The vulnerability of security – a dialogue with Ross Anderson

Man has been acquiring possessions and calling them 'his', has sought to safeguard them from 'others'. He has protected his home with locks and alarms, his wealth with safes and banks, his information with passwords, his inventions with patents, his country with borders and soldiers . . . it is a long list. And technology has made protection more complex. But has it become better?

Ross Anderson's life-work involves establishing security engineering as a discipline. In his book *Security Engineering: A Guide to Building Dependable Distributed Systems*, he ventures a definition: 'Security engineering is about building systems to remain dependable in the face of malice, error or mischance. As a discipline, it focuses on the tools, processes and methods needed to design, implement and test complete systems, and to adapt existing systems as their environment evolves'. In his Indian Institute of Science (IISc) centenary lecture 'What is engineering?' delivered at IISc, Bangalore in December 2009, he spoke of the increasing complexity of socio-technical systems, the accompanying problems of crime, conflict and system failures, and strategies to improve system dependability. He also stressed on the need for learning from and applying tested methods from other disciplines such as economics, systems biology and psychology, and indicates '. . . it is often at boundaries between disciplines that big advances can be made'.

Ross Anderson is Professor of Security Engineering at the Cambridge University Computer Laboratory, where he teaches courses in software engineering, economics and law, and security. He has an MA in mathematics and a Ph D in computer engineering, he worked in avionics, in banking and as a consultant. He was one of the founders of the discipline of security economics and a pioneer of peer-to-peer systems, hardware tamper-resistance and information hiding. He designed the 128-bit block cipher 'Serpent' with Eli Biham and Lars Knudsen, which was a finalist in the competition to find an advanced encryption standard. He is a Fellow of the Royal Society, the Royal Academy of Engineering and the Institute of Physics. He is also the chair for the Foundation for Information Policy Research, UK and an elected member of his university's governing body, the Council, for the period 2003–2010.

The following is an account of the conversation with Ross Anderson on 'vulnerable security' held on 15 December 2009 at the Supercomputer Education and Research Centre (SERC), IISc.

People say security is a 'pain' (privacy, authentication, integrity, non-repudiation). Is it so in the colloquial sense also?

It is better to see security as being, very often, the expression of a power relationship. In the olden days, people used to describe security as simply keeping the bad guys out. But now, it tends to be about regulating behaviour using principles in a system. The big socio-technical systems that are growing today have millions of users who are competing with, trying to control and exploit each other, and indulging in strategic behaviour. So it is no longer possible to say that security is a single indivisible thing. Security is about conflict and power and control. It means different things in different contexts.

In your book, you had quoted Rick Maybury: 'Microsoft could have incorporated effective security measures as standard, but good sense prevailed. Security systems have a nasty habit of backfiring, and there is no doubt they would cause enormous problems'. Can you elaborate on this? Does Murphy's law apply here?

Platform products such as Windows can bring enormous profits to whoever controls them. In order to control such a market, you generally have to be the first. And in a two-sided market where you are appealing not just to the end users but also to the people who develop software for them – then in the beginning it is very important to appeal to the developers. If Microsoft had made Windows more secure in the early stages – Windows 3, Windows 95, Windows 98 – then it would have been more difficult for people to write software for Windows. They might have written software for OS 2 instead, or for Mac, or for Unix. Windows would not have come to enjoy its huge market share. So economics explains why Windows did not have proper security for many years.

Once you have a powerful position in a platform market, it is then rational to add more security to lock your customers in. The security mechanisms of OS 360 were very primitive, then lots more stuff got added as it grew into VM and MVS. Mobile phone operating systems like Symbian system UIQ2 had no security at all really; Symbian UIQ3 brought in many protection mechanisms. Security economics explains this pattern of behaviour – where systems are first insecure and then secure.

Why is it so difficult to design robust systems? Does more complexity mean less reliability?

The dependability of systems is ultimately a result of economic factors and there are many interesting market failures in the way systems evolve. For example, a typical computer or mobile phone has too many features, it is too complex. As a result, it is difficult to use and unreliable; you get software failures.

When the developer is thinking of adding a new feature, the benefit of that feature goes to a small concentrated local group of users; whereas the harm of the extra feature, the added complexity, the added strain on the device's ecology, is spread over all users. So this is just the tragedy of the commons. At the system level, this leads to unreliability; the interfaces between the operating system and applications also become too complex. Operating system vendors make them more complex to stop competitors building compatible systems.

In systems engineering, we are fighting a constant battle against complexity.

We are forever trying to climb up the complexity mountain, until eventually we slip and fall off. Much of what we do as computer scientists is about designing better tools that let us get a little bit further up the mountain. If you push up too far, then your system is too complex, it does not work and you can not ship it. But if you do not push up far enough, then you do not offer as many features as your competitors.

How do we compare the protection incorporated in various operating systems such as Linux and Windows?

Both Windows and Linux are so complex that vulnerabilities are being discovered all the time. However, the bad guys write malware principally for Windows because there are more Windows systems. If each infected machine is going to earn you $1 a week, then obviously you would prefer to infect 100,000 Windows machines rather than 5000 Macs.

But many people are shifting from Windows to Linux these days. . .

If we start seeing Linux in many embedded systems, then it might become the target of choice. For example, many vendors of consumer electronics devices are starting to put Linux in devices such as televisions. If you do not plan in advance to upgrade and patch the software in all the television sets, then you have a problem.

So we need to start working right now on these lines?

What is needed is regulation saying that companies which sell devices that are attached to the network should certify that they are safe by default. Such companies should not be able to disclaim liability by saying that 'if something goes wrong you cannot sue us'.

What should be the mandatory subjects in a course on security engineering? How do we effectively integrate lessons learnt from other disciplines like systems biology, economics and psychology?

There are three traditional core subjects. The first is operating systems security – how things like access controls and capabilities can be used to express security policies such as multi-level security. The second is cryptography – the analysis and design of codes and ciphers. The third is protocols – the rules that bind the cryptography and the operating systems by enabling machines to communicate with each other.

In addition to that, I think it is necessary for people to study many real case histories in security engineering, which is why I wrote my book, with lots of information about how specific military systems, banking systems, monitoring systems and so on work and how they fail. It is also necessary to understand incentives – how these systems work within a business context, and how they often fail because the incentives are wrong. That means we need some understanding of economics. In Cambridge, for example, I have introduced a short service course in economics for all computer science students, not just to support the courses that we do in security but also the courses that we do in e-commerce and in business.

We will probably be adding in some material from the cognitive sciences. For many years there has been a fruitful interchange between the cognitive sciences and computer science. Within computer science, the study of robotics, natural language processing, computer vision and human–computer interaction have drawn heavily on psychological work.

What can we learn from nature?

There are people who do work on learning lessons from conflicts in nature: from conflict within genes, conflict between microorganisms and the larger organisms on which they are parasitical, and the many attack and defence techniques that predators and prey use. One question that is starting to interest me is whether we can learn interesting lessons from systems biology. There are some aspects of the cell's architecture which are very strongly conserved over long periods of time – clusters of genes and genetic mechanisms which have been extremely resistant to any change, with mechanisms of information protection and preservation.

What are the security problems due to man–machine gap?

The first area of research into security and psychology was the field of usability. Alma Whitten, who is now the security manager of *Google*, did a usability survey PGP which then was the most common encryption program, and found that even postgraduate level people could not use it in a safe way. So, the prospect that a normal man would be able to operate this software safely was close to zero. We have done some work on the best ways of advising people to choose passwords, and more recent work indicates that the mechanisms used in online banking discriminate against women because of differences in the way that men and women approach systems. This established that the 'gender HCI issue', as it is called, is also salient in security.

What is electronic and information warfare?

Electronic warfare has been around, I suppose, since the nineteenth century, since the invention of the telegram. It was first used seriously in World War I and very heavily in World War II with radar jamming accompanying the air war. It is common nowadays for air battles to be largely decided by electronic duels – who can jam the other side's radar and still maintain control of their own missiles and direct them at the target? It is of fundamental importance and huge amounts of money get invested in electronic warfare by the major powers.

Some say that information warfare might become big, as people learn to attack the other country's computers and telecommunication systems. This is of sufficient concern for some countries that, as we speak, there are apparently talks going on between America and Russia on a cyber disarmament treaty. The Russians want some assurance that the Americans won't attack them using computer means.

The extent to which this is a problem for India is not clear. Certainly the Chinese have been using hacking techniques to hack PCs in a number of Government agencies in incidents that have become public. And to an extent this is what intelligence agencies have always been doing to each other.

How do firewalls help in internet security?

Some years ago, people started using internet communications or at least IP networks because they were a lot cheaper than the previously used ad hoc mechanisms. It was realized about 10 years ago that this meant that a knowledgeable

attacker could go into, for example, an electricity distribution system and cause chaos. Since then, re-perimeterization has become big: you put a lot of effort into firewalls which 'wall off' the network in your electricity sub-station or your oil refinery or your chemical plant from the outside world, so that people on the outside simply cannot talk to things like sensors and actuators involved in critical process control work. So, firewalls enable you to adapt to rapidly changing network technology when the systems we are trying to protect cannot adapt quickly.

India is now working on a unique identification number (the Unique Identification Authority of India [UIDAI] project). Will this system be secure? What could be the possible vulnerabilities?

In the UK, we had a project to introduce a national identity card. Various reasons were given for the identity card such as cutting back on welfare fraud. But in Britain most welfare frauds are not people lying about their names, it is people lying about their circumstances. Then they said that the ID card would help the police fight crime and the intelligence services fight terrorism. But the policemen and intelligence services know who the bad guys are; they just do not have the evidence and do not know their intentions all the time. So in each case, the identity card does not exactly do what it is advertised to do.

And, a huge part of the project cost had to do with setting up dozens of centres throughout the UK, where everybody would have to go and be finger-printed whenever they renewed their driving license or passport or got their identity card. So, it would have been a huge source of jobs and patronage in the public sector. Colleagues at the London School of Economics estimated that the total cost could have been £ 20 billion. Now, public opinion in Britain has turned against this, and there is a realization that it was just a political ploy whose time is now past.

In America, after 9/11 they started taking two fingerprints from every foreigner entering the country. The idea was that if somebody got into America using a passport with one name and came in again using a passport with another name, they would be able to identify him. But this did not work because the equal error rate with fingerprints is about 1% per finger. So if you take two fingerprints, that means the error rate is 1 in 10,000. And if you have tens of millions of people coming into your country every year, then the rate of false alarms is so high that it swamps the system. So the database that the Americans built for this fingerprint system is basically unserviceable. They are trying to move to four prints or to ten prints. But it is a matter of great confusion, and there was so much political capital invested in it in the Bush era that they have to keep on going forward.

You can use biometrics to disambiguate people, to tell whether somebody who embarked here in Karnataka is the same person who disembarked last week in Kerala. But the appropriate technology to do that is iris scanning, because it has a very much lower error rate than fingerprints. The error rate is better than one in a billion. So, if I were going to build a system for identifying persons uniquely, I would probably limit it to welfare and use more effective technology. Rolling out an identity card to the whole population would probably be technological over-kill. It is better to focus the mechanism on the people who would need it, the actual beneficiaries.

So, we need to focus on special sections of the population, instead of giving everyone a number?

Well, everything that I know about the problems of building such big systems teaches me that. You build a system that works for one place, then you roll it out to another place, and at each stage you pull out the bugs and you examine the costs…you ask yourself whether you are building the right system, whether it actually is cheaper than the paper system you are replacing. You always have to bear in mind that once you put a system into production, it does a lot less well than your initial pilots. With the initial pilots, you are using very highly motivated people. Whereas once you roll it out on a big scale, the operations are being done by general public service employees who are not as well educated or highly motivated.

What are the recent advances in security technology? What is expected in future? Are socio-technical systems going to become more complex, with more problems and conflicts?

I would say that we have had much of the technology now for several years. We have now got decent encryption algorithms, a reasonable repertoire of security protocols, and operating systems with reasonable access control. We also have virtualization – systems such as VMware and Xen. What we are more short of is knowledge and expertise in the methodologies that we need to design, maintain and evolve very large complex systems. Many systems fail basically because the incentives are wrong, because the management is wrong, or because we get the human interface wrong. From my point of view, I think the most interesting problem is that we are going to get more and more tense as systems become bigger and more complex, considering forces that are driving these systems' evolution on the one hand and our ability to maintain that evolution in a secure way on the other hand. The big problems are going to be in information governance and the extent to which interfaces evolve; how systems scale to a global scale. And as the complexity also scales and evolutionary pressures become stronger, how do we manage to keep systems evolving in such a way that they are sufficiently predictable and able to perform critical social, technical and economic functions? That is the challenge that we face today.

ACKNOWLEDGEMENTS. I thank Mr R. Krishna Murthy (Principal Research Scientist, SERC), and Prof. R. Govindarajan (Chairman, SERC) for their guidance in conducting this interview.

Geethanjali Monto (*S. Ramaseshan Fellow*), D-215, D-type Apartments, Indian Institute of Science, Bangalore 560 012, India. e-mail: geethum@hotmail.com