Internet Instability and Disturbance: Goal or Menace?

Richard Ford
Florida Institute of Technology
Dept. of Computer Science
150. W. University Blvd, Melbourne, FL
+1 321 674 7473
rford@fit.edu

Mark Bush
Florida Institute of Technology
Dept. of Biological Sciences
150. W. University Blvd, Melbourne, FL
+1 321 674 7166
mbush@fit.edu

Alex Boulatov
Florida Institute of Technology
Dept. of Computer Science
150. W. University Blvd, Melbourne, FL
+1 321 674 7473
alexb@se.fit.edu

ABSTRACT

Self-replicating code has become an unfortunate part of today's online environment. Viruses and worms have the ability to become pandemic within minutes of first release, and our protection systems are primarily reactive in nature. Thus, there is little or no protection from a new worm which uses a remote exploit in order to spread. Furthermore, such rapidly-moving threats have a documented ability to cause systemic outages; ultimately, such attacks may threaten the overall stability of the Internet itself.

Currently, most exploits leveraged by worms have been well-known and easily solvable *if* the system maintainer had followed best security practices (e.g. deployed a firewall and/or carried out timely patching of vulnerabilities). Thus, actions which drive practitioners toward tighter security are likely to have a positive long-term impact on the overall stability of the global network.

In this paper, we take the unusual position that low-level virus and worm outbreaks are highly *beneficial* to the overall goal of preventing catastrophic Internet failure. To illustrate this position we draw from a biological analogy: the Intermediate Disturbance Hypothesis. This hypothesis argues that within many natural systems it is a continual cycle of disruption which drives diversity... and hence stability and resilience. Finally, we conclude that the deliberate release of viruses and worms that are not threatening holistically may be a necessary approach to protect the Internet from catastrophic outbreaks. This position is supported by empirical evidence from the computer world and by further comparison with biological systems.

Categories and Subject Descriptors

C.2.2 **[Computer-Communication Networks]**: General – *Security and Protection.*

General Terms

Management, Design, Security, Human Factors.

NSPW 2005 Lake Arrowhead CA USA
© 2006 ACM 1-59593-317-4/06/02....$5.00

Keywords

Malicious Code, Viruses, Worms, IDH, Intermediate Disturbance Hypothesis.

1. INTRODUCTION

Computer viruses and worms have become an all-too-familiar aspect of the Internet. Most computer users – especially those using the prevalent Windows platform – are forced to rely on anti-virus software in an attempt to protect machines. However, even the most modern solutions are primarily reactive in nature, and therefore provide little or no protection from new viruses and worms. Because of this flaw, it is possible for worms which evade detection to become pandemic within minutes of their release. An example of such a worm is SQL.Slammer, which had a minimum population doubling time of less than 10 seconds. This outbreak, which occurred on January 25th, 2003, caused widespread network disruption, and even impacted global Internet routing protocols [8].

Despite work carried out on proactive detection of new malicious mobile code [6, 7], the possibility of a massive and catastrophic worm outbreak still exists. In this paper, we propose that the practice of deliberately and periodically *destabilizing* the system temporarily via the release of various forms of Malicious Mobile Code (MMC) may in fact result in higher overall system stability. Drawing from the Intermediate Disturbance Hypothesis, we create a scientific framework for this idea before exploring some of the practical implications of this approach. This approach is in contrast to – but related to – the well-accepted idea that hackers have helped drive computer security improvements. First, in our approach we provide a real scientific model for our observations (via the intermediate disturbance hypothesis). Second, hacking is by its nature a *controlled* operation. The idea of releasing MMC is very different, as the author has no control after its release; it affects the system globally. Finally, MMC has historically been viewed as universally bad, with no redeeming features.

2. THE RISK OF CATASTROPHE ONLINE

Despite the rapid advances in technology, more abundant connectivity and widespread support of standards has created an environment which is almost perfect for the dissemination of MMC. As outlined earlier, we have already observed worms which have the ability to perturb the very heart of the network. Such worms are far from optimal, and there has been significant discussion on ways to improve propagation efficiency (see, for example, [13]). Sadly, current outbreaks seem to be only a foretaste of what is possible.

Given the potential for such rapid spread, it is worth considering the *fragility* of the Internet. Despite the distributed nature of the system, the presence of many infected machines is a powerful force multiplier. Furthermore, there are several critical points on the network that have a broad effect on overall operability. For example, the root Domain Name Servers (DNS) are a *crucial* resource which provides for the translation of human-readable domain names into IP addresses. Without DNS, much of what users consider to be the Internet would not function even though machines would technically still be connected. Similarly, inter-domain peering points represent another critical area: protocols such as Border Gateway Protocol must run in order for traffic to pass between autonomous systems. While there is ongoing research on ways to make the Internet as a whole more tolerant to attacks, despite its apparent resilience, the system is rather more fragile than one might think.

The danger of a widespread worm outbreak is that a very large number of machines – or worse yet, routers – could become infected. Such an infection would provide the perfect mechanism for attacking critical points online, such as the DNS system. Note that such attacks work due to scale: it is difficult for one or even ten machines to overload any of the Internet's critical systems. However, it is easy for ten thousand or one hundred thousand machines to create havoc. Coupled with an attack such as *smurf*, such a network of compromised computers could cause massive disruption [10]. Furthermore, once disruption has begun, it becomes difficult to ameliorate the damage, as communication and distribution of critical fixes is delayed or prevented. Essentially, as the system begins to fail, positive feedback makes it more likely to *continue* to fail.

3. NATURE AND STABILITY

Natural systems are very different from our current virtual environment. Biological systems tend to be richly diverse and highly resilient to change or disturbance. By way of contrast, artificial systems – in particular, our online ecology – tend to be very brittle and susceptible to catastrophic failure. At the code level, for example, there is zero tolerance of errors or change; computers by their nature are binary, and this mindset spills over into the way in which we have designed their operating environment. Protocols and communication must generally occur precisely or there is no communication at all.

In natural systems there is fairly solid evidence that part of the inbuilt resilience is driven by the constant random disturbances these systems face. This concept was formalized in 1987 by Connell [5] as the "intermediate disturbance hypothesis" (IDH). Loosely expressed, this theory argues that an ecosystem maintains its highest species diversity (and therefore maximum resilience to change) under conditions of moderate disturbance.

In Connell's hypothesis, the argument is that all systems are subject to disturbance by events such as fire, disease, trampling by herds of animals, climatic change or volcanic eruptions. These disturbances provide an opening, or a "gap", for other pioneer species to invade. Without such gaps, these pioneer species would become extinct in that environment, lowering species diversity and therefore reducing system stability and resilience.

Consider, for example, a forest. Even the death of an old tree provides a disturbance which can be leveraged by the overall system. The gap formed provides a home for pioneer species. These pioneers are better suited to the gap in the forest canopy, and provide a mechanism for succession whereby the forest can mature. The process of forest recovery frequently requires rebuilding soil nutrient levels through processes such as nitrogen fixation, an attribute of many pioneer species. A further effect of having pioneer species in the community is that they are often less laden with defensive compounds and hence a disproportionately important component of local food webs. Thus a local post-disturbance community rich in pioneer species is an essential step toward restoring both the system and the full diversity of the area.

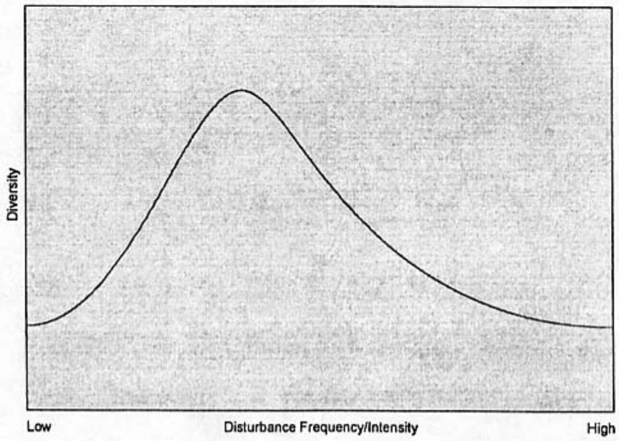

Figure 1: Diagrammatic representation of the effect of disturbance frequency/intensity on species diversity in a biological system.

IDH provides a simple and compelling model of how mature systems maintain high diversity and resilience. Further, it implies that both too high and too low a level of disturbance is detrimental to the diversity of the system.

Continuing with the forest analogy, it is useful to consider a crucial part of forest stability: the disturbances caused by wildfires. Such fires need two components: a source of ignition and a source of fuel. Ignition is typically lightning, and the fuel is the accumulation of organic material such as fallen leaves, twigs and branches. While the dryness of the fuel determines whether a fire will begin, once a fire is established even damp material will ignite.

Once started, it is the amount of fuel per hectare that determines the intensity of the fire. At low fuel levels, the fire creeps along the ground and does not do much more than scorch established trees while allowing animals to flee. Such fires do not cause long-term ecosystem damage, and indeed within fire-adapted systems such as the long-leaf pine barrens of the southeastern United States, these fires help stabilize the system.

By way of contrast, if fuel builds up beyond a certain level, a fire can actually destroy all trees in the forest. Furthermore, organic components within the soil may be broken down, leaving a poor soil structure that is susceptible to erosion. Unlike smaller fires, systems damaged by such intense wildfires may take hundreds of years to recover. Thus, from a system stability standpoint, more frequent but smaller fires provide for the maximum stability of the forest, but allowing for pioneer species and diversity, but without destroying the natural ecosystem overall[1].

To this end, the practice of triggering controlled burns in some environments has become commonplace. The logic behind such an approach is quite simple: to help maintain diversity and stability by introducing smaller perturbations to the system, instead of waiting for a single large perturbation. Given that such an approach works for forests, it is worth considering if similar reasoning can be applied to our online ecosystem.

4. THE BENEFIT OF DISRUPTION

When considering global security, it is important to acknowledge two factors. First, our argument centers primarily on self-replicating threats which are far from the only threat to global computer security: preventing hackers from penetrating important machines is a crucial component of any national security policy. Thus, a vulnerability does not have to be "wormable" (that is, usable by a self-replicating program) to be worth addressing, at least from a national infrastructure perspective.

Second, we are discussing the case where instability is caused by the combined action of a large number of machines. This is not the only failure mode of the Internet – there are several attacks (for example, attacks which target the root DNS Servers) which would significantly impact the network but that would only require a handful of machines to be vulnerable.

With these caveats, we consider Moore's measurement of server patching during the outbreak of the Code-Red worm [12]. Despite the fact that the vulnerability exploited by the worm was well known, Moore measured a significant numbers of hosts and found approximately 80% of IIS installations that were infected during the initial outbreak were later found to be still unpatched. Only after Code-Red began to spread *again* did the patch rate pick up, with finally over 90% of hosts being patched. In essence, the Code-Red outbreak provided a powerful stimulus toward patching, raising the bar of security worldwide.

Similarly, experience in the anti-virus industry has also shown that large malcode outbreaks (or more importantly, perhaps, outbreaks which were widely reported in the popular press) generate significant spurts of virus scanning and improve security globally. For example, Kephart and White [9] show a significant reduction in the number of viruses worldwide for a period of time after the coverage surrounding the Michelangelo virus. Empirically, this phenomenon is well-known within the anti-virus industry, and as such, comes as no surprise to either system administrators or researchers.

A note of caution, however, is that considering the Internet to be analogous to a biological system is just that: an analogy. There are several important differences between online systems and our environment. Perhaps most obviously, our online system is in many ways far *simpler* than a biological system. It lacks the complex interrelationships and rich species diversity. In addition, the Internet is entirely driven by human-selected input; it has been created and is actively managed at every level; it is the result of direct design. Thus, one might expect several limitations in the *details* of the application of IDH. However, at least empirically, our experience has been that the basic premise of the IDH holds true online.

5. CONTROLLED BURNS ONLINE

Based upon our preceding discussion it seems evident that the presence of a large number of remotely-exploitable, poorly protected machines is analogous to a forest where the fuel per hectare has exceeded a critical threshold; that is, where a single fire could wipe out the entire ecosystem.

One interesting aside when considering forest fires, and one possible objection to our analogy, is that in nature, there is no concept of "controlled burns" – natural systems do not actively seek out change. Instead, natural processes slowly move towards optimal solutions by a process driven by evolutionary pressures. Thus, one could argue that in fact IDH is not applicable to our discussion, as the forest fire example does not really describe a *natural* system, but in fact man's attempt to control a system.

In fact, this objection is not valid, as controlled burns usually seek to restore the natural balance that is perturbed by man's impact on the environment. For example, our fire suppression efforts tend to reduce the occurrence of typical small fires; similarly, roads and cuttings create natural barriers to the spread of smaller blazes. A perfect "controlled burn" scenario actually seeks to restore the natural balance, not impose man's will upon a natural system.

To put this discussion in computer terms, we argue that the absence of security threats would lead directly to lax security standards. That is, the amount of "fuel" online for a particular worm could become very large, allowing the potential for a single well-written piece of Malware to create *catastrophic* damage very real. Human nature is such that if there is no perceived threat, little effort will be put into the deployment of protective countermeasures. Traditionally, threat models generally change slowly enough that there is sufficient time for the system to adapt to the new threat without a catastrophic failure. Worms change this by their documented ability to spread worldwide in minutes. In such a situation, there is not time to launch an orchestrated educational campaign or encourage users to patch. Precautions must be taken before such an outbreak.

We believe that based upon an analogy to IDH in nature, it could be in the best interests of global security to deliberately release a worm that spreads via a particular set of vulnerabilities *under certain circumstances*. Such a worm would disrupt the steady state of the system and force administrators to update. Correctly structured, damage from such a worm would be limited and significantly less serious than damage from a malicious worm which used the same replication vector. Thus, deliberate disruption of the system in a controlled manner may be a better strategy than allowing a critical situation to remain unchecked.

Our proposal is therefore to intentionally release a controlled threat in order to drive disruption in the system in certain critical instances. Such a worm would *not* patch susceptible systems, though it would potentially render them unexploitable for a certain period of time (the rationale behind this decision is outlined later). Just like starting a forest fire deliberately in order to prevent large-scale ecosystem disruption later, we propose launching a piece of MMC to prevent a more virulent and damaging piece of MMC being released in an uncontrolled way.

[1] For a more complete overview of this topic, the reader is directed to [4].

Here, perspective is everything. It is vital that the well being of the macroscopic system is considered before that of the microscopic system. In a forest fire, for example, the fire is beneficial at a large scale; at the microscopic scale, such as that of a hapless ant which is incinerated in the blaze, the fire is catastrophic. Similarly, an online "burn" of susceptible machines may well benefit the system holistically but at catastrophic cost locally. We argue that this is a necessary evil – that the local losses are, in effect, inevitable in the case of a large outbreak. This important ethical issue is discussed more completely below.

5.1 Risks of Controlled Burns

One of the challenges with implementing an online campaign of "safe burns" is ensuring that a theoretically benign outbreak does not cause massive global damage. For example, our attempts to "balance" natural ecosystems have not always been successful. A classic example of this is the introduction of the cane toad, *Bufo Marinus*, to Queensland [3]. The toad was intended to control cane grub (native to South and Central America). This introduction proved ineffective as a biological control and the toad instead predated native amphibians, fish, and anything else smaller than itself. It has toxic pouches that emit a deadly neurotoxin if it is bitten and so the natural potential predators, e.g. the quoll (marsupial spotted cat), kookaburras, tiger snakes and goannas, were killed as they mouthed it. The cane toad has now spread throughout the more humid areas of tropical and subtropical Australia, and is poised to invade the World Heritage wetlands of Kakadu.

Thus, tampering with systems is not without significant risk. Anyone deliberately releasing MMC with the intent of improving overall security should do so with considerable caution. At a minimum, we list here several factors that an implementer would have to consider in order to avoid such an occurrence.

First, we must recognize that there is a delicate balance between perceived risk on the part of the end-user (the machine operator) and the actual risk. For example, if the MMC did nothing but spread, causing no disruption in the process, there would be little stimulus toward change. Similarly, if too much damage is caused, the disturbance would be larger than optimal – and potentially worse than the outbreak that we would be attempting to prevent.

Second, it is important to measure the overall susceptibility of the global system to a particular threat. Simulation should provide useful information on expected spread rates and outcomes, but it would be important to be cautious in any approach. Too many attempts at using self-replicating code safely have failed due to underestimation of the tenacity of code once released in the wild.

Third, given that a controlled burn would, by its nature, occur without warning and without attribution, there is an issue of coordination – what is to stop two different groups deciding to employ the same strategy? This concern is real, but is somewhat mitigated by careful monitoring in real time. Unless the disturbances happen almost at the same time, the presence of the *first* disturbance will modify the acceptable parameters of the second: the different groups will limit each other's actions.

Finally, perception of a new threat may be as important as the actual threat. It is impossible to consider computer security without considering the human factors which ultimately control it; thus, perception of risk is as important a factor in determining user action as actual risk. If a deliberately-released MMC sample is seen to be "safe" it may not be a sufficient driver of change. Public knowledge of a "Cyber IDH" program may provide a veneer of safety that renders this approach useless. Therefore a more clandestine effort to perturb the system may ultimately be the best approach. The goal of the perturbation is the continual improvement of security.

Given the preceding discussion it is worthwhile considering the way in which a "controlled burn" decision might be taken.

Any deliberate disturbance to the system would have to be in response to a particular threat. For example, a new vulnerability may have been announced. Once discovered, the next step would be to estimate the overall threat posed by this vulnerability to the overall system. If the exploitability of the bug was very difficult or required special circumstances, it might be determined that no global threat exists due to the problem. If the vulnerability was common and easy to exploit, a risk to the global infrastructure might exist. Furthermore, if patching rates were low, it might be important to drive the patching process by creating a disturbance.

6. ETHICAL ISSUES OF CYBER-IDH

The question of deliberately creating disturbances in order to raise overall security raises several interesting ethical issues. In this section we will examine just two: the question of who has the right to make the decision to release a virus or worm and questions regarding virus writers using this approach to justify their activities.

The primary objection of deliberate MMC release is that it requires some damage to occur in order secure the system as a whole. Who has the right to make such decisions in an online, decentralized community?

In fact, it is the very decentralization of the community which makes this approach critical. There is little or no administrative control between disparate sections of the network. At the macroscopic level, international borders create a patchwork of legislative control that renders many legal approaches to MMC control ineffective. However, even within a single legislative domain, different service providers with different acceptable use policies lead to an environment where there is no single point of control with respect to levels of security. Furthermore, insecurities in another domain can *directly* affect the usability and security of our own domains. Whether we like it or not, machines on the network are tightly coupled with respect to security.

While proponents of a Cyber IDH approach could be accused of "playing God" with the system it is just as valid to argue that *not* following this approach is just as active a choice in terms of outcome. Thus, as in all things, our inaction is itself an active steering of the long-term outcome for the ecosystem.

Examining Cyber IDH in the light of various ethical models yields differing results. For example, Kant's Second Categorical Imperative would tend to suggest that such deliberate release of MMC for the "greater good" was wrong. However, instead of a review of classical ethical models, we turn to the ethical guidelines issued by the ACM, as these seem more immediately applicable at this juncture.

The ACM has two specific sets of ethical guidelines that are apropos. The first, the main ACM Code of Conduct [2] specifically references viruses under a clause titled "Avoid Harm to Others". Here, we read:

1.2 Avoid Harm to Others

...

Harmful actions include intentional destruction or modification of files and programs leading to serious loss of resources or unnecessary expenditure of human resources such as the time and effort required to purge systems of "computer viruses."

Well-intended actions, including those that accomplish assigned duties, may lead to harm unexpectedly. In such an event the responsible person or persons are obligated to undo or mitigate the negative consequences as much as possible.

Note here that Cyber IDH is not an easy decision. One can coherently argue that from one perspective, the release of MMC is harmful. However, it is also entirely reasonable to state that *macroscopically* not to do so is, in fact, allowing harm to befall others.

Similarly, the ACM Software Engineering Code of Ethics [1] specifically instructs SE's to work "in the public interest". It seems reasonable that stabilization of the global network would come under this imperative. Thus, while the issue is certainly ethically complex, it is not obvious that it should be prohibited on ethical grounds.

The second ethical objection to this approach is that virus writers could use its existence to justify their acts of cyber-vandalism as being "beneficial". That is, that the deliberate introduction of disturbance would legitimize certain acts of cybercrime. While such arguments are inevitable, these arguments are specious and easily dismissed.

Continuing with our forest fire analogy, arguing that the virus writer who indiscriminately releases a virus "in the wild" is "helping" is tantamount to arguing that the arsonist (or careless camper) who starts a fire should be applauded if conditions are such that the fire is of overall benefit to the forest. Such an argument is obviously flawed: while the outcome of the act may be beneficial, the arsonist lacks the requisite knowledge to *safely* carry out such a task *even if the intent was to help the overall situation*. Thus, while there is a scientifically justifiable argument that virus writers have *inadvertently* helped stabilize the Internet by preventing a catastrophic outbreak, this outcome does not validate their actions. If an IDH-based approach to Internet stability were ever to be publicly acknowledged, great care would need to be taken in order to handle adverse publicity and control perception in both the general user community and the "blackhat" community. Virus writing would still be irresponsible in most cases, and should be clearly seen as so.

One possibility of avoiding these ethical issues is to consider Cyber IDH on a significantly smaller scale – that is, at a company level, where all machines were under the jurisdiction of a single entity. This approach has two primary drawbacks. First, the macroscopic benefits of the approach are removed. If every entity were sufficiently dedicated to follow this approach, it is likely that a more traditional security technique could be employed. Second, MMC is notoriously difficult to control; a single company would be unlikely to have the expertise necessary to deliberately release MMC within its borders without risking a wider infection.

7. COMPARISON TO OTHER WORK

The concept of releasing worms to stop other more dangerous worms is not entirely new. In [11], Liljenstam and Nicol simulate the result of releasing a "patching" worm which attempts to patch susceptible systems as well as remove worms from already-infected machines. They conclude that such a patching *could* be an effective response to a new worm if deployed quickly and aggressively enough. Their justification is based entirely upon underlying epidemiological models. Similarly, Toyoizumi proposed the introduction of a predator – a self-replicating entity that consumed viruses and controlled outbreaks [14].

While this work is interesting, our proposal takes the idea several important steps further, and is different in several ways. First, our goal is not to fix the underlying problem via MMC, but to force actions which drive security in general. Thus, although our proposed MMC carrier would temporarily render a system invulnerable to a particular attack, ultimately the goal is that the person responsible for the machine would have to fix the underlying problem which allowed for the disturbance. This is a better approach than a patch as the "patching worm" would not necessarily be able to predict the entire outcome of the patch. If such a patch was applied silently and incorrectly it is easy to imagine that system corruption could result. It is better for the operator to *know* that a system has been patched and to be aware of possible problems arising from it. Furthermore, manually recovering the machine from an attack are likely to drive other patches or security enhancements in addition to the one deployed by the worm.

Second, our approach creates a meaningful biological analogy between the stability of natural and artificial systems – in particular, it creates a model for how these systems respond to disturbance. Thus, we believe that we have provided a meaningful framework and approach which is not reactive to a particular exploit, but reactive to a particular vulnerability, in addition to providing overall benefits that help prevent future worms and viruses from spreading.

In addition, the idea of "striking back" at attackers has been suggested before – for example, Welch et al [15] examined some of the issues regarding this approach. However, Cyber IDH goes considerably further: systems potentially compromised are completely unrelated to any current or possibly even past attacks – they are simply vulnerable, and this vulnerability makes them a target for Cyber IDH.

8. FUTURE WORK

Biological systems can be a rich source of inspiration when looking for new ideas in Computer security. In particular, considering those features that provide for stability in biological systems is a good starting point when trying to improve the holistic stability of our network infrastructure. Theories such as IDH can provide a framework for understanding some of the properties of our virtual ecology.

In relation to IDH, there is significant scope for further work in this area. Our current ability to predict the outcome of virus outbreaks is rather low. Despite the fact that spread relies on many random factors, it should be possible to simulate the global threat posed by a particular vulnerability or group of vulnerabilities. In order to do this, more accurate global data on the "exploitability" of our global infrastructure is required. This data includes, but is not limited to, topology, firewall configuration, machine deployment and patching rates. Using these data it should be possible to create a real-time measure of the global vulnerability of the Internet to catastrophic failure. When this measure exceeds some predetermined critical threshold it may be prudent to deliberately trigger an event which perturbs the system. It is our belief that creating such a measurement system that quantifies risk is a critical – and missing – part of our national infrastructure.

9. CONCLUSION

Despite the seeming contradictory approach of releasing MMC to prevent MMC damage, we have presented a strong case for the deliberate perturbation of Internet-connected systems under certain circumstances. In particular, we note that in the absence of any worm or virus outbreaks (i.e. in conditions of low disturbance) system security and countermeasure deployment may become lax. Thus, in these conditions it may be necessary to deliberately drive change within the system.

Our approach is not random and is not worm-specific. Rather, it provides a proactive mechanism for helping prevent a catastrophic failure of our national computing infrastructure based upon measurable quantities. Via careful simulation and data gathering, we believe that a safe and demonstrably effective program could be put in place. Furthermore, such an approach of deliberate disturbance could be used to deliberately shape certain aspects of defensive posture.

The primary drawbacks to our proposed approach are not technical, but ethical and legal. Despite these concerns, the criticality of the problem is such that extreme measures must be taken in order to prevent massive outages. To ignore such a threat is itself unethical: the needs of the many outweigh the needs of the few.

10. ACKNOWLEDGMENTS

The authors wish to thank Dave Ladd (daveladd@microsoft.com) at Microsoft Research for his sponsorship in funding this work.

11. REFERENCES

[1] ACM Software Engineering Code of Ethics and Professional Practice, downloaded August 05 from http://www.acm.org/serving/se/code.htm

[2] ACM Code of Ethics and Professional Conduct, downloaded August 2005 from http://www.acm.org/constitution/code.html

[3] Berra, T.M., *A Natural History of Australia.* UNSW Press, Sydney, 1998

[4] Bush M., *Ecology of a Changing Planet, Third Edition, pp. 241-260.* Prentice Hall, ISBN 0-13-066257-7, 2003

[5] Connell, J.H., *Diversity in tropical rain forests and coral reefs.* Science 199: 1302-1310, 1978

[6] Ford R.A., Wagner M., and Michalske J., *Gatekeeper II: New Approaches to Generic Virus Prevention.* From the proceedings of the International Virus Bulletin Conference, Chicago, 2004

[7] Forrest S., Hofmeyr S., Somayaji A., and Longstaff T., A Sense of Self for Unix Processes, 1996 IEEE Symposium on Security and Privacy, May 06-08, Oakland, CA

[8] Griffin T., and Zhuoqing M., *Internet Routing Streams*, Workshop on the Management and Processing of Data Streams, 2003

[9] Kephart J., and White S., *Measuring and Modeling Computer Virus Prevalence*, Proceedings of the 1993 IEEE Computer Society Symposium on Research in Security and Privacy, pp.2-14

[10] Lau F., Rubin S.H., Smith M.H., and Trajovic L. *Distributed denial of service attacks.* In IEEE International Conference on Systems, Man, and Cybernetics, pages 2275-2280, Nashville, TN, USA, 2000

[11] Liljenstam M., and Nocol D.M., *Comparing Passive and Active Worm Defenses.* In the Proceedings of the First International Conference on the Quantitative Evaluation of Systems, QEST 2004

[12] Moore D., Shanning C., and Claffy K., *Code-Red: a case study on the spread and victims of an Internet worm.* In Proceedings of the 2nd Internet Measurement Workshop, pages 273-284, 2002

[13] Staniford S., Paxson V., and Weaver N., *How to 0wn the Internet in Your Spare Time*, Proc. USENIX Security Symposium, 2002

[14] Toyoizumi H, and Kara A., *Predators: good will mobile codes combat against computer viruses*, Proceedings of the 2002 workshop on New security paradigms, September 23-26, Virginia Beach, Virginia

[15] Welch D.J., Buchheit N., Ruocco A., *Strike Back: Offensive Action in Information Warfare*, Proceedings of the 1999 workshop on New security paradigms, 1999