FUTURE THREATS

John Aycock
Department of Computer Science, University of Calgary, Calgary, Alberta, Canada

Email aycock@cpsc.ucalgary.ca

Alana Maurushat
Faculty of Law, University of New South Wales, Sydney, NSW, Australia

Email a.maurushat@unsw.edu.au

ABSTRACT

'Aren't you just giving the virus writers ideas?'

Research into future security threats is not always taken kindly by all members of the anti-virus community. Understanding the bounds of security research is important, however, with the increasing emphasis on proactive rather than reactive defences.

We begin to examine how far security research can go in terms of looking at future threats, particularly the publication of future threats. Our analysis addresses cultural, ethical and legal issues. We hope to provoke a meaningful discussion on future threat research within the anti-virus community.

INTRODUCTION

Research into computer security threats that have not yet been seen in the wild – potential future threats – is a fixture in the computer security field, for better or worse. But is it better, or is it worse? The whole idea of future threat research does not seem to have been examined thoroughly, even though opinions about it clearly exist; presentations of such work are not always received well.

The major concern is that an adversary, who could be a malware writer, spammer, or someone who manually cracks into systems, could be given ideas and techniques to use that they might not have arrived at on their own. This poses a dilemma to a researcher who has identified a future threat. Should the researcher publish it or not? Note that we use the term 'publish' here in the general sense of making the future threat publicly known, be it at a conference, in a blog entry, a media interview, or a public security mailing list (we will refine the notion of publishing later).

Figure 1 shows the combinations we consider, when both the researcher's and the adversary's actions are taken into account. Two combinations are impossible, because the adversary cannot make use of a published threat if it's not known or published. About half of the combinations result in no effect, as the adversary never uses the threat or the researcher has predicted the threat incorrectly. There is also the case where a researcher has published a threat, but the adversary arrives at the threat and uses it independently of the researcher's publication. In this case, the publication is moot, but by preceding the use of the threat, it has given the opportunity for a threat response to be considered and defences put into place.

Two possible ethical problems are identified. One is where a researcher knows of a threat, but does not make that

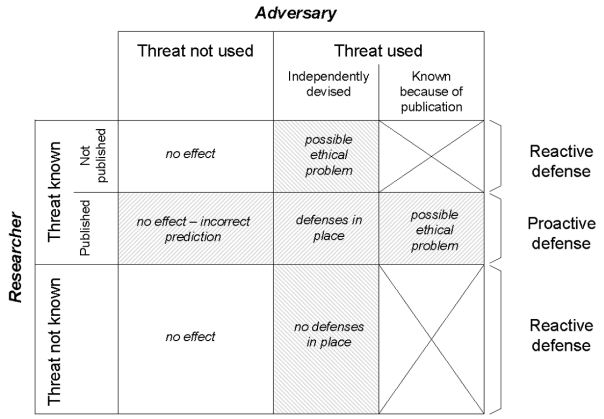

Figure 1: Researcher's and adversary's actions.

knowledge public. An adversary later discovers the threat on their own and uses it. The second is the above-mentioned concern, where the researcher's publication leads to an adversary using the information. Both of these will be examined in the Ethics section. Note that while a situation may not present any ethical or defensive problems, in particular where the threat is not used, it may still create legal issues. Such issues will be explored in the Legal section.

Figure 1 also maps out the difference between proactive and reactive defences (shown by the two types of grey background patterns). We define proactive defences to be ones where a researcher has published a future threat, allowing the possibility of defences being adopted before the adversary uses the threat. Reactive defences are ones where the adversary devises a new threat, and defences are forced to respond after the fact.

Some situations we will deliberately ignore. One is where a researcher knows of a threat, and defends their own systems against it, but does not publish the threat. While the researcher's systems may be protected, other systems are not; in other words, we are only concerned with the defence of the majority of systems. We also do not consider terrorist, military, and other information warfare scenarios, where researching future threats would be an active endeavour for strategic reasons.

In the remainder of this paper, we categorize different types of future threat research and the actions that can be taken by a researcher. We then look at cultural, legal and ethical aspects of future threat work.

TYPES OF RESEARCH

When talking about research into future threats, we distinguish between two types of research: practical and theoretical.

Practical research is research that involves the implementation and testing of a future threat. For example, a researcher might implement some new virus in order to analyse its properties. Interestingly, this is explicitly permitted by the AVIEN code of conduct [1]:

> 'I will not write replicative or destructive code unless I am convinced that it is necessary for internal research or testing purposes as required and defined by my professional activities. If I regard it as necessary to write such code, I will do so under secure and strictly controlled conditions, and I will not publish such code.'

In contrast, we define theoretical future threat research as research without a practical component: no implementation. Here, the crucial element of the work is the idea. In some cases, it may not even be possible to conduct practical research. For instance, a future threat involving a massive botnet cannot be tested in its full glory, simply because a researcher does not have a massive botnet at their disposal.

We make two assumptions in this paper regarding practical research to keep the analysis manageable. First, we assume that implemented threats are handled securely, and that no escapes into the wild occur. Second, we will assume that any code for a future threat is not published or released by the researcher – which is not an uncommon occurrence in computer science research generally. Under these assumptions, the underlying component in both practical and theoretical future threat research is ideas. Are there circumstances where ideas are too dangerous to express publicly?

POSSIBLE ACTIONS

A researcher who has identified a future threat may have a variety of potential actions from which to choose. In practice, the researcher may work for an organization that constrains the actions, but we look at the full range of options here:

- No action. The researcher does not tell anyone about the threat.
- Limited disclosure. Information about the threat is given to a select few, presumably those best situated to act on the information in terms of establishing a defence.
- Publishing. The future threat is made publicly known by the researcher.
- Publishing with warning. The researcher publishes the future threat, but also (possibly prior to publishing) engages in limited disclosure.

This is a high-level classification; further distinctions, like the publishing venue, are made later in the paper.

COMMUNITY NORMS

Examining and publishing new or future threats is common in some communities. So-called 'hacker' conferences like Black Hat, DEFCON and CanSecWest regularly do so; new threats feature in hacker publications too, like *2600* and the currently-maybe-defunct *Phrack*. It is easy to dismiss the hacker community as being a lawless, fringe element, but it seems they are in many ways closer to the overall stance of the broader computer security community than the anti-virus community is. We illustrate this by comparing experiences in the anti-virus community with two other security communities.

Cryptography

The cryptography community has had strong governmental influences on it in the past to keep certain work from being published. Dennis Ritchie, for example, tells the story of a 1978 paper he wrote with some colleagues where the NSA intervened, encouraging them not to publish it in a journal [2]. Whitfield Diffie expresses similar sentiments in his foreword to Schneier's *Applied Cryptography* book [3], and Levy expands on this at length [4].

Now, however, the situation has changed. The test of a new algorithm is whether it has been subjected to public scrutiny and survived. Cryptographers actively look for new attacks on existing, deployed cryptographic algorithms, even if there is no known way to fix the problems yet. A famous recent example of this was at the 'rump' session of the CRYPTO 2004 conference. There, Xiaoyun Wang presented attacks against well-known cryptographic hash functions including MD5, and was given a standing ovation. Finding new attacks is treated as a challenging intellectual pursuit.

Applied, non-AV security

There is no official name for this community, but we are referring to applied security conferences that are not AV-centric, typically with more academic involvement. Like the USENIX Security Symposium, where in 2002 Staniford *et al.* presented an influential paper – *Google Scholar* currently shows over 500 citations – that explains how to construct what they called '"Better" worms' [5]. Perusing the last four years of the ACM WORM archives shows one paper a year discussing new types of worm or improvements to old techniques.

Recently, USENIX published a call for papers for a new workshop called WOOT, the Workshop on Offensive Technologies, which says in part [6]:

> 'Submissions should reflect the state of the art in offensive computer security technology – either surveying previously poorly known areas or presenting entirely new attacks.'

This appears to be a clear attempt to draw upon future threat knowledge from the hacker community.

Anti-virus

While future threat work does appear at anti-virus conferences like VB and EICAR, there are no standing ovations. All of the following cases are based on first-hand observations or have been corroborated; the names have been changed and details left vague to avoid identifying specific people.

- Alice presents a heavily theoretical paper at an anti-virus conference. As part of her work, she implemented a proof-of-concept in her secure lab facility where the implementation work involved creating new malware variants. Upon discovering this, several audience members harshly criticize her on this point, independent of the work itself or the implications the work has in terms of dealing with future threats.
- Bob presents a practical paper at an anti-virus conference, exploring new attacks and variations on old attacks for an up-and-coming platform; this is potentially of great interest to anti-virus companies looking to expand onto that platform. He is taken to task by some audience members for giving the 'bad guys' ideas.
- Carol writes a paper about an existing computer security threat in the field of 'Computology', which is not normally connected to the anti-virus community. She expands on the idea, posing it as a future threat, and uses it to motivate other future threats, these ones directly in computer security. She sends the paper to an applied, non-AV security conference, where it gets good reviews but is ultimately rejected. She then sends it to an anti-virus conference, where the reviews lambaste the

paper because it may encourage Computologists to use the threat.

- Dave discovers that some non-security research work is applicable in malicious ways. He writes a paper about it for an anti-virus conference and the first question Dave is asked after his presentation is 'why are you publishing this? Aren't you just giving the virus writers ideas?'

Obviously there are vocal elements within the anti-virus community who object to research into future threats.

Discussion

In comparing the different communities all under the 'security' banner in this section, we are not implying that any of them are better or worse, nor are we suggesting that a researcher's ethical behaviour should be governed by the community they are in. This would be ethical relativism [7], an idea that can be used to justify deplorable acts such as slavery and the Holocaust.

Instead, we draw attention to the fact that, within the broader security context, the reception of the anti-virus community to future threat research is anomalous. In some respects this reaction is understandable; after all, anti-virus researchers are faced with threats on a daily basis. On one hand, there may be a belief that open disclosure diminishes the ability of an anti-virus company to remain competitive. On the other hand, this could also indicate that the anti-virus community is in the same situation as cryptography was decades ago, and that a greater openness to future threats is a quality that evolves over time.

LEGAL ISSUES

Future threat research involves a number of complex legal issues. We will explore general legal issues which could affect most of the combinations in Figure 2. We will then discuss the salient issues involved with the publication of threat research, as well as touch briefly upon other more tangential legal issues. It must be stressed that the law has had a very limited opportunity to address issues arising from the result of computer security. Most legal provisions have not been sufficiently tested in court to be able to provide definitive

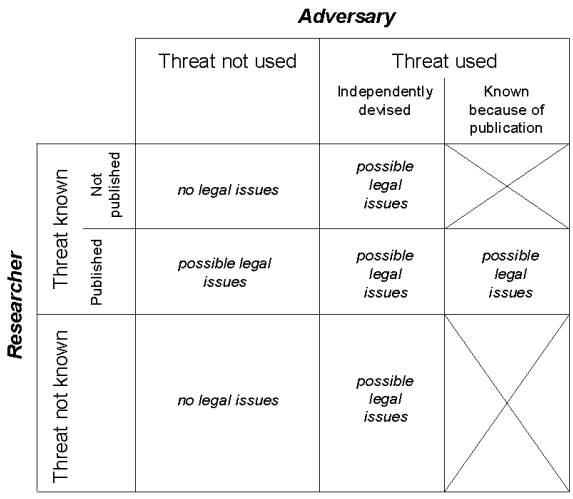

Figure 2: Possible legal issues resulting from researcher's and adversary's action or inaction.

legal advice. This is especially so in the situation of future threat research. Further, laws are jurisdiction-specific and may, therefore, vary vastly from one jurisdiction to another. Because computer hacking[1] and testing is often not contained to actions in one jurisdiction, the laws of potentially all nations may apply to the action.

Freedom of expression

Freedom of expression has typically meant the freedom to publish, which in turn includes the freedom to speak, write and print. More liberal and broad interpretations include the right to communicate as well as to receive (or not receive) content. Freedom of expression is a right protected in a number of international human rights treaties, as well as a legally entrenched domestic right in many jurisdictions in the world, and is considered by many to be a universal moral right. Many theories exist to justify freedom of expression. The utilitarian theory of free speech espouses the idea that speech is a tool to advance truth, democracy and the exchange of ideas [8]. The libertarian model seeks to protect individual self-determination rather than any right [9]. Other frameworks such as Asian values would conceptualize freedom of expression as a narrow concept confined to duties to a community – one which does not extend to individual rights [10]. No matter what theory one subscribes to, freedom of expression has never been absolute. Public, private and self-censorship as well as other limitations look to restrain free speech where it is deemed potentially harmful to society [11].

Is the publication of future threat research harmful to society? The answer remains untested or at least unquantified. Publication may allow defences to be in place before threats are released to the wild. Communication of new and future threats under the 'commons model' may further allow for superior defences to be adopted. Under the commons model, knowledge and information are shared. Sharing of information and resources is based on the premise that, by having the opportunity to build on each other's ideas, rather than duplicate one another's efforts in a 'closed' environment, security vendors are able to produce more efficient and technologically secure products. Conversely, publication may allow for an otherwise unknown threat to be released (if not certainly expedited) to the wild.

One thing is certain: security experts are increasingly expressing the need for a proactive approach to computer security. For example, when commenting on the extent and sophistication of information sharing in the underground organized malware crime world, a United States Secret Service agent at the AusCERT 2007 conference noted that the 'bad guys' were sharing information and computer hacking techniques in a manner allowing them to stay several steps ahead of both the anti-virus and law enforcement community. The agent then commented on the need for those in the security industry to develop a similar culture of information sharing, and the further need to conduct and share future threat research – the results of which would then be communicated to the security community in a timely fashion. There seems to be a growing awareness of the advantage in publishing future threat research.

Freedom-of-expression arguments become interesting where such expression intersects with legislative provisions, and in particular, criminal law. Legislation which unduly impairs freedom of expression may be struck down as

unconstitutional by the courts (not in all jurisdictions but available as a defence in, for example, Canada, the United States and the European Union). When expression, in this case a computer program, becomes an integral part of a crime, a freedom of expression defence will likely be foreclosed unless there is evidence that the expression is directed at ideas remote from the commission of a criminal act [11]. For this reason, a person or group who writes but does not disseminate a virus or trojan, then makes it available to the public, will likely not be able to rely on freedom of expression as a defence. The publication of future threat research at a hacker conference would likely present a closer nexus to becoming part of a future crime, than merely an idea. The publication of future threat research at an AV conference, however, is less ambiguous and would likely be more aligned with presenting ideas. These thoughts are generalizations as all such legal analysis is highly fact-specific.

Criminal acts

A number of criminal acts may be involved with the publication of future threat research. Criminal offences may be divided into two camps: one camp requires *mens rea*, otherwise known as the requisite level of intention to commit a crime, while the other camp is strict liability offences where intention is irrelevant. All law is jurisdiction-specific, but criminal law is particularly so. The following analysis borrows on general principles as seen in a number of both civil and common law jurisdictions.

Computer/data misuse or abuse

Most jurisdictions have enacted computer misuse and abuse criminal provisions. Such criminal provisions generally address situations where any component of a computer (hard drive, software and network) is tampered with allowing for unauthorized access, modification or impairment to data. Most criminal provisions distinguish mere access from modification and impairment of data by looking at the intent and harm caused. The very nature of hacking involves the exploration (and perhaps exploitation) of vulnerabilities which involve unauthorized access to data.

In the past, where a hacker has merely 'looked under the hood of a car without the owner's permission' but has not caused any harm, criminal charges have not been successful. The law in many jurisdictions has been amended to cover a broad range of unauthorized access regardless of intent or harm caused. For example, the Council of the European Union's 2001 Cybercrime Convention was enacted to tackle the growing problem of cybercrime with emphasis on protecting critical infrastructure. Many nations outside of the European Union have also signed the Cybercrime Convention.

The EU has since passed the Framework Decision on Attacks against the Information System 2005/222/JHA to expedite many similar measures imposed by the Cybercrime Convention. The differentiation between mere unauthorized access versus actual modification and impairment of data causing harm lies at the level of penalty. Intent and harm give way to more serious penalties as opposed to less severe penalties where there is a lack of intent and harm (though the penalties are still serious enough to act as a deterrent, e.g. a maximum of ten years in prison).

A principal aim of the Framework Directive is to close legal loopholes for hackers who did not cause sufficiently serious harm or damage which, in the past, allowed them to 'walk free'. The new provisions close the loopholes where 'Minor or trivial conduct is criminalised. This would be in contradiction to the principle of subsidiarity which requires Member States to avoid the risk of over criminalisation. It makes the offence per se criminal, whether or not any harm is intended.' [12, p.387]. The latter statement refers to the offence of instigating, aiding, or abetting in the commission of a crime. There is no exception for security research. Both the Cybercrime Convention and Framework Decision have been criticized for stifling online speech and for disabling effective security research and testing [13]. As the Framework Decision is recent, it remains to be seen if it will be struck down in the European courts as unconstitutional on the grounds of unduly restricting human rights, namely freedom of expression.

Conducting practical security research could be captured by computer misuse and abuse provisions similar to those found in the Framework Decision. Publication of such research may also be considered a criminal activity where the work aided, abetted or facilitated the commission of a crime.

In the event that security research resulted in criminal charges being laid, there is a very strong chance that the provision would be struck down for violation of freedom of expression. This, of course, is no consolation or safety net for security researchers wishing to remain outside of the legal system. However, the criminalization of an activity does not necessarily equate with the political will to prosecute all cases. The practices of some legitimate corporations arguably fall within the parameters of criminal data misuse. For example, products that install rootkits without user authorization are in clear breach of criminal law (e.g. the *Sony* rootkit debacle). There is simply no political will to charge a company such as *Sony*. The same would likely hold true of the lack of political will to prosecute security researchers for activities relating to genuine computer security.

Aiding, abetting or facilitating in the commission of a crime

Most jurisdictions criminalize the act of aiding or facilitating the commission of a crime. While facilitation is a specific crime in most jurisdictions, there may also be crime-specific provisions as is the case in the EU Framework Directive. There is no universal consensus in the meaning and scope of what is considered instigating, aiding, abetting, or facilitating in the commission of a crime. Returning to the discourse found in freedom of expression arguments, the publication of future threat research at a hacker conference is likely closer to becoming part of a future crime than just an idea. The publication of future threat research at an AV conference, however, would likely be more aligned with presenting ideas. In the context of criminal law, the publication venue is less important. The criminal facilitation provision could only potentially be triggered where the publication resulted in an actual crime being committed.

Copyright infringement

The acts involved in both hacking and legitimate security testing may involve several potential breaches to copyright law. Where a copy of even a portion of computer code is made (in some cases even RAM will suffice) the conduct in question may have resulted in copyright infringement. Where

a technological protection measure or digital rights management system is used, the mere act of circumvention will result in an additional infringement.

A famous example of the latter involves Professor Edward Felten and his colleagues when they published and presented their research in the successful breaking of the digital watermark copy prevention on music files otherwise known as SDMI [14]. The circumvention was performed as part of the 'Hack SDMI Challenge' where the recording industry challenged the public to test the security of proposed SDMI copy prevention systems. Felten's team circumvented a number of the SDMI protection mechanisms. In order to claim a prize for successfully breaking these codes, Felten and his team would have had to agree not to disclose the technical details of their circumvention solutions.

The Felten team opted to publish their results instead of accepting the prize. The SDMI member companies sent Felten's team a letter threatening actions under the anti-circumvention measures in the United States Digital Millennium Copyright Act. Concerned that it could be subject to criminal liability if it allowed the Felten paper to be presented at its security conference, the USENIX technical conference organization became involved. After a large amount of negative publicity, the recording industry withdrew their opposition to the presentation of the paper. The lessons to be learned from the Felten story spill into the context of threat research publication.

Most jurisdictions' law on copyright contains what is known as fair use (US – exception rights) or fair dealings (Commonwealth and many civil law countries – defences to copyright infringement). Research and in particular encryption research is an exception or defence to copyright infringement in many jurisdictions, as is the ability to reverse engineer software for *inoperability* purposes. The problem, however, lies in who is entitled to the exception or defence. Universities are clear cut examples of falling within the parameters of fair use or dealings. Hackers who do not have authorization to 'ethically' hack into a system would not be entitled to this exception or defence. As we move towards corporations and AV vendors, however, the line is less clear cut.

To add to this dilemma, many jurisdictions allow companies to contract out of fair dealing provisions. For example, it is common to open up proprietary software containing an end user licence agreement which prevents the user from reverse engineering the code. One countermeasure is found in jurisdictions which have the defence of 'in the public interest'. It could readily be argued that future threat research and the publication of such results in the AV community would be considered as an integral component in effectively combating high-tech crime. It would be difficult to counter this argument where publication is made to a closed venue such as the AV community, where the actual computer code is not revealed, and where there are potentially great benefits from publishing (proactive approach where defences in place).

Copyright law generally contains both civil and criminal provisions. While there may not be political will to prosecute security vendors and those who publish future threat research, the threat of civil liability is less predictable as a wider range of parties may bring suit, all of whom may be motivated by a variety of factors.

The tort of negligence (common law) / delict (civil law)

Where a person or entity presented future threat research he or she knew or ought to have reasonably known would result in its release to the wild (or expedition of), there is the possibility of a civil suit in negligence or delict (quasi-delict if the result is unintentional).

In order for an act in negligence to succeed, it must be shown that there was a duty of care between the parties, and that physical damage was sustained. In the case of publication, there would not normally be a duty of care between the conference presenter and the recipient of a malicious attack under the common law. As one expert writes, 'Under the common law of negligence, a novel duty of care is only usually imposed by the courts where it is reasonably foreseeable that a failure in that duty would cause damage to the person to whom the duty is owed, and where there is no good policy reason to reduce or limit that duty.' [15, p.48; also 16]. This is referred to as the proximity test or remoteness test.

The civil law principle differs in that there is a universal principle of civil responsibility. Expressed in a different way, a duty of care is owed to everyone. Civil law generally recognizes three components: fault, damage and causation. Remoteness or proximity is factored in at the level of causation and the amount of compensation awarded (if any). Publication of future threat research would likely be seen as too remote to have caused a threat to be released into the wild. Whether an adversary uses the publication to develop and release a threat would also generally be unknown and, therefore, difficult to prove in court.

The tort of negligence or delict is also potentially applicable to the recipients of information found in the publication. For example, where a vendor is made aware of a vulnerability as disclosed in a future threat publication (providing the information is reliable and correct), then deliberately chooses to ignore the fact by not developing any defences (or continues to use defences known to be inadequate), negligence could apply. So inaction or failure to act may also trigger negligence. In this situation where the vendor has a direct client relationship with the affected end user, the proximity test is easily made. Further action could possibly also be taken in this circumstance under the law of contracts.

ETHICAL ISSUES

Users are the key component when considering ethical issues involved in future threat research. Users vastly outnumber researchers, and are arguably the raison d'être for researchers; given that, the overall effect of a security threat to users is paramount. Users' computers must be safer or, at the very least, no worse off because of actions taken by a researcher.

Recall that we identified two possible ethical problems in the Introduction. While many ethical issues are potentially raised with threat research, we will restrict our analysis to this limited context.

Underlying assumptions

There are assumptions which must be made in any ethical analysis. Foremost among the assumptions here is the assumption that defensive measures can be taken when a future threat is identified.

The majority of users' computers can now be fairly assumed to have standard security precautions, like anti-virus software. The ubiquity of anti-virus software provides a litmus test: if a meaningful defence against a future threat is anti-virus software (with or without some minor adjustments), then it may be assumed that defensive measures can be taken.

This is definitely not the only test, nor will it be applicable in every case, but it is a strong test. Anti-virus software has demonstrated its ability to update users' computers defensively on a very large scale. This fact, combined with the large number of users who have anti-virus software installed, make this a useful guide for future threat research.

Publishing venue

How are future threats published? It could be argued that the publishing venue is no longer a relevant concern, because even news about major security issues can be broken in informal venues – witness the *Sony* rootkit story first appearing in Russinovich's blog [17]. In comparison to completely public blogs, anti-virus conferences and (applied, non-AV) security conferences are extremely safe publishing venues, since the audience – while possibly containing some 'bad guys' – is primarily the audience of legitimate security researchers that should be addressed. These conferences constitute limited disclosure in that regard. We would not necessarily extend this argument to 'hacker' or 'black hat' conferences.

The following ethical analysis assumes that defences are beneficial and that AV security conferences provide a publication venue allowing for limited disclosure.

To publish or not to publish

For a researcher, it is untenable to base the decision to publish or not publish solely on the possibility of an adversary making use of the information. First, as Figure 1 suggests, this is only one of many possible scenarios, and is not a course of action guaranteed to prevent an adversary from using the threat. Second, whether or not an adversary has used published information when developing a threat is in general unknowable, even after an adversary has used the threat. Even if there is a specific instance where the adversary appears to have used published information, it is also unknowable whether the adversary would have developed the threat themselves anyway.

It seems much more rational to base ethical decisions on information that is known. One aspect that can be known with a relatively high degree of certainty is an adversary's capacity to use a given threat, regardless of whether or not they know about it. There are two extreme ends of the spectrum, which we illustrate with examples from our research.

- A near-term future threat is one where an adversary is currently able to use a threat with little or no effort. For example, in [18] we looked at how more convincing spam could be sent automatically by zombies, using saved email on the zombie computers. Here, it was well known that an adversary would already have access to saved email.
- A long-term future threat, by contrast, is one where the resources to use the threat are not yet generally available to an adversary. [19], for instance, makes an initial attempt to calculate the computing power available in very large botnets, and the attacks that might become possible as a result. This is long-term because adversaries with botnets containing a million zombies are atypical now; this is primarily an academic question at present.

Near-term future threats represent an imminent danger to users that needs to be addressed, and long-term threats are a remote danger that can be safely considered without immediate harm. Either way, it can be argued that the intent of publication is in the best interest of users. The outcome of publication is unknown at the time of presenting. To date, there has been no good empirical work to test the effects of threat research publication.

Inaction is action

In 1942, Asimov published an early version of his Laws of Robotics, which began [20, p.100]:

> 'One, a robot may not injure a human being under any circumstances – and, as a corollary, must not permit a human being to be injured because of inaction on his part.'

The corollary is especially relevant in this context. One could say that a researcher who knows about a future threat but does not publish it has demonstrated inaction. However, barring circumstances where a researcher intended to publish but was as yet unable to do so, not publishing is a deliberate choice by the researcher. Even though the observed effect by the researcher's choice is inaction, the fact that they made a choice constitutes an action subject to ethical analysis.

By not publishing, the researcher has not prevented the threat where an adversary independently discovers a threat and uses it. But the researcher *has* prevented the threat from being considered before its use in the wild, and prevented defences from being established. Users are exposed to the threat, and the proverbial human being is injured through inaction. Because the users' computers are not as safe or safer, we must conclude that not publishing about the future threat is not the correct course of action. Publishing the future threat is the correct action for the researcher in this case.

CONCLUSION

Cryptography observes what is called Kerckhoffs' Law. Originally a design criterion stated in 1883 for military cryptography, it said that a cryptographic method must not require secrecy, and must not cause inconvenience if found out by the enemy [21]. In other words, assume that an adversary knows everything about the cryptosystem.

The anti-virus community is not at that point yet. There is still an undercurrent of security through obscurity, the hubris that adversaries cannot possibly derive new threats themselves. Modern adversaries are driven by a powerful motivator, money, and users' computers are better served by investing time and thought into defending them against current and future threats, instead of hand-wringing over what adversaries know or don't know.

It would be useful for future threat dissemination if there were some single information source where new publications could be announced; there are many security conferences to monitor. Various mailing lists do exist where some announcements are made, like bugtraq, and some researchers have their own notification mechanisms – the first author has

an RSS feed for his publications – but there is no 'one-stop shopping' for future threat work. We recommend that such a single source be established. This would be a positive step towards moving the industry beyond its current practices.

We do not expect that this preliminary examination of future threat research will definitively resolve the matter (it will certainly not convince everyone in the anti-virus community!) but hopefully it will encourage discussion and advance the debate in a meaningful way.

ACKNOWLEDGEMENTS

The first author's research is supported in part by the Natural Sciences and Engineering Research Council of Canada, and he would like to thank his students for discussions about the issue of future threat research. George Ledin pointed out the changes in the cryptographic community, Alfred Menezes confirmed the standing ovation story, and Paul Van Oorschot commented on the cultural aspects. The name 'computology' is from Heather Crawford.

REFERENCES

[1] AVIEN. Anti-virus information exchange network code of conduct. http://www.avien.org/codeconduct.html. Accessed May 2007.

[2] Ritchie, D. M. Dabbling in the cryptographic world – a story. http://cm.bell-labs.com/cm/cs/who/dmr/crypt.html. May 2000.

[3] Schneier, B. Applied Cryptography, 2nd edition. Wiley. 1996.

[4] Levy, S. Crypto: how the code rebels beat the government – saving privacy in the digital age. Viking. 2001.

[5] Staniford, S.; Paxson, V.; Weaver, N. How to 0wn the Internet in your spare time. 11th USENIX Security Symposium. 2002.

[6] USENIX, 'WOOT '07 call for papers', http://www.usenix.org/events/woot07/cfp/, 2007.

[7] Blackburn, S. The Oxford Dictionary of Philosophy, 2nd edition. Oxford University Press. 2005.

[8] Moon, R. The Constitutional Protection of Freedom of Expression. University of Toronto Press. 2000.

[9] Barendt, E. Freedom of Speech, 2nd edition. Oxford University Press. 2005.

[10] Maurushat, A. The benevolent health worm: comparing Western human rights-based ethics and Confucius duty-based moral philosophy. Computer Ethics: Philosophical Enquiry Conference (CEPE 2007), to appear.

[11] Colangelo, A.; Maurushat, A. Exploring the limits of computer code as a protected form of expression: a suggested approach to encryption, computer viruses, and technological protection measures. McGill Law Journal, 51(1):47–96. 2006.

[12] Kierkegaard, S. M. EU Framework decision on cyber attacks: Here come the 'cybernators'!. Computer Law and Security Report, 22(5):381–391, 2006.

[13] EDRI-gram. Council adopts decision on attacks against information systems. March 2005. http://www.edri.org/edrigram/number3.5/attacks.

[14] Kerr, I. R.; Maurushat, A.; Tacit, C. S. Technical protection measures: Tilting at copyright's windmill. Ottawa Law Review, 34(1):6–82, 2002–2003.

[15] Edwards, L. Dawn of the death of distributed denial of service: How to kill zombies. Cardozo Journal of Arts and Entertainment Law, 24(1):23–62. 2006.

[16] Chandler, J. A. Security in cyberspace: combatting distributed denial of service attacks. University of Ottawa Law and Technology Journal, 1(1–2):231–261. 2004.

[17] Russinovich, M. Sony, rootkits and digital rights management gone too far. October 2005, http://blogs.technet.com/markrussinovich/archive/2005/10/31/sony-rootkits-and-digital-rights-management-gone-too-far.aspx.

[18] Aycock, J.; Friess, N. Spam zombies from outer space. 15th Annual EICAR Conference, pp. 164–179. 2006.

[19] Hemmingsen, R. H.; Aycock, J.; Jacobson, M., Jr. Spam, phishing, and the looming challenge of big botnets. EU Spam Symposium, 2007.

[20] Asimov, I. Runaround. Astounding Science-Fiction, XXIX(1). 1942.

[21] Kerckhoffs, A. La cryptographie militaire (part I). Journal des Sciences Militaires, pp. 5–38, January 1883. Available online at http://www.petitcolas.net/fabien/kerckhoffs/crypto_militaire_1.pdf.

FOOTNOTES

[1]In this section, we use the terms 'hacker' and 'hacking' in the popular, widespread sense of the words.