreliable (so buttons would probably still be needed). Many passengers we spoke to said they thought the screen was touch sensitive, and they found out the hard way that it has buttons off to each side. One user was observed to repeatedly press the screen; many press the screen for the final buy ticket step. The designers have applied a widely recognised interface paradigm in an inappropriate context (possibly because the early prototypes they worked on *were* GUIs on their own computers, without a representation of the final context of use).

The actual buttons are placed some distance from the screen (see Figure 1) because reinforcing the machine against vandalism (e.g., on unattended stations) is a major design factor: a steel frame means buttons cannot be closer. It is likely that if the screen information looked less like buttons, and were carefully aligned with the markers linking the buttons to the screen area, the mistake of thinking the screen was touch sensitive (as well as the obvious alignment mistakes) would be less likely. For instance, the rectangular button images could be changed to arrows pointing towards the buttons users must press, in forms like [PURCHASE\_TICKET>>> and possibly animated. Whether this would be an effective idea could easily be checked by experiment, but the point is that the interface needs to be designed for the actual context of use, with all the ensuing constraints.

The top-level screen, which appears either if a user presses [CANCEL] or after a short time out, provides some common destinations, such as King's Cross (a major rail terminus in central London). The following text is shown when the *i* button is pressed when the machine is at its top level:

Using This Machine

1. Select your destination.
2. Select the ticket type required.
3. Select Railcard discount if applicable.
4. Pay by card or cash.

The interaction designers seem to think that ticket type and railcard discount are different things a point we return to below. Note also that this screen is only accessible from the top level at a point where the user has not yet become aware that there is a likely confusion over ticket types and discounts. As is widely recognised, users only read the manual as a last resort, and are therefore unlikely to read this information before starting (this information is not available later in the interaction).



**Figure 1:** The ticket vending machine. The screen (diagonal 25cm, underneath the 1 arrow) is at eye level for a standing adult, its centre 1.47m off the ground. There are five hard buttons on each side of the display screen (looking like triangular arrows) with meanings as displayed on the screen. There are two buttons with fixed meanings, for information (blue, on the left) and for cancelling (red, on the right). There are ten unlabelled buttons beneath the display that do nothing. The user can insert coins, credit cards or paper currency (underneath the 2 arrow). A slot at the lower right, that looks like it might accept proprietary tokens, is apparently used as an alternative slot to return rejected notes. Tickets and change are obtained at the bottom of the machine, under the 3 arrow.

## 2.5 Examples of Use

We now give a detailed list of some design and usability problems encountered with the ticket machine. The point is to illustrate how the interaction takes place at the wrong level for the user, leading to many unnecessary difficulties and breakdowns, and also to set the scene for discussing the potential advantages of PT-based solutions and requirements on their design.

### 2.5.1 Buying a Ticket

First, we present the complex set of steps a user has to perform to achieve the most important goal: buying a ticket.

One of the authors has a routine task of buying a return train ticket for his daily work commute from Welwyn North station (in the countryside) to New Southgate station (in London suburbia). Because he is a regular traveller, he has a Network Card (a sort of loyalty card), which allows certain discounts on certain journeys.

New Southgate is not a destination mentioned on the top level screen of the ticket machine, so the user will select Other destinations and select through the alphabet (through several screens) until New Southgate appears. If the user makes no mistakes, this takes 10 button presses.

None of the choices shown at this stage of the interaction mentions a network card discount, and one button is unused so it would appear that all choices the machine supports must have been catered for. If the user selects the choice More Fares, they get no network card choice in the next screen either. If the help button (*i*) is pressed at this stage, the screen displays nonsense:

Help not implemented for ContextID 0011.

It is difficult to understand why help has not been implemented. As well as being a comment about bad user interface design (including lack of testing), this error is a sad comment about the quality of software development, and poor maintenance (this problem has been present, despite user interface upgrades, for over nine months, and has survived the introduction of multi-lingual

internationalisation a process which one might have supposed would have reviewed all system texts).

The first time the machine was used, it became necessary to ask for help from a rail employee, who asserted that there was a network card option. With help from the official, we persevered, and discovered that from the screen with the spare button, you could choose the option cheap day return and then get more options, *including* discounts. This is an example of unnecessary complexity for the user: discounts are not an after-thought, but a central part of the task of getting the most economical ticket for the journey.

However on the screen providing network card discounts, there is now no previous screen option, so the user cannot change any other part of their choice. The hard key called [CANCEL] remains, of course, but using it at this stage takes the user back to the initial screen. The key could be better called [RESET] or [START\_AGAIN], as it resets the entire interaction rather than cancelling the last step (which the soft key [PREVIOUS\_SCREEN] does on many screens).

If the user selects Railcard discounts from here, the next screen confirms that they have selected a ticket to New Southgate, but does not confirm what sort of ticket it is (is it a single, season, a return, or what?). There are two network card choices, neither of which match our user's network card. The user selects one, and is now ready to pay. The screen does not confirm what sort of ticket the user is getting, and there is no choice to permit a return to a previous screen to revise (or confirm) any previous choices. The user can now put in cash or cards and get the ticket. However, most of the screen stays steady as the cash is inserted, and no special feedback is given to reassure the user that cash is being entered successfully. (We have watched people insert notes, and jump around frustrated, since the machine appears to have eaten their cash but does not appear to be doing anything helpful.)

These usability difficulties indicate a lack of understanding of even the simplest user tasks (as perceived by users), or of the importance of context-relevant feedback on the state of the machine, or of user-relevant help. The task structure reflects the complexity of the organisational task model, rather than the relative simplicity of the user's.

### *2.5.2 Sometimes It Does Not Work*

On several occasions, a few seconds after the ticket machine shows insert payment (right at the end of a 10-or-more key press interaction to buy a ticket) we have found it unable to actually issue a ticket. It just says Your cash is being returned, when no cash has been put in the machine to return. One reason for this problem is that tickets can be printed on two stocks of paper. Suppose the ticket machine has run out of normal ticket paper. The user interface is designed so that type of ticket is one of the last choices a user makes, so it is only at this stage that the machine knows whether it can print the requested ticket. If this is the case, the interaction sequence could be re-designed so that choice of ticket type is made earlier in the interaction, or the machine could display on the initial screen (just as it does when there is no change) that it has restrictions. However one looks at it, this is a poor feature of user interaction design.

Similarly the machine sometimes will only accept credit cards (perhaps due to running out of change or the cash box being full). Originally, the user was only told this at the end of an interaction. Worse, the machine first indicated that all payment methods are available, before deleting those that are not. The delay is long enough for someone to start inserting coins, so that they may no longer be looking at the screen.

Some of these initial problems have now been fixed, but others remain. The greatest problems are exacerbated by ticket regulations.

### *2.5.3 Group Travel and Multiple Ticket Problems*

Suppose the user is part of a family group and wants to buy a ticket for themselves and the party travelling with them. It takes about 13 key presses to buy one ticket. Unfortunately the ticket machine then reverts to the top level screen, so it is not possible to say, as it were, same again. Parents travelling with children cannot easily buy cheaper tickets for them. Each and every individual in a party to the same destination must repeat the whole process again. This will take a considerable time and will provide more opportunities for errors, which in turn will increase the time required to complete their tasks!

The machine would have been easier to use with a same again button, properly integrated with the (existing, but not fully working) previous screen button, so that further tickets could be bought easily, just allowing the user to select whether the new ticket is for an adult or a child.

We understand that features like this, which make it easier for the user, also make fraud easier for thieves. Suppose a stolen credit card is being used: a same again button would permit rapid multiple ticket purchases. From this point of view, it is important to slow down purchases! Losses from just one organised gang amounted to about £25,000 per week, so limiting fraud is a serious design concern.

Changing the ticket machine's user interaction for multiple purchases also impacts on wider issues. The ticket machine must be approved by the Rail Settlement Plan authorities (e.g., regulating how monies are divided between the numerous train operators, for instance as happens with tickets that cover surface and underground travel on different companies). New styles of dispensing multiple tickets is not just a local user interface issue; it could mean seeking new approval; this is a slow and expensive process that has to be offset against the benefits to users (and the increased revenue it may or may not generate).

As in other domains (such as computer security: Adams & Sasse, 1999), the user appears to be viewed as the enemy: the task is made difficult for the many to protect the organisation against the fraudulent behaviour of a few. This entrenched attitude is further exacerbated by the complexity of the change process.

#### *2.5.4 Expensive Errors and Expensive Tickets*

Unlike interfacing to a human operator, mistakes cannot be rectified once a ticket is purchased.

Although the ticket machine knows how ticket prices vary with the time of day, it never tells users that if they wait a few minutes, or if they travel on the next-but-one train, they can buy cheaper tickets – routine advice that human ticket offices give. (Travel before 10am is more expensive because workers have no choice but to travel at those times; tickets are cheaper later in the day when the rail company wants to attract discretionary travel.) We have found that the machine's clock is sometimes wrong (it has been up to 8 minutes slow), and it will therefore sell over-priced tickets. It is tedious for passengers that, although the last train before 10am has left and therefore it is impossible to travel before 10, the ticket machine will happily sell over-priced tickets until it thinks it is after 10.

Some cheaper tickets cannot be bought at all. The ticket machine only sells tickets from the local station. For example, if the user starts travelling before 10am, buying a ticket to an intermediate station and a continuation cheaper ticket from there (for after 10am) to the destination is something human ticket offices provide (saving about £5 on a return trip to London), but cannot be achieved from the machine.

Sometimes a traveller will want to travel as far as they can without exceeding the cash they have available (e.g., to get as close to home as possible when they don't have enough cash). The ticket machine forces a user to choose their destination first. If the destination costs more than the user has, they then have to try other sorts of tickets or other, closer, destinations. This is extremely hard to do with a ticket machine as opposed to a human ticket office (especially as destination stations are organised alphabetically rather than in order along the route).

It might be better if a user could put in as much cash as they can afford for the desired destination, and then ask the machine for cheaper options (e.g., single tickets rather than returns, or for destinations one or two stations short).

The ticket machine knows the time, and gives the user cheaper ticket options when they are apparently permitted. Unfortunately if a train is running late, a user may have bought a ticket at the cheaper rate which is not permitted on the train actually used. The passenger is then subject to a supplement or penalty fare, as well as embarrassment.

In all these cases, we see a particular organisational orientation towards the user's task: the task is framed only as buy a ticket from this station using the next train to a specified destination station, without any recognition that users often frame their tasks in other terms, such as travel from here to there within a reasonable time at low cost, or travel as far as possible in that direction for the available cash. Although simple observational studies in a ticket office or passenger interviews would establish the range of common user tasks, this might be difficult to integrate easily

with the current organisational view, so the user is forced to adapt to the organisational perspective to work at the wrong level.

### 2.5.5 Post-completion Errors

Post-completion errors occur when a user finishes their task before the interaction is complete (Byrne & Bovair, 1997). The classic example occurs in some designs of cash machines (ATMs), where a user, wanting to get cash, inserts their card, types their PIN, collects their cash, and walks away. They have finished but the machine is left holding the card. A simple redesign, widely known, does not give the user their cash until they have retrieved their card: this ensures completing the task completes the interaction.

The ticket machine suffers the same problem. On one occasion we had to wait for a woman to get her ticket, and she walked off with it. We told her to come back and collect her change, which had dropped into the ticket slot the moment after she had turned her back on the machine. This is a standard post-completion error.

The machine itself can deliver change, tickets and receipts in any order; the order used is a software design decision. We understand that one reason for this decision is that return tickets are printed on two separate tickets (the out and return parts), and the designers did not want users to leave part of the ticket behind. Leaving change till last, they must have reasoned, sometimes when there is change! reduces the risk of users leaving without one half of their ticket. Again, trying to improve the user experience at the machine level has implications for the wider design of the entire process: why are return tickets printed on two pieces of paper? Why not simplify?

We interviewed the woman who had bought the ticket. She blamed herself for leaving her change behind. However, her task was to buy a ticket; the ticket machine gave her a ticket, so she picked up the ticket and turned away. If the train she had wanted was already at the station, she would immediately have been running towards it and not even had a chance to hear the change clunk over the noise of the train arriving. She had completed the task she had set out to do (buy a ticket) but the user interface had not finished its part of the task. The ticket machine hardware can drop change before, or at the same time as, the ticket itself.

### 2.5.6 Pre-completion Errors

The converse of a post-completion error might be called a *pre-completion* error: the interface obstructs the user completing the task.

The ticket machine sometimes displays the soft button [PURCHASE\_TICKET], but when it does, it is *not* possible to pay. In fact, the user is supposed to press a button at this point because the ticket machine is otherwise unaware that they are trying to pay (it has no sensors on its coin and credit card slots, which could have been used to avoid the button press). Note, though, that it is difficult for the boundaries of a particular interaction (between one user and the machine) to be clearly identified in a location such as a railway station: so some proximity sensor solutions may not be appropriate.

We understand that this user interface difficulty is caused by regulations: users must confirm before purchasing. One might have thought that trying to pay was tantamount to confirmation but changing the user interface now with financial implications could require a costly new regulatory approval process.

The first time we encountered this problem, we assumed the machine was not accepting credit cards (maybe its communications line had failed) so we hunted around and found a £20 note. By the time we had correctly inserted it into the note slot, the machine had timed out. Despite the user being very active with the machine it times out, because it has no sensors to know the user is still busy with it. (Incidentally, paying by £20 note gets, say, 16 £1 pound coins as change, which takes a substantial time to drop into the ticket slot, and even longer to retrieve as the slot is an awkward pouch without enough room to grab more than a few coins.)

Not only have individual users had these problems, but they can also miss trains because the people in front of them have the same problems, so that the entire queue gets delayed. Again, the user interaction is forced into a particular form by external regulatory conditions that are of no direct concern to the users.

### 2.5.7 *The Machine Does Not Suit Disabled Users*

Users with poor mobility have to get close to the ticket machine, and wheelchair users will find the high screen and high coin slots awkward. Blind or partially sighted users will not be helped that the buttons *always* make a beep, regardless of whether the machine is working or not. (Thus it is pointless rote-learning a sequence of button presses, since there is no way to know whether they are working.)

### 2.5.8 *The Organisational Interface Is Bad Too*

A discussion with a rail employee revealed he had not been trained to use the machine, let alone explain it to anyone else. The ticket machine has an internal user interface when the door is opened. There is an operator's manual, but he hasn't seen it. He has to empty the cash from the ticket machine from time to time; the first time he did this he set off police alarms, since he didn't know how to use the security interface! We understand the machine's manufacturers provide training courses but these have not been widely taken up by rail companies.

One particular point the employee made was that passengers have to *learn* how to use the ticket machine while they are trying desperately to *use* it to get a ticket before the train comes.

## 3 **The Problem Restated**

A computer system has been introduced into a complex system (i.e., the existing railway ticketing system, including its regulations) without taking the opportunity to reappraise the design of the systemic interfaces to the user. The interface design appears to be motivated by providing an easy to use interface to replace the current human interface, and was thus driven by (organisational) functionality rather than by (user) task analysis.

A user interacting with the system still has to contend with all the historical intricacies of the original complex system, as well as coping with (what is for them) a new interface to the GUI itself implemented in a confusing way. In fact, the computer ticket machine has increased the complexity of the interface to the transport system. Fixing the conventional user interface problems that have been created is a new problem, and one that is at or already beyond the limits of the designers' resources (they may not have been funded adequately to do a good job, or they may be unable to deal effectively with the complexity of the design task set them whatever). The overall system goals have got lost in trying to solve the wrong problem at the wrong level, though nevertheless at a level that can be represented as a pushbutton graphical user interface. Unfortunately, the conventional user interface problems are quite interesting in their own right (they are also apolitical), and might divert professionals from tackling the bigger issues of design. It may be that, despite the new interaction problems, the automation of ticketing has obtained some gains in cost-effectiveness for the rail operators: there would then be little incentive to seek better solutions.

New technologies, as we hinted in the introduction, provide new ways of setting the interfaces between users and systems. In this case we shall explore technologies relying on radio communication infrastructure. It is interesting to note that a new infrastructure introduces new stakeholders and the possibility of new investment and new attitudes; it certainly introduces new tradeoffs in design.

### 3.1 *Summary of Concrete Design Problems*

The following points briefly summarise the ticket machine discussion above. In this list we omit topics that could be considered simply as poor design practice, such as inconsistencies and missing help texts:

- The hardware interface is very tightly constrained by concerns of vandalism, robustness and credit card security. We understand some of the hardware components (e.g., credit card verification) cannot be changed.
- The software user interface designers did not consider the users' tasks or their frequency (e.g., the machine does not support buying multiple tickets to the same destination).
- The software user interface encourages users to make errors of well-known and easily avoidable categories (e.g., post-completion errors).

- The software user interface over-serialises: the user has to break down their task into the order the machine specifies rather than an order appropriate for the user's specific task. The machine is not *permissive* (Thimbleby, 2001)
- Rail staff have not been trained.
- The top level screen does not adapt to users who make regular journeys to the same destination.
- The software user interface is misleading for partially-sighted users. (For example, the keys beep regardless of whether they do anything.)
- The interface design increases rail revenue by discounting user needs. (Users who fail using it in the time available but who still travel will be fined; users who buy the wrong ticket can get no refund; and so on.)
- Trying to improve the user experience leads to larger issues and challenging fundamental assumptions on rail travel and regulation.

This list reflects a variety of shortcomings in the design; below, we argue that PTs, together with more rigorous and user-centred design processes, would avoid many of them.

## 4 The Solution

The user's task is to travel; the rail operator's task is to raise revenue from passenger travel. The user can achieve their goals by getting on and off appropriate trains at the appropriate stations; the rail operator requires to know how to bill the user, where they start their journey, where they end it, and some time information (to allow for differential pricing).

Conventional solutions, embedded in the historical necessity of ticket thinking has led to layers of complex interactions between large inflexible (institutionalised) subsystems. For instance, as we saw, eminently justifiable (and very rigid) rules about credit card transactions have a direct and deleterious impact on the user experience. Such rules cannot easily be changed, and certainly cannot be changed within the design brief of the ticket machine design! The crucial problem is that the interface is at the wrong level of the system, and exposes the user to this unnecessary complexity.

Fortunately this level of interaction no longer makes sense when it can be hidden by more appropriate technology for the task.

### 4.1 What Can New Technology Achieve?

We consider the simplest kind of PT: one that simply provides user authentication, and can transmit that authentication data over a distance of, say, 3 metres. With that, we propose a radically different kind of design solution that addresses most of the identified problems. However, as we will show, such a solution introduces new issues: in hiding technical complexity, issues of social complexity emerge as new usability challenges that are masked, or less central, when working with existing technologies.

We assume the user and their PT are registered with the transport operator, with arrangements for billing (e.g. charging to the user's account, or to that of a parent under agreed conditions). There might be agreed limitations on the use of the PT to reduce the risk to the owner of the PT being stolen and used without proper authorisation. We can propose a scenario of use

If there are PT entry channels at stations, the user's starting station can log their entry into the system. At the destination station, the user passes through an exit channel, collecting a receipt showing the journey details and cost, so that they can check the bill when it arrives. If at any point there is a choice of a cheaper or more expensive option, this should be clearly displayed. For example, before the time when fares reduce from peak to off-peak, there should be a clear indication before the user enters the system to show that entering after (say) 9.56am will cost less than entering now. Similarly, if there are alternative train operators offering different services at different rates, there might be clear indications that the passenger is electing to use a more expensive service that will be surcharged. There would have to be additional PT detectors on such trains to bill for the higher grade service; the same principle would apply to those travelling first class. For regular

commuters, there can be discounts (as currently implemented in season tickets, or in other ways, such as rail miles benefits).

Such a solution would address the problems identified earlier, as follows:

- There would be no end-user interface that suffers from problems such as inconsistency of interaction, spelling mistakes, missing help text, etc..
- As noted, the current hardware interface is constrained by concerns of vandalism and security. Radio transmitters can be located in less accessible places and be more easily protected against the weather. They would therefore require less servicing, and no routine servicing (such as repairing breakages, replacing ticket paper, emptying cash boxes). This would remove the need to train station staff in their maintenance.
- With radio communication with the ticket machine, any user interaction could always be completed, so there would be no risk of post-completion errors, provided the user did not have to collect anything physical from the ticket machine (e.g., a hardcopy receipt) at entry: the receipt at exit would only be for the user's records. The user task is defined by the user's actions; a family group can be detected by the system through information stored within the system (on familial relationships and ages of children).
- There would no longer be an issue of over-serialisation of the user task, or over the system adapting to users making regular journeys.
- The system would work as well for blind and partially-sighted users as for anyone else. It would also work for less mobile users who would not need to find and get directly close to the machine itself.
- Users would be charged for the journeys they actually make, and not suffer from buying an incorrect ticket by mistake.

A PT approach is easier to implement and to design thoroughly. A major problem with conventional user interfaces is that programs are serial, but users may work out what they want to do in almost any order (hence the requirement for permissiveness noted above). Often user interfaces are badly implemented, partly because of the combinatorial interaction possibilities. In contrast, the radio protocol for a PT is very simple (above the error correction layers). The main implementation task is maintaining a database of users and billing appropriately — a task that is already performed effectively by most utility (gas, electricity, etc.) companies, and which should therefore be better understood than the current ticketing task.

From the rail operators' perspective, one of the disadvantages of PTs might be a loss of revenue from people making mistakes or buying more expensive tickets than they need to. However, it is likely that this would be more than offset by reduced costs of providing conventional tickets and reduced fraud (particularly fare evasion).

From the user's point of view, such a system presents some new challenges:

- There may be situations where the PT does not work. One might be wishing to treat someone else to a rail outing (for example, a children's outing): a conventional ticket would have to be bought for people not covered by PTs.
- The user may lack confidence in the system — for example, about billing or about whether they are actually authorised to travel — and, without a paper ticket, may feel vulnerable to over-charging or other error.
- Probably most fundamentally, such a solution raises privacy issues: as users are tracked through the system, their current locations may be traceable. Whether the issues are much greater for this use of PT than they are for established technologies such as mobile telephones remains to be seen.

The current ticket machines are intended to address 80% of user needs, leaving the remaining 20% to be dealt with in the traditional way by a clerk. One difficulty a user has is knowing from the outset whether their current task is an 80% one or a 20% one: the user may try to use the machine and eventually give up, or may go straight to the clerk, which defeats the point of providing the machines. If the same approach is taken to PTs, they are not being proposed as the total answer to all travel billing needs, but as a way of dealing with most traveller needs, while leaving some to

the ticket office. The ticket office might sell single use PTs or what look like traditional tickets but with embedded disposable chips.

There are two areas in which the system may break down: failures associated with users, or associated with the technical infrastructure. Currently, the rail companies can cope with passengers who have mislaid their season ticket: such passengers can formally request a refund of the fine (but only a limited number of times over the life of the season ticket). Clearly the non-functional PT excuse will have to be treated at least as leniently. Using network communications, the official should at least be able to check whether the passenger claiming to have left their PT at home in principle does have a PT. In the second case, where the PT infrastructure fails, simple safeguards (e.g., as are used in automatic bank transactions) can ensure the passenger is not charged, is refunded, or obtains compensation, as the case may be.

## ***4.2 Existing Personal Technology Solutions***

There are several trial schemes in operation, often using mobile telephony technology, to test the use of PTs in some of the ways proposed here. One example involves linking a ticket device together with local radio communication and a mobile phone to pay parking fees ([http://www.ericsson.com/mobileinternet/stories/mcom/case\\_5.shtml](http://www.ericsson.com/mobileinternet/stories/mcom/case_5.shtml)).

Although such schemes have been widely reported in popular, marketing and trade press, there has been little study of the usability issues. Rather, the schemes seem very technology-led, with manufacturers and communication network providers keen to demonstrate applications for their technology, rather than improved usability.

## ***4.3 Conclusions about Personal Technologies***

Despite more than thirty years of usability research, the public interactive kiosks that fill the world, exemplified by ticket machines, are difficult to use, causing the public frustration and worse. We have analysed one such state-of-the-art machine to show how fundamental and far reaching the problems can be.

*PTs can solve or avoid conventional usability problems.* We have shown that unsophisticated PT can be used to solve numerous user interface problems.

*PTs highlight complex social problems.* The simplest PT solutions raise privacy problems, which have no obvious solution apart from the public becoming accustomed to the out-weighting benefits of simple and reliable technology. Developers of more complex infrastructures need to be aware of the far-reaching social implications their designs may have.

*On the basis of current implementation standards, simply getting PTs (or anything else) to work reliably remains a problem.* Unfortunately many of the schemes proposed in the PT literature are more complex than the unreliable, unusable ticketing devices we have described. We believe the solution proposed here, based on a basic understanding of the key task of the user, which is to travel economically but reliably from place to place, is simple enough to overcome many implementation difficulties. It may, however, force rail operators to re-design the fares structure to ease implementation.

*PTs will help transform society.* Even within the limited scope of the case study it is clear that many conventional commercial working practices (such as the use of complex fare structures) will become counter-productive as PTs become more widely taken up. Travel, in particular, gives people benefits. Users may be willing to trade the obvious increased usability for the more hidden surveillance by PTs: the PT infrastructure is likely to become much wider than transport, to provide other business or even state benefits (e.g., combining the PT with a national ID card). For international travel, now with heightened security concerns, the tradeoffs are stark. The central and right concern of HCI, valuing the user, clearly now has a very important role to play in leading to a more peaceful and better world.

## **5 Conclusions**

The invisible computer has become a popular slogan, and personal technologies certainly make computers invisible to users. Books have been written on the invisible computer (Norman, 1998) and on the invisible future (Denning, 2002). From our arguments it should be clear that invisible computers are a side-effect of some good design, not a requirement. We might better talk about the

invisible system (or even the invisible implementation of the socio-technical system), for we want to support the user's tasks and experience, with as little irrelevant interference from the inner details and workings of the system as possible. Our approach is much deeper than just making the computer invisible: it means making unnecessary and irrelevant complexity invisible. We support the user's task by defining the system interface, not to control the computer (as a GUI does so well) but to support the operations required by the user. In the case of travel, the user's task is obvious, and—as the case study made clear—personal technologies can directly support the user's task and almost completely hide the intricacies of the rest of the system.

In short, user interface design is not about building interfaces (such as GUIs) to the existing system, or to a computerised version of the existing system (we have argued that the system is *much* larger than the PC that may run the user interface): user interface design is about defining interfaces, in the systems sense, that appropriately support the user's tasks. Doing so may result in a specialised interface that does one or few things well, but therefore is much easier to implement reliably and easier to achieve satisfaction for the user.

### **Acknowledgements**

Harold Thimbleby is a Royal Society-Wolfson Research Merit Award Holder and acknowledges their generous support. We are grateful for some very helpful meetings with ASCOM Autelca of Gmligen-Berne, Switzerland, who manufacture the hardware but not, in this case, its software. This research was partly supported by EPSRC Grant GR/R71467/01.

### **References**

- A. Adams & M. A. Sasse, The user is not the enemy, In *Communications of ACM*. 40—46, Dec. 1999.
- M. D. Byrne & S. Bovair, A Working Memory Model of a Common Procedural Error, *Cognitive Science*, **21**(1), 31—61, 1997.
- K. Cheverst, N. Davies, K. Mitchell, A. Friday & C. Efstratiou, Developing a context-aware electronic tourist guide: Some issues and experiences, *Proc. ACM CHI 2000 (Conference on Human Factors and Computing Systems)*, 17—24, 2000.
- A. D. Christian & B. L. Avery, Digital smart kiosk project, *Proc. ACM CHI 1998 (Conference on Human Factors and Computing Systems)*, 18—23, 1998.
- P. J. Denning, ed., *The Invisible Future*, McGraw Hill, 2002.
- R. M. Friedhoff & M. S. Peercy, *Visual Computing*, Scientific American Library, 2000.
- D. A. Norman, Cognitive Engineering. In D.A. Norman and S.W. Draper, Eds. *User Centered System Design*, Hillsdale NJ: Lawrence Erlbaum, 31—62, 1986.
- D. A. Norman, *The Invisible Computer*, MIT Press, 1998.
- N. A. Stanton & C. Baber, Rewritable routines in human interaction with public technology, *Proc. of ECCS 97: European Conference on Cognitive Science*, ed. E. Hollnagel, 20—25, Paris: EACE, 1997.
- H. W. Thimbleby, Permissive User Interfaces, *International Journal of Human-Computer Studies*, **54**(3):333—350, 2001.