This product is part of the RAND Corporation monograph series. RAND monographs present major research findings that address the challenges facing the public and private sectors. All RAND monographs undergo rigorous peer review to ensure high standards for research quality and objectivity.

CYBERDETERRENCE AND CYBERWAR

MARTIN C. LIBICKI

Prepared for the United States Air Force

PROJECT AIR FORCE

The research described in this report was sponsored by the United States Air Force under Contract FA7014-06-C-0001. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

Library of Congress Cataloging-in-Publication Data

Libicki, Martin C.
Cyberdeterrence and cyberwar / Martin C. Libicki.
p. cm.
Includes bibliographical references.
ISBN 978-0-8330-4734-2 (pbk. : alk. paper)
1. Information warfare—United States. 2. Cyberterrorism—United States—Prevention. 3. Cyberspace—Security measures. 4. Computer networks—Security measures—United States. 5. Civil defense—United States. I. Title.

U163.L539 2009
355.3'43—dc22

2009030055

Cover design by Carol Earnest. Associated Press photo with overlay.

Published 2009 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665
RAND URL: http://www.rand.org
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Summary

The establishment of the 24th Air Force and U.S. Cyber Command marks the ascent of cyberspace as a military domain. As such, it joins the historic domains of land, sea, air, and space. All this might lead to a belief that the historic constructs of war—force, offense, defense, deterrence—can be applied to cyberspace with little modification. Not so. Instead, cyberspace must be understood in its own terms, and policy decisions being made for these and other new commands must reflect such understanding. Attempts to transfer policy constructs from other forms of warfare will not only fail but also hinder policy and planning.

What follows focuses on the policy dimensions of cyberwar: what it means, what it entails, and whether threats can deter it or defense can mitigate its effects. The Air Force must consider these issues as it creates new capabilities.

Cyberattacks Are Possible Only Because Systems Have Flaws

As long as nations rely on computer networks as a foundation for military and economic power and as long as such computer networks are accessible to the outside, they are at risk. Hackers can steal information, issue phony commands to information systems to cause them to malfunction, and inject phony information to lead men and machines to reach false conclusions and make bad (or no) decisions.

Yet system vulnerabilities do not result from immutable physical laws. They occur because of a gap between theory and practice. In theory, a system should do only what its designers and operators want it to. In practice, it does exactly what its code (and settings) tells it to. The difference exists because systems are complex and growing more so.

In all this lies a saving grace. Errors can be corrected, especially if cyberattacks expose vulnerabilities that need attention. The degree to which and the terms by which computer networks can be accessed from the outside (where almost all adversaries are) can also be specified. There is, in the end, no forced entry in cyberspace. Whoever gets in enters through pathways produced by the system itself.[1] It is only a modest exaggeration to say that organizations are vulnerable to cyberattack only to the extent they want to be. In no other domain of warfare can such a statement be made.

Operational Cyberwar Has an Important Niche Role, but Only That

For operational cyberwar—acting against military targets during a war—to work, its targets have to be accessible and have vulnerabilities. These vulnerabilities have to be exploited in ways the attacker finds useful. It also helps if effects can be monitored.

Certainty in predicting the effects of cyberattacks is undermined by the same complexity that makes cyberattacks possible in the first place. Investigation may reveal that a particular system has a particular vulnerability. Predicting what an attack can do requires knowing how the system and its operators will respond to signs of dysfunction and knowing the behavior of processes and systems associated with the system being attacked. Even then, cyberwar operations neither directly harm individuals nor destroy equipment (albeit with some exceptions). At best, these operations can confuse and frustrate operators of mili-

[1] Distributed denial-of-service attacks, the sort perpetrated against Estonia in 2007, are a partial exception. They clog the entryways to the system, rather than get into it. However, such attacks are, at worst, a minor nuisance to organizations (e.g., the military, electric power producers) that can run without interacting with the public at large.

tary systems, and then only temporarily. Thus, cyberwar can only be a support function for other elements of warfare, for instance, in disarming the enemy.

The salient characteristics of cyberattacks—temporary effects and the way attacks impel countermeasures—suggest that they be used sparingly and precisely. They are better suited to one-shot strikes (e.g., to silence a surface-to-air missile system and allow aircraft to destroy a nuclear facility under construction) than to long campaigns (e.g., to put constant pressure on a nation's capital). Attempting a cyberattack in the hopes that success will facilitate a combat operation may be prudent; betting the operation's success on a particular set of results may not be.

Strategic Cyberwar Is Unlikely to Be Decisive

No one knows how destructive any one strategic cyberwar attack would be. Estimates of the damage from *today's* cyberattacks within the United States range from hundreds of billions of dollars to just a few billion dollars per year.

The higher dollar figures suggest that cyberattacks on enemy civilian infrastructures—strategic cyberwar—may be rationalized as a way to assist military efforts or as a way to coerce the other side to yield to prevent further suffering. But can strategic cyberwar induce political compliance the way, say, strategic airpower would? Airpower tends to succeed when societies are convinced that matters will only get worse. With cyberattacks, the opposite is more likely. As systems are attacked, vulnerabilities are revealed and repaired or routed around. As systems become more hardened, societies become less vulnerable and are likely to become more, rather than less, resistant to further coercion.

Those who would attempt strategic cyberwar also have to worry about escalation to violence, even strategic violence. War termination is also not trivial: With attribution so difficult and with capable third parties abounding (see below), will it be clear when one side has stopped attacking another?

Cyberdeterrence May Not Work as Well as Nuclear Deterrence

The ambiguities of cyberdeterrence contrast starkly with the clarities of nuclear deterrence. In the Cold War nuclear realm, attribution of attack was not a problem; the prospect of battle damage was clear; the 1,000th bomb could be as powerful as the first; counterforce was possible; there were no third parties to worry about; private firms were not expected to defend themselves; any hostile nuclear use crossed an acknowledged threshold; no higher levels of war existed; and both sides always had a lot to lose. Although the threat of retaliation may dissuade cyberattackers, the difficulties and risks suggest the perils of making threats to respond, at least in kind. Indeed, an explicit deterrence posture that encounters a cyberattack with obvious effect but nonobvious source creates a painful dilemma: respond and maybe get it wrong, or refrain and see other deterrence postures lose credibility.

The case for cyberdeterrence generally rests on the assumption that cyberattacks are cheap and that cyberdefense is expensive. If cyberattacks can be conducted with impunity, the attacker has little reason to stop. Besides, nuclear deterrence prevented the outbreak of nuclear conflict during the Cold War. What is there about cyberspace that would prevent a similar posture from working similarly well? Plenty, as it turns out. Questions that simply do not crop up with nuclear or even conventional deterrence matter in cyberspace whenever the target of an attack (the "we") contemplates retaliation.

Will we know who did it? Cyberattacks can be launched from literally anywhere, including cybercafés, open Wi-Fi nodes, and suborned third-party computers. They do not require expensive or rare machinery. They leave next to no unique physical trace. Thus, attribution is often guesswork. True, ironclad attribution is not necessary for deterrence as long as attackers can be persuaded that their actions may provoke retaliation. Yet some proof may be necessary given (1) that the attacker may believe it can shake the retaliator's belief that it got attribution right by doing nothing different ("who, me?") in response to retaliation, (2) that mistaken attribution makes new enemies, and (3) that neutral observers may need to be convinced that retaliation is not aggression.

Can retaliators hold assets at risk? It is possible to understand the target's architecture and test attack software in vivo and still not know how the target will respond under attack. Systems vary by the microsecond. Undiscovered system processes may detect and override errant operations or alert human operators. How long a system malfunctions (and thus how costly the attack is) will depend on how well its administrators understand what went wrong and can respond to the problem. Furthermore, there is no guarantee that attackers in cyberspace will have assets that can be put at risk through cyberspace.

Can they do so repeatedly? It is difficult to imagine an act of cyberretaliation that is prospectively so awesome that no potential attacker would run the risk of being hit (a crucial feature of nuclear retaliation). Repeated application may be necessary but is not necessarily possible. Even successful retaliation may not be convincing if the attacker tells itself it will be less vulnerable the next time around.

Can cyberattacks disarm cyberattackers? In a world of cheap computing, ubiquitous networking, and hackers who could be anywhere, the answer is no.

Will third parties stay out of the way? Cyberattack tools are widely available. If nonstate actors jump into such confrontations, they could complicate attribution or determining whether retaliation made the original attackers back off.

Might retaliation send the wrong message? Most of the critical U.S. infrastructure is private. An explicit deterrence policy may frame cyberattacks as acts of war, which would indemnify infrastructure owners from third-party liability, thereby reducing their incentive to invest in cybersecurity.

Can states set thresholds for response? Unless a state declares that all cyberattacks, no matter how minor, merit retaliation, it needs to define an actionable threshold. Proving that any one attack crossed it, however, may be tricky.

Can escalation be avoided? Even if retaliation is in kind, counterretaliation may not be. A fight that begins in cyberspace may spill over into the real world with grievous consequences.

Responses to Cyberattack Must Weigh Many Factors

In many ways, cyberwar is the manipulation of ambiguity. Not only do successful cyberattacks threaten the credibility of untouched systems (who knows that they have not been corrupted?) but the entire enterprise is beset with ambiguities. Questions arise in cyberwar that have few counterparts in other media.

What was the attacker trying to achieve? Because cyberwar can rarely break things much less take things, the more-obvious motives of war do not apply. If the attacker means to coerce but keep its identity hidden, will the message be clear? If the attack was meant to disarm its target but does so only temporarily, what did the attacker want to accomplish in the interim? Can an attack and its aftermath be used as part of a competitive strategy in commercial or political markets? What role might a cyberattack play in the attacker's master narrative?

What should the target reveal about the attack? Many attacks—corruption attacks, disruption attacks on systems deep within an organization—have effects that are not generally obvious. Revealing what happened is more honest and necessary to justify public retaliation. However, silence might mitigate panic, preserve confidence in systems as they are being fixed, and support nonconfrontational strategies (e.g., private exposure followed by diplomatic threats) or nonpublic retaliatory strategies. Whether and when to name the attacker also deserve thought. Premature revelation can be embarrassing, but if revelation comes long after the attack, the link between retaliation and the original attack may lose credibility. Revelation too far in advance of retaliation gives the attacker time to ward off a retaliatory attack through better defenses, counterthreats, or mobilizing opinion to its side.

How should states respond to freelance attacks? Establishing that a state is protecting the attackers creates yet another hurdle to attribution, but what if the hackers were just not sought with sufficient vigor? It is hard to know whether retaliating against such a state would energize its prosecutorial energies—or backfire.

Should deterrence be extended to allies? Figuring out who actually hit the ally's system and with what effect requires poking into their systems, something they may balk at ("don't you trust us?"). Allies may also have ulterior motives for fingering one particular attacker.

Military Cyberdefense Is Like but Not Equal to Civilian Cyberdefense

Because military networks mostly use the same hardware and software as civilian networks, they have mostly the same vulnerabilities. Their defense resembles nothing so much as the defense of civilian networks—a well-practiced art. But military networks have unique features: real enemies, specific cyberthreats, and many closed systems.

The primary goal for the military is to function as well under cyberattack as it does on a day-to-day basis—after all, performance under military attack is how militaries are measured. Robustness is key, but it goes beyond network security engineering to encompass all measures that permit the broader system (the military itself) to work when its subsystems do not. The military must pay more attention than others do to the failure modes that are likely to be the most damaging or most prevalent.

Because the effects of cyberattack are temporary, the military's first priority in the wake of a major successful attack is to figure out whether a physical attack is coming to take advantage of the systems being crippled. The second priority, assuming the attacker is monitoring such systems before deciding whether to attack, is to make it look as though little damage has been done. The third is to achieve recovery. Everything else (including retaliation) follows later.

Implications for the Air Force

The United States and, by extension, the U.S. Air Force, should not make strategic cyberwar a priority investment area. Strategic cyberwar, by itself, would annoy but not disarm an adversary. Any adversary that merits a strategic cyberwar campaign to be subdued also likely possesses the capability to strike back in ways that may be more than annoying.

Similar caution is necessary when contemplating cyberdeterrence. Attribution, predictable response, the ability to continue attack, and the lack of a counterforce option are all significant barriers. The United States may want to exhaust other approaches first: diplomatic, economic, and prosecutorial.

Operational cyberwar has the potential to contribute to warfare—how much is unknown and, to a large extent, unknowable. Because a devastating cyberattack may facilitate or amplify physical operations and because an operational cyberwar capability is relatively inexpensive, it is worth developing. That noted, success at cyberwar is not only a matter of technique but also requires understanding the adversary's networks in the technical sense and, even more, in the operational sense (how potential adversaries use information to wage war). The Air Force should also recognize that the best cyberattacks have a limited shelf life and should be used sparingly.

Throughout all this, cyberdefense remains the Air Force's most important activity within cyberspace. Although most of what it takes to defend a military network can be learned from what it takes to defend a civilian network, the former differ from the latter in important ways. Thus, the Air Force must think hard as it crafts its cyberdefense goals, architectures, policies, strategies, and operations.