Interactive to Proactive: Computer Ethics in the past and the future

By L Venter, MS Olivier and JJ Britz

Abstract

The integration of mobile technology, wireless networks, ubiquitous computing and artificial intelligence with thousands of embedded devices such as sensors and actuators may result in networks that can proactively monitor and respond to human behaviour without human interaction and with little supervision. Decisions that can influence or alter the environment will be made at faster-than-human speeds. Ethics as applied to current interactive computer systems will not be adequate. The paper proposes that the principles of clinical ethics can be applied to proactive computing, and argues that removal of the autonomy of the users of proactive systems may make such systems inherently unethical. We note that clinical ethics is accepted in the medical profession but that no such institutionalization and internalization exist in the computer industry. We conclude that although possible principles exist, significant change in the community is needed to apply them in practice.

1. Introduction

Computing technology is becoming more pervasive by the day. With the proliferation of small computing devices such as cellular phones, personal digital assistants (PDAs), smart cards and other devices, people are carrying on their person many devices capable of communicating with other devices over wireless networks. Further development of small processing chips allows embedding of computing devices into almost anything. Smart clothes, wearable computers and the like will make ubiquitous computing a reality, with a person interacting with hundreds or even thousands of devices embedded in the environment.

Interaction on a one-to-one basis with devices on this scale is humanly impossible. A single user cannot possibly configure and maintain every embedded device manually. The ideal situation would be that the devices configure themselves and decide for themselves what actions should be taken given a certain situation.

This leads to the concept of proactive computing [Tennenhouse, 2000, Want *et al*, 2003]. Proactive computing envisions networks of computing devices, sensors and actuators that dynamically configure and maintain themselves, monitor the environment and respond to or even adapt to the environment. They may even change the environment itself. These networks of devices will operate with very little human supervision. As such they will no longer be interactive and can therefore operate at faster-than-human speeds.

This paper looks at the way our perception of computer ethics must change to accommodate the next trend in computing devices. We consider ethical principles that can be applied to such systems, in particular the principle of autonomy, and ponder about the ethicality of a

system that denies its users autonomy. We note that the lack of internalised ethical principles in the computing industry may be the cause of unethical proactive systems.

Section 2 discusses the current view of computer ethics in an interactive computer world, and introduces the concept of proactive computing. Section 3 strives to identify ethical principles that can be applied to proactive computing, using examples and hypothetical cases to demonstrate their validity. Section 4 concludes the paper.

2. Background

Computers and computing devices as currently known are for the most part designed to be interactive with their users. The user will give a command; the computer will execute it and then wait for the next command. Most of the CPU time of personal computers is spent running the system "idle" process, even when the user is running several programs at once.

This makes a computer far less effective than it has the potential to be. Interactive computers are by nature limited to human speeds. They are also limited to act only on human commands. Expert systems may give suggestions after analysing trends in data, but the decisions still rest on human users.

The ethics of interactive computer systems focuses on how such systems are (or can be) used by human users. Most scholarly work focus on the following aspects [Johnson, 2001, Bynum and Rogerson, 2003, Erman and Shauf, 2002, Edgar, 1997]:

- Policies
- Moral and legal issues
- Professional ethics and responsibility
- Hackers and hacker ethics
- Netiquette
- Privacy
- Property rights
- Social and democratic issues
- Free expression
- Accountability and liability

All of these issues deal strictly with how humans can use or abuse computer technology in their social context. It is clear that current computer ethics accept the fact that computers only act on instructions from human users. But what about computers making decisions autonomously?

The concept of Artificial Intelligence has spawned numerous debates on the possibility of creating intelligent computers that can make autonomous decisions and "think" better than humans can. Opinions on this topic vary from the optimistic [Furse, 1999] to the paranoid [Joy, 2000] to the sceptical [Narin, 1993]. For the purpose of this paper, we will consider computers as simply following their programming, not as intelligent beings with free will and possibly malicious intent. Artificial intelligence is seen as just another tool used in the solving of computational problems, specifically that of analysing large amounts of data gathered by thousands of embedded sensors and other devices.

It is expected that in the not-so-distant future, computers and computing devices embedded in everyday objects will outnumber the people on this planet [Tennenhouse, 2000]. Where computers to date have been designed and programmed to be interactive with human users, it will become impossible for human users to supervise the working of thousands of devices per person. A growing field of research is therefore proactive computing; that is, the operation of networked computing devices without human interaction. Intel Corporation has set this as their long-term research goal, and has identified several key points to be addressed to practically implement this goal [Tennenhouse, 2000].

- Getting physical: That is, connecting pervasive networks directly to their environments.
- Getting real: Allowing devices to respond to stimuli at faster-than-human speeds.
- Getting out: Removing human interaction and having devices function autonomously.

Each of these points present technical challenges that ongoing research is addressing, but they also present issues regarding their use in the context of human society. Technology creates the opportunity for positive use but also for misuse of resources, which poses ethical questions that need to be addressed.

Proactive computing will need integration of several technologies currently available or in development. Technologies like energy harvesting sensor and actuator networks [Kansal and Srivastava, 2003] with wireless connectivity [Bose, 1999, Brunette, 2003] will be implemented using artificial intelligence [Kansal, 2004], distributed computing protocols [Poor, 2003] and security protocols [Undercoffer *et al.*, 2003] to create networks of devices that can monitor their environment, respond to changes and even change their environment as needed [Noury *et al.*, 2003]. Other key fields that will be related to proactive computing are mobile computing, pervasive or ubiquitous networks, context-aware networks, embedded networks and location-aware computing.

Brunette *et al*. [2003] describe several simple proactive applications making use of small wireless sensors. One possible application is a reminder system that logs the position and proximity of other sensors. By learning a user's normal behaviour, it might for example warn the user if he is leaving his office without his house keys, because that will deviate from the normal pattern. Other applications could be to monitor logistics in a factory or workroom, thereby enabling users to streamline the use of equipment and facilities.

Another possible application for proactive computing may be a traffic control system that monitors the flow of traffic using sensors embedded along the traffic grid, and transponders placed in vehicles. The system might increase the time for traffic light cycles to allow more congested roads to clear. Since the transponders in vehicles would respond to the system, it would also be easy to track specific vehicles, assisting with the recovery of stolen vehicles. Taking it even further, the vehicle might have a computer system itself, capable of communicating with the traffic system. The traffic system could therefore instruct a vehicle to slow down, taking control of the vehicle effectively out of the driver's hands.

The technologies to implement such systems already exist; the challenge currently is how to integrate the different technologies on a scale that makes it practical.

Proactive systems will make decisions without direct human supervision. If the traffic control system decides to hold up certain roads to make others flow faster, it would not consider the inconvenience to some road users. It would make the decision based on how the flow of cars can be optimized through the grid. If someone is trying to urgently get to a hospital, the system's decision may very well have a life-and-death impact on the user.

The ethical implications of decisions made by human users of interactive systems have been discussed at length since computers were invented. Now it becomes necessary to consider the implications of decisions made by computers. Principles are needed to judge the ethicality of such decisions. In the next part we propose some principles, and identify a dilemma.

3. Ethical principles for proactive computing?

People in positions of authority can make decisions that will influence other people. In a democratic government, larger groups of people elect representatives to make those decisions – decisions that are debated and discussed and made in concert with other representatives to benefit the country. Another position of authority is that of a doctor interacting with a patient. A doctor, using specialized knowledge and technology, is in the position to make decisions that influence a patient's health and life. The patient is dependent on the doctor to make those decisions because of the patient's lack of specialized knowledge. Even though the patient does have the right to refuse treatment, that decision can only be made if the doctor gave the patient the correct information.

A proactive computer system can be seen as being in a position of authority or power similar to that of a medical practitioner, since the decisions that it makes can influence the lives of the users of the system, even if only on a small scale. The users of a proactive system will be dependent on the system to make those decisions because of the sheer amount of information that the system is able to gather and process, but which the user won't be able to comprehend.

To identify ethical principles that may be applied to proactive computing, it is therefore helpful to look at the principles already identified for other situations where people are placed in positions of authority.

One example is the Intelligent Habitat for Health [Noury *et al.*, 2003]. This is described as an apartment fitted with "sensors and actuators which cooperate among themselves and with an information system" to monitor the physiological state of the patient, detect symptoms or events that indicate medical problems and to assist the patient by automatically controlling devices in the apartment. By constant monitoring of the patient, the system can identify crisis events and, rating the severity of the crises, decide who to contact. It may alert a neighbour if the patient got stuck in the bath, or alert emergency services if the patient's heart stops beating. This system is proposed as an alternative to keeping chronically ill or elderly people in hospital under constant human observation.

Since the system must analyse sensor data and on a minute to minute basis evaluate the patient's status and make a diagnosis based on that status, it plays much the same role as a doctor or a nurse observing a patient in a hospital. The principles of clinical ethics can therefore apply to this case.

3.1 The principles of clinical ethics

Beauchamp and Childress [1979] compiled an outline of ethical principles for use in the medical profession. They are personal autonomy, veracity, nonmaleficence, beneficence, confidentiality and justice.

Personal autonomy implies the right to act on one's own behalf. This requires information to make informed decisions as well as the freedom to make those decisions. Patients therefore should have the right to know all relevant information about their cases, to refuse treatment or to participate in a procedure and most of all, the right to have their autonomous decision implemented as they wished.

Veracity is closely tied to personal autonomy as it is the right to be told the truth about a situation. This is necessary to make informed decisions. Nonmaleficence requires the doctor not to inflict harm and to strive to prevent harm where possible. Beneficence implies that whatever is done to the patient should be to his benefit. Confidentiality requires that the doctor should not make public the details of a patient's case without consent from the patient. Finally justice requires that "like cases should be treated alike" [Francoeur, 1983]. This principle calls for the fair distribution of health care and services.

In the context of proactive computing, we propose that these principles can be applied as follows: autonomy refers to the right of the user to decide what should happen as result of a given situation. Veracity implies that the user knows exactly what information the system gathers about the user. Beneficence is the expectation that the use of the system will be for doing good; nonmaleficence is the expectation that the system will not be used with bad intent. Confidentiality ensures that the information gathered by the system will not be freely available. Justice is the expectation that the system's decisions will be fair.

Applying these principles

In considering proactive systems one should not fall into the trap of ascribing thought, intentions and emotions to the system. A proactive system will have no more consciousness than a personal computer does now. Any intent in the use thereof should therefore be considered the domain of the designer and/or the owner of the system. This is more in line with conventional computer ethics which could be deemed adequate to cover the issues involved. However, since people will place proactive systems in positions of "authority" over other human beings, it may still be useful to see how each of the bioethics principles can be applied to a proactive computing context.

A proactive system will be designed and developed for a specific application, just as interactive systems are designed now. The application or purpose of the system could therefore be measured against the bioethics principles to gain an understanding of the ethicality of the system.

A traffic control system, developed to ensure smooth traffic flow and to reduce accidents, would have a beneficent purpose. Any decisions that the system makes, for example to hold up certain streets to clear others, would be non-maleficent and would be made based on a straightforward cost-benefit analysis: for the cost of delaying a few data points in the grid,

many others would have optimised flow. And in any case, as soon as a held up street becomes too congested, the system will reassess and give preference to those data points.

If the system was developed with ill intent, if the designer for some reason deliberately programmed it to cause accidents rather than prevent them, or if the owner of the system (the local government) deliberately misuses the system to spy on people, then the purpose may not be beneficent and non-maleficent. But in that case the intent lies with the designer or owner, and not the system itself, and the situation would be no different from people misusing computers and technology now.

Proactive systems will gather huge amounts of data with their sensor networks. This data will be used to make the decisions that implement the purpose of the system. Just as data can be misused now, data gathered by proactive systems could be intercepted and used for other purposes. For example, the police can use the traffic control system to trace stolen vehicles or suspects. In the same manner, a terrorist group could tap into the system and use the data to trace vehicles of political leaders, or a group of thieves could trace a money transport vehicle to plan a heist. Confidentiality of the data gathered by the system is therefore just as important in proactive systems than in any computer system. Guaranteeing confidentiality is the domain of computer and information security [Pfleeger, 2003] and current technologies like intrusion detection, access control and encryption would have to be integrated on a large scale.

The just and fair decisions of a proactive system would once again depend on its purpose and the way it was designed. Coming back to the traffic control system, it would make sense to implement the system in such a way that certain vehicles, such as ambulances or police cars, would have a higher priority in the grid, and the system would ensure their easy movement through the grid. There could therefore be levels of importance assigned to users, but in a fair system, all users on the same level would be treated alike. Possible misuse of such an implementation could be that users could somehow assign themselves a higher level, thereby getting an unfair advantage in the system.

Veracity, the right of the user to know exactly what information is being gathered by the system and what it is used for, is closely tied to confidentiality. As a user of the traffic control system, one should be told exactly what the system is doing with the data. Is it merely a transponder signal that signals a data point in the grid? Or does the transponder uniquely identify each vehicle? Is the data used in real time and then discarded, or is it kept to build up a history of each user's movements? The owner of system that is designed for a benevolent purpose and is not being misused should have no problem releasing details on what is done with the data, without making the data itself available.

Proactive systems can be designed with one or many users in mind. The Habitat for Health described in a previous section would have one user at a time, and can therefore be personalized for that user. User preferences and habits can be "learned" to make the system as user-friendly as possible. This doesn't give the user control over the system. If the system is designed to execute certain actions, such as to keep the temperature in the room at a certain level, then the user will be not able to change it. The user won't be able to tell the system to stop monitoring location or life signs, even if the user might want some "privacy". User autonomy is therefore restricted in a small-scale system.

Another example suggested by IBM Research [Want *et al*, 2003] is a system that proactively monitors and adjusts a home central heating system. The system could optimise energy usage

to save on fuel or electricity. With fuel shortages likely to increase in the future [Duncan, 1996, 2000], it is conceivable that governments will require their citizens to use such systems to conserve energy. They could enforce quotas, causing the system to switch off lights and heating at prescribed times. The system could even report energy waste to the authorities. The home owner will have little autonomy with regards to the running of the system, which could act as a virtual energy dictator in the home.

In a larger system such as the traffic control system, there will be little personalization possible, as the system couldn't practically adapt to suit all the drivers in the city all the time. There will be no way for a driver to control the system. Even if one could remove the transponder from a vehicle, to become "invisible" to the system, then one would still be influenced by the system's control over the traffic lights. Short of getting out and walking, there would be no way of avoiding the system's control. User autonomy is therefore nonexistent in the context of a large-scale system.

Just as every doctor has to make ethical decisions on a case-by-case basis, and does so by measuring them against the bioethical principles, so one can look at each proactive system and measure its decisions against the principles. In doing so it becomes obvious that a proactive system can be benevolent, non-maleficent, truthful, confidential and just – but by its very nature it denies its users full autonomy.

Interestingly, it is possible to transfer the ethical responsibility of five of the principles (beneficence, nonmaleficence, veracity, confidentiality and justice) to the designer/creator of a proactive system, but user autonomy can not be transferred.

3.2 Proactive computing is inherently unethical?

Personal autonomy is seen as being fundamental to our value as human beings [Johnson, 2001]. Without the ability and freedom to make our own choices, we cannot function as independent human beings. Living in a society with established laws, rules and expectations, we do place ourselves under certain authorities that limit our autonomy to ensure that we comply with the rules of the society. But we expect the right and freedom to make decisions (and act on them) that influence our own lives in a way that satisfies us.

Proactive systems will place another level of authority over people, making decisions that can impact their environment and their lives. Using principles of ethics such as that of bioethics, one can judge whether a proactive system is making ethical decisions up to a point, but since the inherent nature of proactive systems is to remove humans from the control loop as much as possible [Want *et al.*, 2003], the principle of personal autonomy will be hard to prove to be present.

A case could therefore be made that proactive computing will be inherently unethical. After all, if a doctor makes a decision about a patient without taking the patient's wishes into consideration, that would be an unethical decision. Proactive systems will do that as a matter of course. But as people are under the authority of governments, councils, company management and other authorities at all times, and their personal autonomy is being restricted by these authorities in any case, one could conclude that authority is inherently unethical.

Since authority is necessary for an orderly society, we accept those restrictions on our autonomy, and do not consider it as unethical (unless in a government misusing its power to suppress the people). We therefore choose to live under such restrictions for the benefit of imposed order to society.

Similarly, if a proactive system, by imposing some restrictions, ultimately benefits its users, then the users may choose to accept the restrictions for the service that the system provides. If the restrictions imposed by a proactive system can be seen as contributing to the order of the society, or to the user's benefit, without harming the user physically or emotionally, then one can say they are not more unethical than any other imposed laws. The main difference is that the proactive system can enforce its "laws" directly, while the government has to rely on the justice system to do so.

Proactive systems are therefore not inherently unethical, but to determine the ethicality of a system would require review of every system on a case to case basis in light of accepted ethical principles. The next section discusses the role of professionalism in this process.

3.3 Institutionalised and internalised principles

The principles of medical or bioethics have a long history [Kuhse and Singer, 2001] and have been widely accepted in the medical profession. As professionals, doctors and health specialists are required to register as medical practitioners, and they are subject to discipline if they should act in a manner that contravenes the principles as accepted by the profession. The principles have effectively been internalised and institutionalised in the medical profession, so that the Charter of Medical Professionalism states "The profession as a whole must strive to see that all of its members are competent and must ensure that the appropriate mechanisms are available for physicians to accomplish this goal" and also "Physicians are expected to … participate in the process of self-regulation, including remediation and discipline of members who have failed to meet professional standards" [Sox, 2002].

In contrast, computing is not yet seen as a profession [Johnson, 2001], although there are many professional groups in the computing field. Although each of these groups have codes of ethics and standards to which their members should keep, membership to these groups are not compulsory and failure to meet the standards have no real disciplinary result. As such, programmers and other IT specialists can work in the computing field with no regulatory structure to ensure that the work is done in an ethical and appropriate manner.

Proactive systems will be making decisions that will influence people's lives, and they will be designed by people who may not consider the ethical issues that the systems may cause. As stated in a previous section, the ethical principles impacting such systems can for the most part be transferred to their designers, whose intent and competence will determine in which way the systems will influence their user's lives. If a programmer does not care about the ethical implications of a system, then the system may very well be unethical in the end.

It will therefore be necessary for IT professionals to reassess computing as a profession, because even if it is not necessary right now, it will become paramount in the future to ensure the safety and wellbeing of people dependent on the systems we build.

4. Conclusion

In conclusion, the development of proactive systems will:

- Provide users with extensive services with minimal human supervision
- Restrict user autonomy by taking humans out of the control loop
- Force a reassessment of professional structures in the computing field to ensure compliance to appropriate standards

Future work will be to look at government and public responsibility and how proactive systems might impact the open society due to the fact that their functionality may be hidden from the users.

5. References

Bose, V. and Wetherall, D. and Guttag, J. (1999), Next century challenges: Radioactive networks, Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking, 242-248

Beauchamp, T.L. and Childress, J.F. (1979), Principles of Biomedical Ethics, New York: Oxford University Press

Brunette, W., et al. (2003), Systems, platforms and applications: Proximity interactions between wireless sensors and their application, Proceedings of the 2nd ACM international conference on Wireless sensor networks and applications, 30-37

Bynum, T.W. and Rogerson, S. (ed) (2003), Computer ethics and professional responsibility: Introductory text and readings, Blackwell Publishers

Cram, N. and Wheeler, J.(1995), Ethical issues of life-sustaining technology, IEEE Technology and Society Magazine, Spring 1995

Duncan, R.C. (1996), The Olduvai Theory: Sliding towards a Post-Industrial Stone Age, online at http://www.dieoff.org/page125.htm accessed 13.06.2005

Duncan, R.C. (2000), The peak of world oil production and the road to the Olduvai gorge, online at http://www.dieoff.org/page224.htm accessed 13.06.2005

Edgar, S.L. (1997), Perspectives on Computer Ethics, Jones & Bartlett Publishers

Erman, M.D. and Shauf, M.S. (ed) (2002), Computers, ethics and society, 3rd Edition, Oxford University Press

Francoeur, R.T. (1983), Biomedical Ethics – A guide to decision making, John Wiley & Sons

Furse, E. (1999), Arguments for strong AI, online at http://www.comp.glam.ac.uk/pages/staff/efurse/Theology-of-Robots/Arguments-for-Strong-AI.html, accessed 13.06.2005

Johnson, D.G. (2001), Computer Ethics, 3rd edition, Prentice-Hall

Joy, B. (2000), Why the future doesn't need us, Wired 8.04, online at http://www.wired.com/wired/archive/8.04/joy_pr.html, accessed 13.06.2005

Kalapriya, K., et al (2004), A framework for resource discovery in pervasive computing for mobile aware task execution, Proceedings of the first conference on Computing frontiers, 70-77

Kansal, A. and Somasundara, A.A. and Jea, D.D. and Srivastava, M.B. (2004), Intelligent fluid infrastructure for embedded networks, Proceedings of the 2nd international conference on Mobile systems, applications, and services, 111-124

Kansal, A. and Srivastava, M.B. (2003), An environmental energy harvesting framework for sensor networks, Proceedings of the 2003 international symposium on Low power electronics and design, 481-486

Kuhse, H. and Singer, P. (ed) (2001), A companion to bioethics, Blackwell Publishers

Maybury, M.T. (1990), The mind matters: Artificial intelligence and its societal implications, IEEE Technology and Society Magazine, June/July 1990

Narin, A. (1993), The myths of Artificial Intelligence, online at http://www.narin.com/attila/ai.htm, accessed 13.06.2005

Noury, N. and Virone, G. and Barralon, P and Ye, J and Rialle, V. and Demongeot, J. (2003), New trends in health smart homes, ITBM-RBM (RBM), 24:3, 122-135

Pfleeger, C.P. and Lawrence Pfleeger, S. (2003), Security in computing, 3rd Edition, Prentice Hall

Poor, R. and Auburn, C.B. and Bowman, C. (2003), Self-healing networks, QUEUE, May 2003, 52-59

Sox, H.C. (ed) (2002), Medical Professionalism in the New Millennium: A Physician Charter, Annals of Internal Medicine, 136:3, 243-246

Tennenhouse, D. (2000), Proactive Computing. Communications of the ACM, 43(5):43-50

Undercoffer, J. and Perich, P. and Cedelnik, A. and Kagal, L. and Joshi, A (2003), A secure infrastructure for service discovery and access in pervasive computing, Mobile Networks and Applications, 8:2, 113-125