Research Report 314

Multi-Agent Systems: A Risk to Freedom

Iain D Craig

RR314

Multi-agent systems are being proposed for a number of contexts where autonomous agents are required to act in concert with other agents and with other software in order to perform certain tasks. For example, it has been proposed that autonomous software agents be employed to gather information from the Internet: these agents would acess databases and bulletin boards, collecting information that is of potential relevence and sending to other agents who determine its true relevance. There is a clear sense in which agents can be used to restrict civil liberties: one can forsee, for example, agents being used in surveillance of citizens, particularly when television, telephone and computer are connected. Given the increasingly computerised nature of the world and the increasing number of networked systems, such information gathering becomes a real possibility. There are the possibilities for restricting freedom of movement, freedom of choice, freedom of thought. It becomes possible for an unscrupulous agency to tamper with records with the effect that one simply disappears - literally becomes a non-person.

Department of Computer Science
University of Warwick
Coventry CV4 7AL
United Kingdom

September 1996

Multi-Agent Systems: A Risk to Freedom?

Iain D. Craig
Department of Computer Science
Warwick University
Coventry CV4 7AL
UK
email: Iain.Craig@warwick.ac.uk

March 27, 1995

Abstract

Multi-agent systems are being proposed for a number of contexts where autonomous agents are required to act in concert with other agents and with other software in order to perform certain tasks. For example, it has been proposed that autonomous software agents be employed to gather information from the Internet: these agents would access databases and bulletin boards, collecting information that is of potential relevance and sending to other agents who determine its true relevance. There is a clear sense in which agents can be used to restrict civil liberties: one can foresee, for example, agents being used in surveillance of citizens, particularly when television, telephone and computer are connected. Given the increasingly computerised nature of the world and the increasing number of networked systems, such information gathering becomes a real possibility. There are the possibilities for restricting freedom of movement, freedom of choice, freedom of thought. It becomes possible for an unscrupulous agency to tamper with records with the effect that one simply disappears—literally becomes a non-person.

1 Introduction

Technology is usually presented as improving the quality of life. The convergence between computing and telecommunications and the development of the Internet have, as yet, hardly made much of an impact on the lives of most people. If the telecommunications companies and other organisations, equipment, software and services vendors, with interests in this convergence manage to produce the products they are currently talking about at prices that are low enough, the impact of this convergence is potentially enormous. In addition to the basic Internet operations,

there will emerge new companies that will provide services and products which we have not yet even dreamt of.

Already AT&T has produced promotional videos in which customised 'information agents' which will assist the customer in finding out about the best bargains offered by the remote shopping services provided by the network of the future. These agents will also be able to find out interesting things such as where to find services such as car maintenance, insurance, financial advice, and interior decoration. It will be possible to buy a new house without moving from the old one: multi-media presentations will be provided by estate agents, mortgages will be arranged over the net (with or without the help of the personal assistants we have just mentioned).

There is already a slight penetration of 'teleworking'. The teleworker is the worker of the future, it has been claimed. He or she will stay at home to work; they will switch on their home workstation and will go to work electronically. Meetings will be held via electronic means, even if the meeting would have required people to travel half-way round the world; all files will be held on file servers—whether at the office or local. Pollution will be reduced because fewer people will travel to work, thus lowering the demand for cars as well as public transport. Less paper will be consumed, so fewer trees will need to be cut down, and fewer toxic chemicals will be needed to process the wood pulp into paper. Teleworking emphasises service industries, so this form of employment, it is claimed, will become the most naturally important in the electronic society of the future; structural and economic changes might result from the integration of telecommunications and computing.

There is another possible effect. Political systems will be able to change. A big problem with the democracy of today is that its citizens have to elect representatives who then perform the tasks of government *on behalf of* their electorate. Elected government is open to various abuses and is, in any case, far from ideal. Elected members can opt not to respect the wishes of the majority that elected them and can go their own way if they so choose. Citizens of the UK who follow current affairs will have heard local parties in Conservative constituencies complaining that their elected representative, their M. P., does not respect the electors' views on Europe. It works both ways. Sometimes local parties that are anti-Europe complain about pro-European M.P.s; sometimes, the local party is pro-Europe and the M.P. anti-Europe. In each case, it boils down to the same thing: the elected representative does not share the opinions of the majority and has gone his or her own way. There is, it is claimed, a lack of accountability; the issue need not be the future of European integration, it could be something else (higher taxes, increased education, anything will do).

The current system is deeply flawed *precisely* because people are unable to have their say. If we could all participate in parliamentary or senate debates, government would come to a standstill. We are not, they say, even able to influence our elected representatives because communicating with them is so difficult. The new convergence between telecommunications and computing is the answer. The aver-

age citizen can send electronic mail to their representatives and *tell* them what they think. Eventually, elected representatives will become unnecessary because we will all be able to engage in debates; we will govern ourselves in exactly the way that the original Greeks governed their city states.

If we are to be realistic, this picture considerably overstates the case. It is fine as a publicity gimmick, but as a serious picture of the future political system, it fails by being hopelessly optimistic. Anyone who has taught a class of 80 students at a university will know what happens when the students deluge you with email—it shouldn't take too much imagination to figure out what. In my department, we frequently have discussions about policy by email (email can be broadcast to groups of people, remember, it can be open to all). This sounds a good idea in principle. There are sixteen full-time academics in the department. What happens is that nothing is ever decided. Issues are discussed for a short time, and often what amount side-issues come to predominate. A large amount of time is spent reading and posting messages, calming down people who have become worried that they will be discriminated against in some way or another; some people dominate the discussion whereas others say nothing (perhaps feeling that there is no point in their saying anything because it will not affect their life in any positive way in any case—it is a department that operates by making life painful for other people). After a while, traffic begins to reduce; everyone becomes bored with the issue because the discussion is approaching no conclusion and no consensus is being reached. Eventually, the matter is decided in another fashion: the decision is imposed on the department.

There are sixteen full-time academics in the department. Consider what would happen if twenty million people were able to participate in a debate in a way which allowed each and every one to make as many contributions as they like. Government would probably stop completely.

So-called "open government" along these lines will probably remain an ideal rather than a practicality, but there is no doubt that media like the Internet can serve educational purposes extremely well; entertainment—and best of all, the combination of education with entertainment—could also be served. The first step is for the home computer to be networked so that it can access bank accounts, medical records (there have already been experiments in the UK with remote medical diagnosis, not on the Internet, but there is no reason for it not to be), and other important information sources. Following that, there is no reason why the telephone and fax machines should not be added: digitization is already common. Next, the television could be added so that high-definition images can be received from home. Videophones might replace the voice telephone so we can have images as well as sounds. The intelligent building is another new topic: the home could be monitored and controlled using a home computer. The vision of the future that some would have us believe is one in which home computers control everything in the house and also give us extensive access to the outside world: they could do our accounts, make

regular purchases, take over the role of direct debit and authorise other payments. In short, the home computer takes over many of those dreary tasks that bother us so much *as well as* giving us access to an entire universe outside the home.

2 Intelligent Agents

What is an agent?

The *Shorter Oxford English Dictionary* defines an agent, amongst other things as:

> *n.* **1** A person who or thing which produces an effect:
>
> **2** A person who acts for another in business, politics, etc.

These definitions will serve present purposes very well: I do not want to get into the technical details of agent-oriented programs or of multi-agent systems. An agent is an entity that acts on behalf of some other agency or entity and is able to alter the world in various and appropriate ways. What is needed is some explanation of how such effects are produced, but I would prefer to leave that for other places, and rely upon everyone's intuitions about when and why to act.

What happens when intelligent software agents are released on the global information network? As we have already noted, companies like AT&T are already investigating the roles of such agents; multi-agent systems is one of the fastest growing research areas in Artificial Intelligence and Computer Science. If intelligent agents that are able autonomously to make decisions for us, the increase in power we can expect from computers will be enormous. Instead of telling the machine of every single action we want it to take, it will be able to do these things for itself. For example, an agent might monitor the recording quality on the VCR and order new tapes when their quality reduces below some threshold. Another agent might notice that the house insurance will soon expire and arranges for either a renewal or a new policy (one that is cheaper or provides better cover for the same cost, or perhaps for a little more). Another agent might determine one's nutritional needs for the coming week, figure out what we are to eat and then order the food from the on-line supermarket.

It might be countered that something like this already happens in the dealing rooms of the world's financial markets. Here, programs routinely trade in stocks, commodities, and so on. These programs are, though, entirely driven by values. When the value of a stock falls below some threshold value, a program might buy it; conversely, when the value drives above some value, the program will sell. This is not particularly 'intelligent' behaviour, but the mindless testing of two cases, and performance of the enabled action:

- If the price, v, of stock, s, is such that $v < \theta_1$, buy s.

- If $v > \theta_2$, sell s.

The rules are extremely simple (but see below). It is worth noting that another rule that might be used is the following:

> If v drops below a special threshold so that $v < \theta_d$, then sell all of s because it is worthless. (A rule for dumping worthless stock.)

To make any of these rules viable, it seems, some form of trend analysis would have to be employed. However, we are only presenting an argument, not attempting an accurate account.

Even though agents are not necessarily particularly powerful cognitively, they could improve upon this simplistic behaviour. For example, an agent could change the value of its θ_1 and θ_2 thresholds as a function of other values. An agent could also perform some trend analysis and notice that the price always rises on Thursday mornings; this analysis will not uncover the reasons *why* this happens. (Perhaps the company's weekly stock of raw materials runs out by mid-morning on Thursday; perhaps the owner of the company always buys his stock—or buys his food for the next week—on Thursday. This might sound silly, but stranger things have influenced stock prices). It doesn't matter *what* makes the price change, but the agent will discover this movement and will alter its behaviour so that it doesn't dispose of s on Thursdays. A more intelligent agent might take into account the prices of the other stock in its portfolio: in such a case, its goal might be to maximise the profit over a more extended period of time, so it could afford to ignore some price fluctuations.

The kinds of behaviour I have just suggested as alternatives to the extremely simple one that is exhibited by current dealing software could be characterised by the term 'rationality.' These agents perform rational acts, acts which can be explained and which can be described as having a rational basis. One reason why they are rational is that they respond to the situation at hand. The last agent mentioned might, in order to ensure that it maximises its profit, might attempt to exploit its best growing stock (say by buying more and, perhaps, borrowing money to help it do so) after it has suffered a modest loss; it will exploit strategies that support the achievement of goals over long(ish) periods of time rather than simply selling because the price has exceeded some *a priori*, and probably totally artificial, number.

Banking and finance are, naturally enough, deeply concerned with money; the cost of anything must be justified thoroughly. Why should a bank or dealing house engage in agent-based technology when current generation dealing software (which is supposed to be costly and complicated enough) seems to work very well? Will not agent-based technology be expensive, more complex and probably more unreliable than the software that is currently available? It must immediately be asked whether current software is reliable: there are measures of correctness against which it has not been compared—it is 'reliable' because it hasn't broken yet (this is a useful definition of correctness that applies to most currently available software). Perhaps agent-oriented software will be more complex, but it exhibits rational behaviour and

should be able to explain itself, at least somehow. An agent that applies the long-term strategy mentioned above will be more complex than a routine that applies the simple rules given above. When questioned, a strategy-following agent should be able to produce some kind of understandable explanation as to what it has been doing. In this case, it could point to its goal of maximising profit over the long term and to the loss that it has suffered; then it could point to the strategy that it has adopted and point out that this strategy should lead to profits in the long run. (None of this is impossible, even with current technology.)

However, there is a significant problem with the simple rules I gave above. They can lead to *chaotic behaviour*. Many systems that exhibit complex dynamics are based upon simple rules. Rules such as the three that cause the buying and selling of stock based upon thresholds could easily lead to chaotic dynamics. Some would argue that the stock market collapses of the late 1980s and early 1990s have been due, at least in part, to the chaotic properties of systems based on these rules, or at least on rules very much like them. The interactions between systems that employ these rules can be highly complex; in addition, the possibility of explaining behaviour in fine detail is remote if not impossible. Some might argue that the very simplicity of the rules is all that we need in order to explain behaviour, but this, given these arguments, seems a rather slender argument.

Dealing programs show no rationality. A rational agent will analyse the situation and will decide how to act and when to act on the basis of available evidence *together with* other factors[1]. In many cases, these additional factors will be of as much importance as the information available from the current situation. A rational agent can decide *not* to act if it so decides; rational agents have reasons for acting or not—they can report these reasons to others when necessary. Agents, it would appear, are much better than conventional software. Systems composed of agents, together with conventional software, are almost cerainly much more powerful than conventional software alone. There are reasons other than rationality for supposing this: for example, *open systems*, systems whose components can vary with time in number and function, are more easily constructed using agent-oriented technology because it is easier to connect agents—get them to talk to each other—than is the case with conventional software. Systems built from agents can also exploit the potential for increased speed due to the fact that agents run in parallel with each other, so more than one thing (job, task) can be done at any time. Agent-based software also allows the prospect of extremely reliable systems because agents can *monitor* other agents and, if they detect errors or deviant behaviour, they can take appropriate action to correct matters.

Agents are a good idea. They allow us to build flexible software that is more responsive and responsible than the conventional stuff.

[1] Of course, the recent collapse of Baring's bank was caused by rational agents. Perfection is only for the religious or the naive.

3 Future Negative?

Are agents *necessarily* a good idea? Is the Internet *necessarily* a wonderful idea?

Let us consider a world we described above: a world in which cards with magnetic stripes have becom rare: smart cards are more common. A world in which telecommunications are integrated with home computers so you can surf across TV channels on your PC and also surf on the (expanded) Internet. A world in which homes are smart buildings. A world in which your work is done at a workstation. A world in which your car's computer talks to the building control computer and to the Internet machine. A world not too far away from ours, but far enough to make it a bit different: this is a world in which computers are as connected as telephones and much more common. I don't find this particularly far-fetched or futuristic.

Given this world, the purpose of this section is to ask what sorts of abuses are possible and by whom. A second purpose of the section is to ask how abuses of civil liberties can be enhanced by the use of agent-based technology. Some of the ways I can see of doing this are already possible, provided things are connected over relatively accessible (i.e., interconnected) networks; what I hope to indicate is that matters become very much worse when agents are used. The main areas in which I, after a few minutes' thought, can see potential for abuse can be roughly characterised as:

1. Financial.
2. Medical.
3. Domestic and censorship.

Under 'Domestic and censorship,' I include such things as movement control and the regulation of what is seen on television. I will consider each area in turn. I originally had a section on political abuses, but it became clear that most abuses could be put to political ends.

Financial

These days, a person is worth more in financial terms than the amount of money they have in the bank: they have a credit rating. Credit ratings determine how "good" a risk they are. Banks are already interconnected to some extent. In the future, greater interconnection can be expected. This year saw the introduction of an automatic cheque clearing system for the UK clearing banks: this system arranges for the automatic transfer of money between banks on the basis of the movement of cheques. Hacking into individual bank accounts has gone on for years, so one's own account information is not as secure as one might wish. If banks are more highly interconnected, it becomes easier for individual accounts to become available to other agencies than the bank, and it becomes easier to access other parts of the bank's operation.

Interconnection already allows credit card sales to be, in principle, checked against credit worthiness. Checking against bank account balance is also possible and somewhat easier today than it was, say, fifteen years ago. This information, once it is in the possession of a third party, can be used in a great many ways. It is already the case that credit card companies sell information about their clients to third parties. For example, credit card companies keep records of the spending patterns of their clients. This information is sold to vendors; it could also be sold to market research organisations; it could also be sold to political parties. Vendors might be interested because people who tend to buy a certain kind of item might be persuaded to become customers. Market research organisations want to know about credit ratings and spending patterns for reasons that should be clear. Less clear are the reasons why political parties should be interested.

First, it needs to be remembered that the aim of any political party is power: its purpose is to become the governing party. If it can obtain information about the populace whom it might govern, it can, at the very least, shape policy. It can also target campaigning activities on potentially fruitful sectors and areas. It is a waste of effort trying to win over clear supporters of the opposition and clear supporters of one's own side: effort should be expended on those who are likely to be won over.

Governments can become interested in matters such as credit rating and spending patterns. The latter can provide information about political affiliation, whether direct or indirect. Direct affiliation would be evidenced by payments to a political party: indirect by payments to an organisation that is associated with a political organisation either overtly or covertly. Membership of CND during the 1980s was taken by the authorities as evidence of affiliation to a subversive organisation; donations to Irish charities might also be construed in a similar way. Membership of a political organisation that is considered to be unfriendly can also be a sign of a subversion. Governments have a duty to protect their people, so gathering information of this kind is justified in terms of the national good, the good of the people. In a similar fashion, drug trafficking poses a threat to society, so it must be discovered and stamped out. There are many examples of criminal activity whose detection could be aided by monitoring of the kind considered here. Advertising on, for example, healthy eating is best aimed at those who do not eat in what is considered to be a healthy fashion. With the convergence of telecommunications, computers and broadcasting, pin-point aiming becomes possible, of course.

Information about purchasing can lead to inferences about political affiliation as well: those who buy environmentally-friendly and non-exploitative goods possibly hold political opinions that are either to be suppressed, ridiculed, tolerated or exploited. Lifestyle information would also be of interest to insurance companies.

Information about the purchase of goods can lead to inferences about the kind of person one is. It can also be of interest in gathering detailed information about how, for example, a national trade deficit is being created (some might remember the "I'm backing Britain" campaign of the 1960s). An example of the opposite effect is that

of convincing the population at large that some industrial section produces goods that are bad or unacceptable (one way to do this is to use misleading statistics). One reason for doing this would be to destroy, for example, a trade union, just as the National Union of Mineworkers was destroyed by slowly closing down Britain's coal mines. The Union had to be destroyed, it could be argued, because of its strike in the mid-1980s, a strike which massively opposed the rule of the governing party and its leader; the strike threatened to destabilise government by threatening essential services such as electricity generation.

There have been reports of credit rating information being divulged to persons who are not employees of the credit card companies or the banks. For example, institutions in the business of giving loans might be interested in people who are in trouble keeping up with their card payments (and with other loans, of course). People who are having difficulties with their bank accounts might also be of interest to such companies.

Information about financial matters is important, but so is control. If it could be shown that someone had a bad record in clearing debt or had a record of debt, that person could be refused employment, could be implicated in crime, or could be kept under surveillance. Falsifying records is just as possible as inspecting or duplicating them if easy access is possible. Falsification of financial records could destroy a career: for example, a politician's, a banker's, or a police officer's career. Falsification of accounts could discredit charities or political parties, and it can cause companies to become insolvent.

An interesting property of the falsification of electronic records is that, unless precautions are taken (e.g., extensive backups), falsification can lead to alterations that come to represent the truth. For example, if someone falsifies a bank account's details and no-one notices until some considerable time after the act of falsification, the state of the account represented by the falsified figures becomes the *actual* state because of the lack of reliable audit trails. If the period is sufficiently long, it can be the case that there is *no way* in which a return to the original state is possible: that state cannot be regained because there is no way to restore it and then apply all the transactions that have occurred. In some cases, it will even be possible to falsify the backups (e.g., if they are also stored on disk on machines that are currently networked or, in order to perform the backup, connect to the network—all one need do is wait for them to connect and then perform the false updates).

Medical

We expect our medical records to be confidential. Our medical history is important to such matters as employment, earnings potential, and insurance status. Imagine that you have been diagnosed as having a long-term illness, a fact which is bad enough in itself. How would you feel if, the following day, you were approached by a representative of the company that insures your life or house who wants to invalidate your policy or to charge you an increased premium because of some statistical con-

nection between your illness and something else. How would you feel if you started to receive junk mail from drug companies that tried to get you to persuade your doctor to prescribe their products and to get you to buy other products made by them?[2]

These examples could be dismissed as unethical conduct on the part of the insurance of drug company. It could be countered that insurance companies have a significant interest in knowing one's state of health; this might be a strong argument for a medical insurance company like BUPA. Drug companies could argue that detailed information on the spread of a disease, detailed epidemiological information could be of critical importance in understanding the nature of a disease and of its vector. This argument is almost undoubtedly true, but the response by the company and the *precise* nature of the information it gathers should be of concern to us. Health authorities need information about medical matters in order to plan health care provision. Of concern also is the attitude of the medical practitioner who provides the information: it is, of course, conceivable that the information is clandestinely obtained.

These cases might be considered not to be particularly severe. There are clear cases in which *unauthorised access to raw medical records*, and it is this that concerns us, can pose a threat to the individual. Employment might be denied if a potential employer were informed of a previous illness or medical intervention. One might become a social outcast if knowledge of an illness or hospitalisation became known. Consider the cases of mental illness, abortion, sexually transmitted diseases. All of these have been cited recently as causes for concern. A cabinet minister recently had to defend herself when it became known that she had had an abortion (in the particular case, she was criticised for hypocrisy). Insurance companies are trying to introduce special clauses about HIV infection: the risk of HIV infection is clearly of significance to surgeons and dentists. Mental illness raises horrific spectres and can pose a real barrier to employment: employers claim that they have to protect their workers and have to ensure that the prospective employee will not become ill while at work: there is an erroneous image of mental illness and violence being closely associated.

Careers have been destroyed because of revelations about previous medical history. The cabinet minister mentioned above did not fall from office, but others have. There are certain medical conditions that society considers to be signs of degeneracy, anti-social behaviour, or taboo. These attitudes can be engendered by government, by pressure groups or both. It has been argued that view of HIV as a new plague that threatens the extinction of mankind (a view that was supported by quite a few during the 1980s) was a propaganda exercise undertaken by political groups in order to change the populations' sexual habits (and discredit homosexual political groups as a side effect). Consider the case of abortion in the USA at present: the

[2]It could be worse. Imagine calling out your GP because you are incapacitated by a gastro-intestinal infection and the next day being approached by a bicycle clip manufacturer!

moral 'majority' is killing abortionists and turning abortion—abortion for any reason, including medical—into a new taboo. The sanctity of life has been taken over by a pressure group for political ends.

Worse than disclosing medical information is falsification. It is one thing to be exposed (if that is the right term) as having suffered from some illness, or even from suffering an illness: it is quite another to be falsely accused of suffering from it. In some cases, one might not want anyone to know that one is ill: share prices have fallen because of the medical state of a chief executive, governments have been forced out of office because of a leader's illness. As we noted above, with electronically stored records, it is possible for falsehood eventually to become the authenticated truth.

Some infections are considered to be so severe that sufferers and carriers are isolated from society until they recover or die. If someone's records were falsified so that they recorded that they suffer from some highly infectious disease, it would be possible for them to be held captive for an almost unlimited time. Infectious diseases are particularly good in this respect: in the UK, we have *habeas corpus* to control imprisonment, and we have the various mental health acts to control detention in psychiatric hospitals, but there is nothing, as far as I am aware, that controls the length of stay in an isolation ward. Habeas corpus could be invoked (as it is in the play *Who's Life Is It Anyway?*), but if the disease is sufficiently rare, and if the medical records have been falsified with an appropriate degree of care, who would be in a position to gainsay the medical authorities detaining the 'patient'?

In a way similar to medical records, criminal, social security and employment records could be obtained and falsified. Consequences of varying severity would ensue from disclosure or falsification. For reasons of space and for fear of being boring, I will not spend time on them and leave them to the reader's imagination.

Domestic and censorship

If the super-highway becomes part of the home, we can expect abuses there, too.

We noted above that a legitimate (at least, not a completely illegal) use of financial information was to determine the purchasing habits of individuals. This allows marketing organisations to target their activities and it also helps manufacturers in improving their products. Additional information can be gathered from the viewing habits of TV viewers or the listening habits of the radio audience. Correlations between the broadcasting of particular commercials and increases (and decreases) in purchases could be made more accurate. People tend to watch commercials that deal with products they intend to buy, so knowledge of this could lead to sales.

Demographic information could also be obtained by monitoring what is watched on TV and VCR, what is listened to on radio and CD. For example, in the case of video tape and CDs, it would be possible to correlate the number of purchases of either tapes or CD with the number of times they are played (perhaps the time of day when played could be correlated with other information that could be gathered

from monitoring the home). Whether one would *really* want to be monitored while listening to a CD or watching an opera on TV is another issue—I suspect that many people, myself included, would find this an invasion of privacy even though, to a certain extent it is harmless.

The phrase "to a certain extent" is key. The fact that I bought a copy of *La Bohème* some weeks ago does not necessarily mean that I particularly like Puccini, Romantic Italian opera, or the singers appearing in the recording; I have bought other CDs on the same label. I would not like to be bombarded with information about any of the above. If I want the CD publisher's catalogue, I will order one: I do not necessarily want one sent to me automatically, nor do I want my intelligent TV to show lots of advertisements for the label.

Television viewing can reveal a lot in addition to the commercial information just described. People can reveal things about themselves by their viewing or listening habits. For example, tuning into a channel that has a high proportion of programmes about environmental matters could lead to the conclusion that the viewer is concerned about the state of the planet. A viewer who changes channel when a particular politician appears can be inferred as having political views. Someone who often watches programmes with a certain political slant can be inferred as being sympathetic to that view. On their own, these items of information might not add up to much, but when integrated with information from other sources, a picture could emerge.

A further refinement comes when the telephone is used to eavesdrop; microphones on other equipment, for example camcorders or tape recorders, could also be used if the equipment is connected to the net. If this is done, conversations can be monitored and evidence of opinions, intent, etc., can be obtained directly from the horse's mouth, so to speak. In a similar fashion, email can be monitored, as can telephone and fax communications; network user group subscriptions can be obtained and investigated. This monitoring need not be restricted to the home: monitoring at work is also possible, and for the same reasons and using the same techniques.

Monitoring is not the only thing that can be done. Viewing and listening can be censored. Contacts via telephone, fax and email can become unavailable. A smart television can be instructed to fault or to switch off, or can be instructed to show an alternative programme to the one selected.

Censorship could be extended to any form of electronic communication. Telephone calls could be stopped; certain telephone or fax numbers could become permanently unobtainable; email messages could bounce or fail to be delivered; web pages could become unobtainable or could direct the user somewhere else. With increased intelligence in a system, any of these could be performed with ease. Furthermore, because the form of each of these types of communication is identical—a bit stream—it is possible to store communications for later analysis, it is even possible to transform them in various ways.

Next, imagine that domestic alarm systems are connected. In addition to the integrated telecommunications and computing system, information is processed about movements within a house, and when the people who live there are at home. It is also conceivable that such a system could monitor energy consumption within the building. Such a system could provide valuable information: it could inform a third-party of when the owners are out. Such information could be used to plan clandestine access to the property (say for planting additional bugs, replacing failed or failing equipment). Information of this kind could also be gathered from sensors on other equipment (e.g., the telephone). Information about energy consumption could be used to regulate the amount used: a "green police" could enforce consumption limits. Alternatively, users could be encouraged to use more energy as a result of monitoring their consumption patterns.

It might seem that control of population movement is pretty far away from agent-based systems and the concept of the Internet. Unfortunately, it is not as the following examples suggest.

Assume that smart cards have replaced credit cards. The credit card companies are already examining this option to avoid fraud: smart cards are claimed to be less prone to fraudulent use than are conventional cards with magnetic stripes. One suggestion in the UK parliament recently was that the national identity card (which some politicians favour) could be combined with a social security card, driving license, and all the credit and charge cards that one uses; it would represent a multi-function card. For the first couple of examples that follow, the fact that the smart card also functions as a national government identity card doesn't matter; its use as money is adequate for our purposes.

Consider toll roads. The UK government is trying to introduce them as an alternative to publicly funded roads and claims that they will reduce congestion. If main roads are toll roads, motorists must pay to use them. Now, if a smart card has to be used to pay for using the road, information about the journey can be recorded on the card. It is necessary, in order to compute the toll, to know when the motorist joined and left the road; the time between these two events can also be recorded. (Use of a smart card to purchase all services along a toll road makes sense because everything could then be charged automatically to a central agency: this seems to make for increased efficiency.) If payments at service stations along the way are required to be made using smart cards, the points at which the motorist stopped will also be recorded. Unless motorists can stop at arbitrary points on the road (which is something that the authorities try to prevent for reasons such as safety), it is possible to determine their movements along the road, their average speed (so they could be fined for speeding or dawdling), as the point where they joined the road and the point where they left. If the network of toll roads is extensive, it would be possible to track someone's movements across the country.

One proposal that has been put forward by the UK government is to require motorists to fit transponders to their vehicles. When the vehicle enters a toll road,

it passes over a strip that exchanges data with the transponder and records the fact that the vehicle has passed on to the road, the time and the point of entry. The idea is that this device would be easier to use than requiring motorists to obtain a ticket or some kind of token from a booth at the point of joining. The use of a transponder would involve equipping the road and requiring all vehicles to be fitted; the costs could be extremely high (but governments have done unpopular things in the past and have survived). A smart card is a cheaper way of doing things (an on-board computer could also be made to do this, but on-board machines might also be expensive). It is not even necessary for the card to be inserted into an interface machine. A signal could be emitted by a strip under the road when the vehicle passes over it (a simple inductance loop will do); the signal is received by the card which then records the relevant information. At the other end, payment has to be made. Charging motorists at the end of each month is one option, but payment at the roadside via the smart card is also possible.

It has also been proposed that vehicles entering city centre areas should be charged. Devices similar to that just described have been proposed for this purpose.

A practical question: without transponders on vehicles, how does the system know how much to charge? A smart card on its own cannot determine whether it is in a car or a truck; furthermore, what happens if the card is carried by a passenger on a bus? These problems are not insuperable, however. If the system is concerned with knowing where people are going, there is no problem: the problem is the minor one of determining what sort of vehicle is being used. Perhaps vehicle operators would be required to issue a card for each vehicle; personal data about the driver would then be copied onto the card—what happens if the driver forgets and has the wrong card? Perhaps we can assume that the penalties for carrying the wrong card would be sufficiently high to make this a rare case; perhaps a technological solution can be found. It matters not whether the actual device is a smart card, a transponder, a computer, or a combination of them; the argument rests upon the *ability* to track any vehicle at any time and with no additional effort.

Imagine what could happen if there were an agent associated with your national identity card. This is where things can take on an additional level. Telephone companies have historically had an intimate relationship with central government. They have to provide and maintain telecommunications links with secret installations and have to ensure security for all information communicated by government agencies. In addition, telephone and telex taps have to be applied at the request of government agencies and the law enforcement agencies.

Now suppose that a government wanted to know of people's movements down to the buildings that they frequent. Most buildings are equipped with telephones. In public buildings, one does not always pay to use telephones, nor does one pay in cash or by card each time one uses the phone at home. It is possible, though, for smart cards to record information about their whereabouts in response to signals sent to them through either the telephone network or through some device "attached" to

telecommunications cables (i.e., attached to every building in a clandestine fashion). (It is, of course, possible to legislate the every building has an identifying device fitted, but this would need careful justification.) The point about the national identity card is that it would be compulsory to carry the card. Every time someone enters and leaves a building of any sort, the card can be triggered to record the fact.

It is possible, and this should be noted, for concerns other than central government to be interested in movements. These days, most telecommunications agencies are privately owned and, presumably, open to hire by anyone who can pay the fee: all that is required is enough money to buy the services on offer. Anyone who is sufficiently motivated and who can pay could track movements. This would require legislation by governments, but the necessary laws are either already in place or else being contemplated.

The above has been concerned with *monitoring* the movements of the population. What happens if *controlling* movements is the aim? This poses a few more problems; these problems concern the physical control of movement. In some cases, such control is very easy. For example, air travel could be prevented with relative ease (smart cards on passports, check-ins reporting passport numbers to other places, etc.). If there was a need to prevent someone from travelling by train or bus, it could be done with relative ease. If payment of fare has to be made using a smart card, or, in the case of rail travel, cheques have to be guaranteed using a smart card in place of conventional guarantee cards, it is an easy matter to arrange for the card to be rejected or, as has been discussed, for the would-be traveller's account to claim that it is empty or overdrawn. Travel without a ticket is not permitted, so the person is unable to go to another place; furthermore, it could be made a crime to use the card in this way, so the police could be called and the person detained.

A similar approach could be adopted at services along toll roads or at filling stations. These methods, however, do nothing to prevent movements by motor vehicles with full fuel tanks and by other means (we will ignore bicycles and small boats; travel via ports could be monitored and controlled with some ease). Motor vehicles have to cross intersections, pass traffic lights, cross level crossings, pass pedestrian crossing controllers, even pass telegraph poles and street lamps. Each and any of these could be fitted with transponders. Detection and tracking do not constitute the real problem: the real problem is prevention of movements and this depends upon manpower.

If it were desired to determine the identities of those people in a particular area, the following approach could be adopted. It would be relatively simple to emit signals that are detected by the smart cards (national identity cards perhaps) in a given area; this could be done by sending signal over the telecommunications system. The smart cards then transmit signals to the nearest transponder to convey their current location; alternatively, the signal could be stored until a later date and transmitted when the card is next used (innocently used). Remember that the cards can process information, so they could record the triggering signal and encode

it in a particular fashion. If the telecommunications system is used, each location (dwelling, hotel room, etc.) would presumably have a unique identifier which could be stored by the card or could be transmitted to some other place at the request of the card. The upshot is that the location of all cards can be determined.

3.1 Agent-oriented Programs

Above, we have considered various possibilities for invading people's privacy and denying them liberties. These possibilities could become realities if there is an easily accessible convergence between telecommunications and computing and if there is greater interconnection. It is already possible to tap telephones invisibly: there is no need to go to an exchange and physically connect to the circuit. Instead, a remote computer can instruct the exchange to send it all calls from the specified number. In the past, it was also necessary for internal security organisations to have large regional and central organisations so that the necessary staff and archives could be housed. The need for large archives is already a thing of the past because records can be put on computer. However, a large staff can still be required to correlate and interpret information.

The East German secret polic had an enormous staff and not only because it needed to maintain paper archives. It needed a large staff to correlate information, interpret bugs, telephone and telex intercepts, follow people, burgle their homes and offices. The East Germans were monitoring *everyone* from foreign visitors to ordinary citizens; similar monitoring took place in Bulgaria and other former soviet-bloc states.

The amount of effort needed to keep the entire population of a country under surveillance is enormous. With the interconnected systems we have considered, the information gathering effort becomes is reduced because it is automated and, for the most part, is already in the best position to gather data. Data gathering is what I have concentrated on when considering the various possibilities because it is the easiest to describe. Furthermore, it is necessary to convince people that such information gathering is *feasible* before addressing the next stage. The next stage is the integration of information from the many possible sources into a coherent global picture. Previously, this has required human intervention: in the near future, parts of the task will be automated. It would be unwise to suggest that the entire task will be automated for, presumably, checks will be needed.

Agents are a relatively inexpensive technique. Once an agent has been constructed, it can be copied as many times as one wants and at almost no cost. The cost of deploying twenty million agents of a particular type, say an agent that interacts with smart identity cards, is extremely small; even the cost of constructing the agent will be relatively small compared with other potential costs. Next, agents can be active all day, every day, twenty-four hours each day, seven days a week, fifty-two weeks each year: they never tire and their accuracy remains the same (unless there is a software fault). Third, agents can be made undetectable, or at least can

be made to be very hard to detect. Remember that we are considering *real-time* systems that have to be running all the time. When a conventional virus hits a PC, the PC has to be halted and its disks sterilised; this cannot happen in a real-time system because down time cannot be tolerated. The real-time nature of the systems and their scale make the task of finding an agent very difficult. Furthermore, if an agent is monitoring the programmes you watch on television, that agent could be on a machine that is physically remote from the one that controls your television set. The agent could be at the television station or on a satellite: what it does is record the fact that you have selected the channel, that you have selected the channel on your decoder; similarly for cable. (Terrestrial television can be controlled more easily, so its output could be arranged to be bland and uncontroversial.)

The justifications for agents given above can also be given for other kinds of software. There are differences, though, of a purely technical nature. For example, an agent-oriented program is an *open* program; that is, its components can vary with time, so it is easy to introduce or to remove components. This flexibility implies that an agent-oriented program can be dynamically structured to respond to the needs of a particular situation or phase of a problem. As resources are required, they can be added (up to a maximum, for resources are never infinite): resources can be local or physically remote. This ability of an agent-oriented program to distribute resources over different processors in different locations can also contribute to the difficulty in detecting them. Because they are naturally distributed, agent-oriented programs can make use of considerable resources, with only a small load being recorded at any site. (Of course, agents communicate with each other: this might prove to be a way of detecting their presence.) The distributed nature of an agent-oriented program makes it more secure: there is no single place that records all important information, and there is no single place in which the state of the system is stored (state information can be useful if one wants to clobber a system).

A significant justification for the adoption of agents is that they are more 'intelligent' than conventional software. An object-oriented program is a program that is built from abstractions called 'objects'. Each object represents a kind of entity with which the program deals. Objects are a way of structuring programs, but they are not the same as agents. Agents can be described in terms of mental states and intentions as well as content and action; objects can only be described in terms of their content and possible actions. Furthermore, agents are able to communicate using high-level languages, whereas objects communicate at the level of program structures (and do not make inferences about communications).

In the last section, I explained how an agent-based share dealing system could improve upon the conventional version. The ability of an agent-based system to engage in rational behaviour and to perform 'intelligent' tasks (some of which we will mention below). Many of the activities I have mentioned in this section require the rational integration of information from a variety of sources before a coherent picture can emerge: someone is not a subversive because they want to see the

eradication of nuclear weapons, what makes them a subversive is their affiliation to subversive groups *and* their participation in acts of a particular nature. Participating in public demonstrations might be seen as further evidence, but it does not necessarily demonstrate intent. Of course, participating in demonstrations might be ample evidence to some, but to others, the final verdict would only be pronounced after more information has been adduced. This is the kind of reasoning that agents can perform: they can integrate information from diverse sources and reason about it.

Transcription of telephone traffic is a difficult task. The speech has to be converted into a written form; next, someone has to read the transcript and evaluate it. If the digitised speech train that is processed in digital exchanges, or sent across so-called 'value-added' networks, can be analysed in real time for interesting content (e.g., the presence of keywords), time and money have been saved (the message can still be archived for future reference and/or processing); if a telephone call can be transcribed automatically (perhaps not perfectly, but the sound and text can be presented synchronously on the transcriber/editor's screen to make correction easier), time is saved. If patterns in someone's movements can be detected or confirmed given clues already given by other information sources, predictions can be made about their future conduct. If there are many people who are subject to such surveillance, automation makes the task far easier than having someone following them. Similar trends could be detected in purchases and other shopping patterns.

In addition to mere trend detection, inferences can be made about the person targeted by the agents. One particularly useful kind of inference would be the assessment of the threat or risk the person poses. A problem with this issue is that traditional methods are incapable of sifting through all the information quickly enough to ensure that the latest and most accurate assessment is available. With agent-oriented techniques, this can be done.

Instead of just processing information, new things can be learned and justified with reference to what is already known. Inference can be employed to make explicit what has previously been implicit, and it can be used in making predictions. This is a further characteristic of agent-oriented programs, or can be, and it makes them more powerful than conventional software.

Inferences can be made to form, justify or reject hypotheses. For example, if an agent has information that someone subscribes to a particular electronic newspaper and communicates with people who are known to be associated with someone of particular interest, an agent could make inferences about the nature of their relationship with this person, and it could make inferences about the person on the basis of what it knows about the interesting person. It could also arrange for further action to be taken either to gather information or to detain those involved.

In a similar fashion, knowledge of which newsgroups someone subscribes to and knowledge of their spending habits and available funds might be useful in deciding whether they would like to receive publicity information about your latest product.

The complexity of the inference-formation processes required to perform many of the actions described above is quite high. These are typically non-trivial tasks. By spreading the load across agents whose job it is to form hypotheses about particular aspects of the problem, or to process particular kinds of information, or to transduce information from one form to another, the complexity of the problem can be factored into more manageable amounts. Furthermore, as noted above, the processing demands of a task can vary with time and agent-oriented programs are very well-suited to this. Agents could be given the task of monitoring diseases, when one of them finds an instance of an interesting disease, it collaborates with other agents to analyse the lifestyle of the person suffering from it. This analysis could propagate outwards to insurance companies or to other organisations and agencies. The process of propagating from a single agent to a collection of many agents is an example of the dynamic nature of this kind of system.

3.2 More Thoughts

Before we end, it is worth pointing out that the abuses that I have discussed above have been couched in terms of the individual. It is possible for many of the abuses to be directed at commercial organisations, even governments. The rights of the individual concern us most for they are the most precious and open to abuse by those in power (who probably have commercial interests anyway).

Furthermore, I have chosen to concentrate on the kinds of abuse that national governments could engage in because they are the most obvious sponsors of such activities. Governments have long histories of black propaganda, dis- and misinformation, blackmail and abuse of civil liberties; they spy on each other and, almost paranoid, they spy on themselves. Governments already have the apparatus necessary to do the things I have described above: they might not have the technical infrastructure in place at the moment, but they have the means and the money. However, it is just as possible to replace the phrases "brand loyalty" or "company loyalty" where I have use terms referring to political belief or affiliation. In some cases, the analogies become far-fetched: for example, it does not seem reasonable to expect companies to monitor and control population movements (they might have an interest in demography, but that is a different issue). The relationship between business and government can be close, however (consider the recent revelations in the UK concerning newly privatised utilities employing the ministers who privatised them).

4 Conclusions

Above, I have outlined some ways in which the massive interconnection of computers coupled with the telecommunications system could be subverted to various illegal and semi-legal ends. The basic technology I have assumed is not too extreme: it

amounts to the kinds of equipment that we can expect in our houses in about ten years' time.

The Internet is rapidly expanding. In the last four weeks, there was an intergovernmental conference on regulation of the Internet at which European governments discussed ways of controlling the traffic on the net. The reasons that were given for this conference were concerns about pornography (particularly pornography that could be accessed by children) and the use of the net by anarchist and other subversive/anti-social groups. Shortly after this conference, the *Sunday Times* ran an article on use of the Internet by anarchist groups. It claimed that groups in the UK and Holland were exchanging information on how to infiltrate and subvert demonstrations, how to make explosives, and on other such matters. This, governments will insist, is one very good reason to control the Internet.

It has to be remembered that the US government pushed the adoption of the DES. This is a data encryption standard; the algorithm which was pushed by government was not the best of the available options, but it had the pleasant property that the NSA was able, even at the time the standard was introducted, to decipher any traffic that was encrypted using the encryption algorithm. If governments are allowed to monitor the Internet (and other networks) for pornography and for information that they consider prejudicial to the security of the state, they will need to read all traffic. Picture formats, encrypted messages (which some academics use for sensitive transmissions, which businesses use for commercially confidential transmissions), sound encoding formats will all need to be standardised so that they can be read. Commercial information will pour into network monitoring stations: this could be to the detriment of the sending company. Posters of material of a dubious or critical nature will be open to investigation, and perhaps the techniques outlined in the last section might be used against them.

Masses of information will need to be searched if governments get their way and censor the Internet. In the limit, this could be enhanced by bans on the purchase of satellite dishes and decoders because they could be used to obtain private feeds from the Internet's comsats (one way of avoiding the national Internet nodes). The amount of information that would have to be sifted through would be enormous: agent-oriented techniques would be of benefit here and for reasons already given.

It should be clear that the kind of system that I have described above is not something that can be had today off the shelf, at least not in its entirety. There are scientific questions that need to be asked: these are questions of genuine scientific interest, so they should be asked and answers should be sought. There are, though, engineering issues, applied science and technology issues. The engineering and technology questions are those that convert science into working systems on the scale that would be required to control population movements, for example. One case in which work of a technological nature is required is that of the smart card; at present, smart cards are limited in storage and processing power. The solution to these problems is merely technological: there are no scientific problems. There are

always people who will work on such problems, even when they know that the result will not be for the general good (I can think of at least one academic who would work on such a project because it represents a quantity of money that will be recorded on a university document, thus making him look good to 'senior management').

I am not suggesting that the science should not be done, nor am I *necessarily* suggesting that the engineering and technology should not be done or provided. I *am* suggesting that there are ethical issues which must be confronted and the relationship between funding and political ends (of all kinds) be considered more critically.

Any organisation that musters its software agents in the ways described above must do so in secret. Perhaps secrecy does not matter too much: after all, even if you know that there are agents spying on you from the Internet what can you do? Agents can move physical location, so you can try very hard and still fail to find the process image that represents the agent. Agents can perform remote sensing, so you could try to switch devices off. This could be *Catch-22*: you might need the devices on to regulate your life: it could be *2001* in the sense that you have to reduce the functionality of your system to an unacceptable level in order to reduce the risk to an insignificant level. Openly admitting that information gained in this way might lead to commercial problems, but what if government and commerce are in the game together? The world's intelligence services are now re-orienting themselves towards economic intelligence, so such collusion might become common. It might be the case that government either relies upon commerce to such an extent (say for political funding or support in the sense of privatised function) that it cannot, perhaps dare not, regulate and police. Any kind of relationship between commercial interests and government is possible, as is the possibility that government engages in these activities on its own.

By its very nature, abuse of the kind proposes above will be very hard to detect and very hard to police, even if there is the political will.

Agent-oriented programs, as I have suggested, are harder to detect and harder to destroy. Their distributed state and their dynamic reconfigurability make for more robust programs that can reconstruct themselves in the event of failure. Their presence in real-time environments also implies that intervention with the aim of halting them must take the form of detective work. This is clearly a difficult task to perform and one that could easily fail. When an 'enemy' agent is detected (say, by means of its data emissions) what can be done? Could another agent in some way 'kill' it? One thing that could be done is to have agents intercept the 'enemy' agent's input data streams (but that would need knowledge of the agent's structure); alternatively, all resources in a system could be protected by agents that are partly concerned with enforcing security policies. In this case, we would have agents directly opposing agents: a kind of fight or subversion might ensue. If agents can destroy agents, though, there is a risk that entire systems will be destroyed in the sense that access to crucial resources will not be possible because their guardian

agents have been damaged or destroyed. Perhaps we have a case of Mutual Assured Destruction here that will prevent the construction of such intrusive systems: if the system is built, people will defend themselves and this might lead to the eventual destruction of the system, not control of the populace.

For those who find all of the above far-fetched, and who feel that the Internet and technological convergence is necessarily for the good of mankind, I will end with the following thoughts.

> *What is technically feasible today will become reality tomorrow if there is money or power to be gained from it. Money and politics are frequently mutually supportive, especially when politicians conspire in the abuse of power.*