

# “Good” Worms and Human Rights

John Aycock  
Department of Computer Science  
University of Calgary  
2500 University Drive N.W.  
Calgary, Alberta, Canada T2N 1N4  
aycock@cpsc.ucalgary.ca

Alana Maurushat  
Faculty of Law  
University of New South Wales  
Sydney NSW 2052  
Australia  
amaurushat@yahoo.com

TR 2006-846-39, October 2006

## Abstract

The extent of Internet censorship in countries like China is regularly tested, but the testing methods used from within a censored country can entail risk for humans. A benevolent worm can be used for testing instead: the worm’s self-replication, long the bane of suggested benevolent viruses and worms, is shown to be essential here. We describe the design of this benevolent worm, along with some other related applications for it. A full technical, ethical, and legal analysis is provided.

**Disclaimer: the following paper discusses a novel type of computer worm. Release of such a worm, and possibly even its creation, could result in severe legal penalties. We do not advocate the creation and release of this worm, but present it here for research purposes only.**

## 1 Introduction

China is well-known for protecting its citizens from the perils of the Internet. Known evocatively as the “Great Firewall of China,” technology is used to limit certain material flowing into or out of China, such as information about Tiananmen Square.

On one hand, this stance by the Chinese government is to be expected – China has a historical tradition of censorship [26]. On the other hand, it is hard to imagine

marshaling the technology and resources to successfully censor the Internet, yet the Chinese government has billions of dollars invested to try and do just that [8].

In the case of the Great Firewall, even the extent of the censorship is not apparent. Attempts to access forbidden material yield results akin to network or server problems [22, 32]. Accurate glimpses into the censorship mechanism are rare, like the discovery of a list of banned words shipped with Chinese instant messaging software [25]. Moreover, the consensus is that Chinese censorship is a dynamic work in progress, and subject to frequent changes [22, 32].

Groups with interests in human rights, freedom of expression, and privacy monitor the extent of Internet censorship in China and elsewhere. For China, the current methods of testing are listed below. All the tests originate outside China unless otherwise noted.

- Fetch URLs containing forbidden terms from Chinese web servers [4]. This testing is based on the supposition that the Firewall's operation is symmetric, and censors the same material coming and going. It is not a complete test because coarse-grained censorship like blocking of IP addresses is not examined.
- Fetch URLs whose web pages possibly contain sensitive content, via dialup modem to Chinese ISPs. This method was eventually made unusable [32].
- Fetch URLs whose web pages possibly contain sensitive content, through Chinese open proxy servers [22, 32].
- Examine the results from Chinese search engines, when searching for particular web sites and keywords [14]. Here, the testing was done from both the U.S. with a U.S. ISP, and from China using a Chinese ISP.
- From within China, fetch URLs entered manually or fetch URLs *en masse* using a program. The URLs were entered, and the program was run, by volunteers [22].

Where applicable, controls are used to distinguish censorship from legitimate network and server failures [14, 22, 32].

These tests are not without their share of problems. They can suffer from 'limited scope' [22, page 23]. Differences have been observed between proxy server tests and in-state tests [22]; given that over 70% of Chinese in a survey claim not to use proxy servers anyway [18], in-state tests are really the best way to get an accurate idea of what the typical user sees (and doesn't see). However, in-state testing entails risk for the humans who perform it.

In the remainder of this paper, we propose an alternative way to perform this in-state testing with reduced risk to humans, by using a benevolent worm. We lay the groundwork in Section 2 by surveying other benevolent viruses and worms. Our "human rights worm" is presented in Section 3, along with other applications, followed by a detailed analysis in Section 4. Finally, we give our conclusions in Section 5.

## 2 Benevolent Viruses and Worms

The idea of “good” viruses and worms that have a beneficial effect has been around since the earliest academic virus and worm research. For example:

- A virus could be written that compresses executable files to save disk space [5]. Infected/compressed files would be automatically decompressed by the virus as needed. This idea was realized by the Cruncher virus in 1993 [17].
- The KOH virus encrypts floppy disks and hard disk partitions for security reasons [19]. A legitimate user would know the decryption key and could access the files, i.e., KOH was not “ransomware” being used for extortion.
- Early worm research implemented a distributed computing framework at Xerox PARC [28]. After solving some problems controlling the worms, a variety of applications were built including network diagnostics and computing frames of a computer animation.
- A virus could perform system maintenance, like upgrading outdated versions of programs [6].
- Predator worms are revisited periodically, the somewhat romantic notion that good worms can hunt down and destroy bad worms, or that good worms can find and patch vulnerable machines [1, 10, 30]. Real attempts at predator worms, such as the Welchia worm which tried to clean up after Blaster [24], have generally proven disastrous and have resulted in more trouble than the original worm caused.

Although we defer a detailed analysis of any benevolent viruses and worms until Section 4, most of these suggested examples suffer from one basic problem. There is a way to accomplish each of the above tasks without the risk of using hard-to-control virus/worm propagation mechanisms.

## 3 The Human Rights Worm

Computers and robots have long been used in environments where it is too dangerous, hostile, or difficult for humans to perform tasks. But what about cases where the danger stems from fellow human beings?

Right or wrong, the Great Firewall of China *is* tested by people from within the borders of China. The people who do this testing undertake substantial risk unto themselves and their families located in China. Although no one has been prosecuted yet for such testing, it is well-known that Chinese law is deliberately ambiguous and general in a manner consistent with its unpredictable application. The Chinese government has historically used vague legal drafting as a form of inflicting fear of persecution; the best known example is the area of state secrets and subversion [13]. Added to this is a growing trend where ‘police have begun detaining more “ordinary” users’ [31, page

33] of the Internet. Taken together, in-state testing of the Great Firewall is a dangerous pastime.

We propose that a computer worm can perform the task of testing automatically, avoiding the danger posed to humans who take part in testing. We call this the *human rights worm*. The self-replication mechanism of worms is ideal in this case, because a person whose computer is infected takes part in the testing but has perfect deniability: they performed no deliberate action and have no knowledge of the worm.

The human rights worm would have three primary characteristics. First, the worm would be slow-spreading by design, so as not to wreak havoc on normal network operations. Second, it would perform targeted infections of computers within China; before an infection attempt, the worm could identify a target IP address as Chinese or not using geolocation [21]. Even a crude geolocation method like performing a reverse DNS lookup of a target IP address to see if its domain name ends in `.cn` would be sufficient to limit the worm's spread. Third, the worm's payload would perform the Firewall testing periodically and report the results to interested parties – not necessarily the same people who created and released the human rights worm.

Who would create the worm? There is a strong psychological and political element to this question. A worm created inside China would have the distinct advantage of appearing to be change from within; a worm created outside China might seem to be external meddling, imperialism, or worse, an act of war.

The human rights worm has a variety of other applications, most just a simple matter of changing the worm's payload:

**Chaff.** Not surprisingly, the Great Firewall's filtering is not able to detect banned content in encrypted traffic; it is not possible for the Firewall to decrypt the traffic. However, it has been suggested that the Chinese government may eventually detect the presence of encrypted traffic [4]. Even if the traffic cannot be decrypted, a person detected using encryption software inside China may not be looked upon kindly.

The human rights worm could help address this problem, by feeding encrypted chaff to the Firewall. The real encrypted traffic will be lost in the noise generated by the human rights worm, and any attempts at detecting encrypted traffic will be met with a steady stream of meaningless alerts.

**Information delivery.** One strange aspect of Internet censorship is how it inverts the application of technology. Spam is sent into China by 'overseas dissidents and free-speech advocates' [3, page 29]; obfuscations used by spammers are used to avoid detection by authorities [12]; anti-spam techniques are used by the government to censor content [12].

Access to information may involve more than simply freedom of expression; timely information can be a matter of life and death. Indeed, there are two specific areas where censorship and a lack of accurate information distributed in a timely manner have had unrefutable dire consequences in China in recent history:

- **AIDS:** The Chinese government has and continues to suppress information on the spread of HIV/AIDS. By 1987 the government had reported only

four known cases, claiming that AIDS was a foreigners' disease [27]. The reality is that China has one of the highest HIV/AIDS rates in the world outside of Africa. While the impact of accurate and timely information in this epidemic is unknown, it is certainly plausible that access to such important information could have reduced the rate of infection.

- SARS: Similar to AIDS, the Chinese government withheld critical information on severe acute respiratory syndrome in 2002 [16]. This allowed the disease to spread more readily from Guangdong province to other provinces in China and eventually the rest of the world. The SARS health crisis can be partially attributed to the nondisclosure of pertinent information.

While bird flu has not yet reached crisis levels, history indicates that any information provided by Chinese officials should be treated with caution.

As China's filtering/anti-spam technology evolves, the use of spam to disseminate information will become less effective. A new mechanism will be required for large-scale information delivery, and the human rights worm provides one possible solution. Infected machines could display information in pop-up windows, for instance, or override a user's default web page with one displaying information.

**Anonymity.** An anonymity network is a means by which a user can hide what they are connecting to – an attempt at accessing forbidden content might be detected, but an anonymity network would make it prohibitively difficult to trace the request back to its source.

A practical problem arises if mere use of a well-known anonymity network is enough to raise suspicion. The Tor anonymity network [7], for example, supplies a list of Tor servers' IP addresses and ports [29]; a connection to any of those is a clear signal that a bid for anonymity is being made.

Previous work has stated that malicious software can be used to automatically establish an anonymity network [11]. The human rights worm could build such an anonymity network to provide anonymity service temporarily until filtering was changed to detect it.

**Other countries.** Although we have been singling out China in this paper, there are other countries that censor access to the Internet. For example, Vietnam's filtering of Internet content is also tested [23]. The human rights worm would work equally well in these other countries where the political and legal landscape create risks for in-state testers.

## 4 Analysis

Bontchev [2] has compiled the most detailed list to date of the criteria that a benevolent virus must have in order to be useful. Most of these are equally applicable to worms as well as viruses. We present Bontchev's criteria in abbreviated form (set in *italics* below) and analyze the human rights worm in this light.

## 4.1 Technical Issues

To separate technical matters, assume for the moment that it is desirable for the human rights worm to spread and operate as designed. What are the technical issues?

- *Viruses do not spread in a controlled way.*

As described in Section 3, the human rights worm would limit its spread to machines in China. The worm would not attempt to infect any machines that it could not place inside China through geolocation.

- *Viruses could find themselves in an unexpected environment where they could do inadvertent damage.*

A worm that spreads by successfully exploiting a specific bug arguably has a very good understanding of the target environment; buffer overflow exploits can be very fragile, for example.

The human rights worm may encounter systems that were unknown at the time of the worm's release, like an operating system upgrade that the worm was not tested with. To mitigate this sort of problem, the worm can be designed to shut itself down after a reasonable period of time: new, incompatible software is typically not released often or without warning.

- *It is not possible for anti-virus software to distinguish between "good" and "bad" viruses.*

Let us assume that anti-virus software is able to detect the human rights worm. Moreover, let us further assume that anti-virus vendors have *chosen* to detect the human rights worm, having analyzed it to see what it does. (This latter point may raise interesting moral issues for some virus analysts.)

Bontchev's criteria were polarized in that a virus was either good or bad in its entirety. We tend to side more with MidNyte, who argued that there could be a 'bad' virus working for a 'greater good' [20]. In other words, we can more thoroughly analyze the human rights worm by weighing its actions separately from its effects.

We consider the effects of the human rights worm in Section 4.2. As for the worm's actions, they are clearly malicious: infecting computers that do not belong to the worm author, acting without the permission or informed consent of the computer's owner. Anti-virus software can comfortably label this worm as malicious, and does not need to distinguish good from bad.

Interestingly, the reactive nature of much anti-virus software may be a saving grace. If the human rights worm is going to shut itself down anyway to avoid compatibility problems, by the time anti-virus vendors capture a worm sample, create and test an anti-virus update, and have customers install the update, the human rights worm may already have had enough time to accomplish its goals. Any moral issues in terms of detection are quietly skirted.

- *Viruses waste computer resources.*

The slow-spreading human rights worm would occasionally make attempts to infect other machines and periodically perform Firewall testing by trying to access various URLs. Given that the human rights worm is concerned with overall web site accessibility, the success of a basic HTTP transaction would be sufficient for probing purposes. The traffic and resource demands would thus be much lower than a typical user's web browsing.

- *Viruses can contain bugs.*

This is somewhat of a red herring, because all software can contain bugs. In fact, even anti-virus software has had its share of catastrophic errors [9, 15]. There is no reason to believe that the human rights worm could not be debugged and tested to a professional standard prior to release.

- *The parasitic action of viruses can cause compatibility problems.*

A worm is not parasitic and does not attach itself onto existing code in the way that viruses do, so this concern is not applicable to the human rights worm.

- *The same task could be performed without self-replication.*

As discussed in Section 3, the self-replication of the human rights worm is critical, as it firmly establishes plausible deniability for the owner of an infected machine.

## 4.2 Ethical and Legal Issues

Bontchev presents a number of ethical and legal issues worth refuting. The human rights worm poses some additional issues in this regard which we discuss at the end of this section.

- *Unauthorized data modification is unethical and illegal.*

Bontchev is correct to assert that data modification is illegal in many jurisdictions; it may attract civil liability and/or criminal sanctions. This does not, however, automatically lead to the conclusion that a benevolent worm or virus would be unethical. Not everything ethical is legal, and not everything legal is ethical.

Most, if not all, ethical theories allow for the breaking of rules or law where it would be construed as ethical to do so. The use of the human rights worm conceivably falls within the realm of ethically acceptable action under many ethical theories. This is easiest to see if we consider the "information delivery" payload.

Consequentialism would examine the consequences of an action; the consequence of disseminating illegal information on AIDS may lead to a jail sentence but it could also prevent the spread of the disease and reduce the rate of infection, thereby maximizing both health and welfare. Deontology advocates a duty to moral rules, allowing for the possibility that moral actions may be illegal as in the case of the human rights worm. Virtue ethics looks to maximize benevolence. Confucian ethics points to the rights and well-being of a community – the use of a human rights worm to disseminate important information contributes such well-being in the community.

- *A virus can create copyright and ownership problems.*

Bontchev argues that modification of a program could result in a loss of copyright and ownership. There is no legal basis for this claim. Copyright is not lost when someone modifies a work. The modification of a work (e.g., copying, making available to the public without the copyright holder's consent) is in direct violation of nearly every copyright statute in the world. By no means would ownership of a software program affected by a virus or worm be compromised as a matter of law. By the same account, a benevolent worm may violate copyright, but this does not naturally lead to the conclusion that it would be unethical. And in many jurisdictions, illegal acts performed in the interest and welfare of the public is a complete defense.

Bontchev further argues that technical support rights for programs affected by viruses or worms could be voided. The latter sentiment may be true but it is one of internal corporate policy, not a matter of law. It may be the case that a contractual provision which canceled technical support in the presence of malware would itself be null and void under the law (e.g., an unconscionable provision, or contrary to consumer protection law). Further, if the software used was poorly designed to be vulnerable to malware, the software vendor may themselves in turn be liable.

In the context of the human rights worm, vulnerability in software is an asset for its successful propagation. One potential concern is whether the Chinese government might impose liability on software vendors who produce software vulnerable to such attacks – in other words, shift the onus onto private corporations.

- *A “good” virus can be altered to carry “bad” code.*

As discussed in Section 4.1, the human rights worm should be treated as malicious. It enjoys no special status on infected computers, and so malicious alteration of the worm is a moot point.

- *Allowing “good” viruses would justify writing viruses of any kind.*

We would not argue that all viruses contain elements of benefit and research. However, we do advocate mature reflection on the social values presented by new kinds of viruses and worms on a case-by-case basis.

In the case of the human rights worm, there are a number of persuasive arguments as to why its propagation reflects the greater public interest, is ethical, and is a responsible approach to some of the problems currently faced in China and other areas of the world.

The greatest dilemma faced by an individual or group in China considering writing and distributing a human rights worm is the harsh penalties they might face if caught. It is plausible that probing the Firewall, exposing vulnerabilities in the system, and propagating illegal material could be categorized as an act of terrorism and a breach of national security. The penalties in China for these crimes are severe and, in extreme circumstances, may even include the death penalty. While a worm created from within



China could be viewed as change from within, the potential risk to Chinese citizens makes worm creation outside China more palatable and certainly less risky.

### 4.3 Psychological Issues

Bontchev lists two psychological issues with respect to viruses. First, users feel they cannot trust the virus code. Second, the very term “virus” has a negative connotation (the term “worm” does not fare much better).

These psychological issues would be of substantial concern *if* the user was a voluntary participant in the execution of the human rights worm. From a higher level, the adversaries are the creator of the human rights worm and the Chinese government; the users and their computers are the fabric upon which their adversarial battle is played out. The negative connotation of a “worm” may even favor the worm’s creator in this case!

The problem of trust might further be addressed if the worm originated from a trusted source, like an established human rights NGO. In any case, the psychological issues are minor considerations when stacked against the higher arguments in favor of the human rights worm.

## 5 Conclusion

The extent of Internet censorship in China and other countries can be automatically tested on a large scale, in-state, without risk to humans. The key is to employ a benevolent worm – the human rights worm – to self-replicate and perform the testing. Such a worm has other anti-censorship applications too, like supplying pertinent and timely information. From our analysis, the human rights worm is technically sound, and while illegal, is ethically justifiable.

## 6 Acknowledgments

The first author’s research is supported by a grant from the Natural Sciences and Engineering Research Council of Canada. The second author’s research is supported by a grant from the PROCURE – France/Hong Kong Joint Research Scheme. Andreas Hirt provided technical advice on anonymity networks.

## References

- [1] D. Aitel. Nematodes – beneficial worms. <http://www.immunityinc.com/downloads/nematodes.pdf>, 2006.
- [2] V. Bontchev. Are “good” computer viruses still a bad idea? In *Proceedings of the EICAR '94 Conference*, pages 25–47, 1994.

- [3] M. S. Chase. *You've Got Dissent!: Chinese Dissident Use of the Internet and Beijing's Counter-Strategies*. RAND, National Security Research Division Center for Asia Pacific Policy, 2002.
- [4] R. Clayton, S. J. Murdoch, and R. N. M. Watson. Ignoring the great firewall of China. In *6th Workshop on Privacy Enhancing Technologies*, 2006.
- [5] F. Cohen. Computer viruses: Theory and experiments. *Computers & Security*, 6(1):22–35, 1987.
- [6] F. B. Cohen. *A Short Course on Computer Viruses*. Wiley, second edition, 1994.
- [7] R. Dingedine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *13th USENIX Security Symposium*, pages 303–320, 2004.
- [8] Economist Intelligence Unit. Big brother learns to surf. *Business China*, 29(1):3–5, 2003.
- [9] J. Evers. McAfee update terminates Excel. CNET News.com, 10 March 2006.
- [10] A. Gupta and D. C. DuVarney. Using predators to combat worms and viruses: A simulation-based study. In *20th Annual Computer Security Applications Conference*, 2004.
- [11] A. Hirt and J. Aycock. Anonymous and malicious. In *15th Virus Bulletin International Conference*, pages 2–8, 2005.
- [12] S. Hom and A. Tai. Human rights and spam: A China case study. In *Spam 2005: Technology, Law and Policy*, pages 63–69, 2005.
- [13] F. Hualing. Counter-revolutionaries, subversives, and terrorists: China's evolving national security law. In *National Security and Fundamental Freedoms: Hong Kong's Article 23 Under Scrutiny*, pages 63–91. Hong Kong University Press, 2005.
- [14] Human Rights Watch. Race to the bottom: Corporate complicity in Chinese Internet censorship, 2006.
- [15] Japan Times. Bug in antivirus software hits LANs at JR East, some media, 24 April 2005.
- [16] S. Kalathil. Battling SARS: China's silence costs lives. Originally appeared in the *International Herald Tribune*, 3 April 2003.
- [17] E. Kaspersky. Cruncher – the first beneficial virus? *Virus Bulletin*, pages 8–9, June 1993.
- [18] G. Liang. The CASS Internet report 2005: Surveying Internet usage and impact in five Chinese cities, 2005.
- [19] M. Ludwig. *The Giant Black Book of Computer Viruses*. American Eagle, second edition, 1998.

- [20] MidNyte. Argument for a ‘good’ virus, 1999.
- [21] J. A. Muir and P. C. van Oorschot. Internet geolocation and evasion. Technical Report TR 06-05, School of Computer Science, Carleton University, Apr. 2006.
- [22] OpenNet Initiative. Internet filtering in China in 2004–2005: A country study. <http://www.opennet.net/china>, Apr. 2005.
- [23] OpenNet Initiative. Internet filtering in Vietnam in 2005–2006: A country study. <http://www.opennet.net/vietnam>, Aug. 2006.
- [24] F. Perriot and D. Knowles. W32.Welchia.Worm. Symantec Security Response, 28 July 2004.
- [25] X. Qiang. The words you never see in Chinese cyberspace. *China Digital Times*, 30 August 2004.
- [26] K. M. Reed. From the great firewall of China to the Berlin firewall: The cost of content regulation on Internet commerce. *The Transnational Lawyer*, 13(2):451–476, 2000.
- [27] E. Settle. AIDS in China: An annotated chronology 1985–2003. [http://www.casy.org/chron/AIDSchron\\_111603.pdf](http://www.casy.org/chron/AIDSchron_111603.pdf), 2003.
- [28] J. F. Shoch and J. A. Hupp. The “worm” programs – early experience with a distributed computation. *Commun. ACM*, 25(3):172–180, 1982.
- [29] Tor. Tor directory protocol, version 2. <http://tor.eff.org/svn/trunk/doc/dir-spec.txt>, 2006.
- [30] H. Toyozumi and A. Kara. Predators: Good will mobile codes combat against computer viruses. In *Proceedings of the 2002 Workshop of New Security Paradigms*, pages 11–17, 2002.
- [31] B. Wong. The tug-of-war for control of China’s Internet. *China Rights Forum*, (1):32–34, 2004.
- [32] J. Zittrain and B. Edelman. Internet filtering in China. *IEEE Internet Computing*, pages 70–77, March/April 2003.