1

Introduction

Despite its roots in the U.S. Department of Defense (DoD), the global Internet has primarily, although not exclusively, been an avenue and arena of peaceful commerce. With every year, an increasing percentage of the world's economy has migrated from physical media, or older electronic media such as telephones and telegraphs, to the public Internet and to private or semipublic internets. Systems that were once inaccessible to persons off-premises, such as power plant controls, are now theoretically accessible to anyone around the world. Other hitherto self-contained networks, such as those that transferred money, are now commingled with the larger, more public networks such as the Internet or the international phone system.

Indeed, its very success is what has turned the Internet into a potential venue of warfare. It is not only that defense systems of advanced militaries are being knit into more powerful systems of systems – thereby becoming the militaries' new center of gravity. The real impetus is that the more cyberspace is critical to a nation's economy and defense, the more attractive to enemies is the prospect of crippling either or both via attacks on or through it. Hackers can and do attack information systems through cyberspace. They can attack the cyberspace itself through operations against the networks that provide the basis for this new medium. Defenders thus must keep these hackers out of their systems. If hackers get in, they could wreak great damage. At a minimum they might steal information. Worse, they can make systems go haywire. Worst, they could inject phony information into systems to distort what users think they absorb when they deal with systems. Hackers might take over any machine (such as a pump) controlled by a networked computer system and use it according to their ends and not those of its owners.

None of this requires mass, just guile. For that reason, attacks in cyberspace do not need the same government backing as attacks in older media do. Any group, or even individual, can play – even, perhaps especially, terrorists. Prior to 9/11, in fact, it was difficult to conceive of a strategic attack on the U.S. homeland by nonstate actors *except* through the medium of cyberspace. Such would be a bloodless attack from afar that left no traces but could cause the systems we rely on to crash mysteriously. The President's Commission on Critical Infrastructure Protection argued in 1996 that the capability to launch such an attack did not yet exist – but given five years (that is, by 2001), it very well might.

Perhaps needless to add, although advanced nations have more at stake in cyberspace than developing nations do, the latter are increasingly being drawn into its domain. Thus, they too are vulnerable to attacks from what are, in general, the larger and more sophisticated cohorts of hackers from the first world.

By such means, cyberspace has joined air and outer space as a new medium of conflict.[1] Granted, evidence that it has become a *significant* medium of conflict is sparse. This may be because the last three wars in which cyberspace could have played a role – Kosovo, Afghanistan, and Iraq, respectively – were against countries with minimal presence in cyberspace. They had little that the United States could attack, or at least attack more efficiently than conventional means already permitted it to do. So far, other countries have lacked the sophistication and will to do much damage to the U.S. use of cyberspace. But since participation

[1] The 2001 Department of Defense Quadrennial Defense Review Report listed four "Key Military–Technical Trends." The third was "Emergence of new arenas of military competition":

> Technological advances create the potential that competitions will develop in space and cyber space. Space and information operations have become the backbone of networked, highly distributed commercial civilian and military capabilities. This opens up the possibility that space control – the exploitation of space and the denial of the use of space to adversaries – will become a key objective in future military competition. Similarly, states will likely develop offensive information operations and be compelled to devote resources to protecting critical information infrastructure from disruption, either physically or through cyber space (p. 7).

in and dependence on cyberspace is growing, the odds of consequential conflict, and thus hostile conquest, must certainly be rising.

Lost in this clamor about the threat from hackers is another route to conquest in cyberspace, not through disruption and destruction but through seduction leading to asymmetric dependence. The seducer, for instance, could have an information system attractive enough to entice other individuals or institutions to interact with it by, for instance, exchanging information or being granted access. This exchange would be considered valuable; the value would be worth keeping. Over time, one side, typically the dominant system owner, would enjoy more discretion and influence over the relationship, with the other side becoming increasingly dependent. Sometimes the victim has cause to regret entering the relationship; sometimes all the victim regrets is not receiving its fair share of the joint benefits. But if the "friendly" conquest is successful, the conqueror is clearly even better off.

The central contention of this work is that the possibilities of hostile conquest may be less consequential than meets the eye while the possibilities of friendly conquest ought to be better appreciated. The current obsession with hostile conquest fosters a tilt toward closed systems, at least among those who have powerful systems to begin with. Those with the most attractive systems – in terms of information, knowledge, services, and reach – have an inherent advantage whose benefits they might deny themselves by concentrating on the threat to themselves. This is particularly so for the national and homeland security community (including law enforcement, homeland defense, and infrastructure). By taking a more open approach to cyberspace, they may extend their influence and the influence of their values more certainly than they would by taking a closed approach.

In a sense, this argument echoes the distinction made by Joseph Nye between a nation's hard power and its soft power.[2] Hard power is embodied in military force, soft power in its culture. Hard power, like hostile conquest in cyberspace, ultimately entails one nation doing to another what the other would prefer it not do. It is involuntary. Soft power, like

[2] Joseph Nye, *Bound to Lead: The Changing Nature of American Power*, New York (Basic Books), 1990.

friendly conquest in cyberspace, describes the process of enticement. It is voluntary, at least at first. In the case of soft power, the elites of the affected country may find themselves unable to roll back the tide of imported cultural and economic mores without facing resistance and revolt. But rarely can one nation control or manipulate the instruments of soft power to create such a dependency; more often, it works independent of national strategy. With friendly conquest in cyberspace, however, the seducer retains part of the leverage precisely because the controls over the seductive system are not relinquished.

Hence the choices, many of them public choices. Hence, too, the orientation of this work, one to be understood in its policy and management rather than technical context. It is aimed at educated individuals who are interested in public policy. Admittedly, issues of cyberspace can become quite technical, and so the text tries to clarify some key concepts. Cyberspace issues are not unique in that regard. It can be hard to understand, say, the pros and cons of strategic ballistic missile defense without some understanding of physics. Nevertheless, arguments about strategic defense are not entirely technical ones. Similarly, arguments about the proper use and exploitation of cyberspace are not entirely technical. Readers who happen to be information security experts may appreciate reading this or that point of view; they are unlikely to add much to their technical knowledge of their craft by reading this.

1.1 What Does Conquest Mean in Cyberspace?

This work is entitled not "The Conquest *of* Cyberspace" but "Conquest *in* Cyberspace" for a reason. To emphasize the "of" is to suggest that there is, in fact, *a* cyberspace that exists in the same sense that the oceans do. It has distinct parameters and perimeters, and one can define conquest within this space. This leaves the only interesting question one of determining who has, in fact, taken possession of what part of cyberspace and how they accomplished such feats. Emphasizing "in," by contrast, reflects the fact that while something akin to conquest can be defined for cyberspace, cyberspace itself cannot be conquered in any conventional sense.

To understand why, it helps to understand what cyberspace itself means. Ironically, that process is best begun by discussing what cyberspace does *not* mean – or at least does not mean yet.

The term "cyberspace" was coined in William Gibson's 1984 classic, *Neuromancer.* The concept was further described in compelling detail in Neil Stephenson's 1989 *Snow Crash.* Both portrayed it as an alternative universe that people could participate in ("jack into" pace Gibson). It may be seen, particularly in some movies, as being just on the other side of the twenty-first century's version of Alice's looking-glass. Cyberspace, so defined, may be evoked through a text-only medium such as a chat room, but it can also be evoked more tangibly by a virtual reality simulation in which what one sees, hears, and, to some extent, feels is all synthesized on the spot. Computer power and fat networks make this illusion easier to generate with every passing year.

This often attractive concept should not lead one to imagine cyberspace as being *the* parallel universe – as if a mapping of this reality into another dimension. Four tenets suggest why cyberspace should be understood on its own merits.

First, cyberspace is a replicable construct. Being replicable, it exists in multiple locations at once. Because it is replicable, it is also reparable.

By contrast, only confusion can follow the unconscious assumption that there is *one* cyberspace in the sense that there is, say, one outer space. The existence of a single something called outer space derives from the simple fact that there is a planet earth and that every point on or above the planet has a unique location relative to it. This uniqueness is firmly rooted in physical law. The planet, for instance, has only one geosynchronous belt, and locations[3] in it are carefully allocated for every satellite (of a given broadcast frequency). There is also one spectrum, uses of which are governed by international conventions such as the World Radio Conference. From a military perspective, one nation's fleets of hunter-killer satellites can keep another country from establishing its own constellation. Control in space, can, in theory, be exclusive.

Cyberspace, by contrast, is built, not born. Every system and every network can hold its own cyberspace – indeed, it can hold a limitless number of quasi-independent spaces. Cyberspace can appear in multiple,

[3] Satellites in geosynchronous orbit appear to linger above a single point on the equator. Satellites in such orbits have to be separated from each other by a certain arc length if they broadcast in the same frequencies. As such, there are a finite number of such orbits and each is assigned on a global basis.

almost infinite, manifestations and forms. Even shared spaces can be indefinitely replicated. This is apparent, for instance, in multiplayer games (such as the *Sims Online* or *Rise of Nations*). Since the number of players in any one game is small, there have to be many of them to accommodate everyone who wants to play.

Not only is cyberspace a construct, but the rules of cyberspace are largely constructs – there is little hard-and-fast physics of the sort that dictates what can and cannot be done in, say, outer space. What can and cannot be done in cyberspace need not follow the laws of physics or the laws of man – although violating the latter may have real-world repercussions. There is no inherent "there" there except as mutually accepted.[4] Even larger games (massive multiplayer online role-playing games [MMOPRGs]) such as *EverQuest* or the popular South Korean multiplayer game *Lineage*, although more unified, exist because they have been constructed for that purpose, often a commercial one. Admittedly, some people are so hooked by these games that they actually pay real money to acquire virtual goods, useful only online.[5] Yet, they are not inherently fixed; should something more attractive come along, they could be easily abandoned. Not so for outer space or the oceans; they will always be there.

Second, to exist in cyberspace, your interactions must be recognized there. To show why, consider a distributed interactive simulation of the sort used in military training. If such a simulation is to work at all, there must be a synthetic universe into which all player attributes and actions are mapped. Supposedly, all players could factor in everyone's moves and initial attributes in their own unique way (for example, what you see as driving, others see as flying), but inevitably the result would resemble nothing so much as the argument of children: "you're dead," "no, I'm not," "yes, you are," "no, I'm not." So there has to be a master set of rules for any given space. Messages (that is, byte-strings) that do not accord to the rules are invariably rejected as meaningless, regardless of

[4] "Mutually accepted" is not meant to imply commonly understood. People may think the game has certain rules when, through simple misunderstanding or subterfuge, it turns out to have quite different ones.

[5] See Julian Dibbell, "The Unreal Estate Boom," *Wired*, January 2003, 11, 1, pp. 106–13.

how earnestly or maliciously put forward. The intrusion of a third party must be reflected in a change in the game's state; again, whether calculated centrally and broadcast or, instead, calculated individually in an identical way is secondary. Unrecognized actions or the actions of unrecognized parties do little harm except for perhaps clogging the lines. And, of course, not every player need be human; they can be machines.

Third, some aspects of cyberspace nevertheless tend to be persistent. A few rules of cyberspace, such as the laws of cryptography, derive from mathematics. Others are artifacts of well-accepted conventions (such as TCP/IP) or reflect the dominance of certain products in the marketplace (such as Microsoft Windows). One can construct a cyberspace without them, but most information systems adhere just the same. Such rules can come from many places, such as from those who write the software or from the community that maintains the environment in which the software runs.[6] So, while these rules remain constructs, they are constructs in which people have invested value.

Certain systems, as well, are persistent. There is, for instance, only one Internet, and it has certain conventions such as a hierarchy of routers as well as a set of recognized names corresponding to a set of recognized addresses.[7] But there are also internets (small "i") that use the same ubiquitous and richly supported communications protocols as the Internet but are not connected to the Internet. There are yet others, which are connected but in ways that make it very difficult for the innocent public or not-so-innocent hackers to get into them.

Even at a macro level, as Lawrence Lessig[8] has argued, cyberspace has nearly no inherent properties and only a few strong tendencies; everything else is imposed by those with the power to do so. The oft-cited aphorism that the Internet interprets censorship as network damage and routes around it has been used to imply that the inherent qualities of the Internet

[6] It would be harder to change unilaterally games such as Dungeons and Dragons, whose rules have evolved organically over time.

[7] To illustrate that even constructs have value that can be captured and traded, note the large amounts of money associated with certain domain names that serve as beacons in a fog of potential URLs. Nevertheless, the tendency to type "socks.com" in order to begin shopping for socks is being replaced by that of typing "socks" into Google.

[8] Lawrence Lessig, *Code, and Other Laws of Cyberspace*, New York (Basic Books), 1999.

have made free speech inevitable in that medium. But this runs up against the real-world constraints that governments such as China have largely imposed successfully on its Web users. There is, Lessig goes on to argue, a substantial capability to express social norms, hitherto reified only in legal code, as computer code to achieve roughly the same ends. To say that cyberspace is a "commons" or a "market" presupposes some expression in cyberspace of social norms and, in some cases, legal enforcement that permits commons or markets to function in real space. It does not arise from the nature of cyberspace or always come out in the same way. For instance, EBay, the online auction site, is a global market, and there are mechanisms (for example, a global reputation registry) that work on EBay to provide other or more efficient ways of enforcing commercial norms that exist in the physical world. But supposing such markets would work absent any mechanisms and social norms whatsoever is naive.

Fourth, cyberspace has separate layers, the conquest of each of which has vastly different meaning. Stated briefly, and discussed in much greater detail in Chapter 10, one can define three layers in cyberspace with their parallels in linguistics: the physical, the syntactic, and the semantic.[9]

The physical layer – including such things as wires, routers, and switches – is the foundation of cyberspace in the tangible world.[10] Conquest that takes place here could be understood in terms of physical control over the infrastructure – frustrated only by the ease with which most of the infrastructure can be replicated if necessary.

The syntactic layer reflects both the format of information in cyberspace and how the various information systems from which cyberspace is built are instructed and controlled.[11] As explained further

[9] Chapter 10 also discusses a fourth layer, pragmatics – essentially the intentions that lie behind the speech acts. Until such time – and it may be coming – that one can usefully impute intentions (or goals) to programs and machines, the pragmatic layer applies only to person-to-person interaction mediated through cyberspace. Thus "conquest" at this layer is very hard to define.

[10] It is not impossible to build a functioning cyberspace atop a biological stratum and use it to convey analog and/or fuzzy information (as today's nervous system does). All the software needs to know are the system's abstracted basic parameters (for example, how fast, how reliable, how ubiquitous, how much capacity).

[11] One of the great engineering successes of the TCP/IP protocol (and the Internet conventions that rest on them) has been to push the intelligence into the periphery of the system rather than concentrating it in the control infrastructure. J. H.

later in this chapter, the syntactic layer can itself be divided into sublayers; one canonical model, Open System Interconnection (OSI), identifies seven of them. Control here is often a matter of mastery: Can my knowledge of the rules overcome your knowledge to get machines to do what I, rather than you, want? And who writes the rules?

The semantic layer contains the information meaningful to humans or connected devices (for example, machine tools). Here the issue is one of influence: can I present to you a different version of reality that others take to be true?

So, conquest works differently at different layers. Physical access (that is, connectivity) does not mean syntactic access. Syntactic access does not mean meaningful semantic access. And semantic access does not necessarily result in meaningful change in what people believe about the world or even about cyberspace.

The layers of cyberspace may be likened to a party hall with private rooms. All these rooms are part of the same physical structure and they are mutually accessible, but that does not mean that what goes on in one room says much about what goes on in the next. To get into any one room may require a key (in cyberspace terms, knowledge of the network address and the password). Those who make their way in still have no guarantee of meaningful interchange with any of the participants. One may be simply ignored or not understood. Becoming a meaningful part of the conversation has three requirements: getting to the party hall, finding a key to the private room, and being accepted by those who are conversing. Some party rooms are better than others, in part because of better physical facilities (in cyberspace terms, faster connections, better data stores, more sophisticated support services, and so on). Some conclaves such as chat rooms are open to everyone.[12] Others, such as The Well (a Sausalito-based

Saltzer, D. P. Reed, and D. D. Clark, "End-to-End Arguments in System Design" *ACM Transactions in Computer Systems,* November 1984, 2, 4, pp. 277–88, also available at web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf; see also David Isenberg, "The Rise of the Stupid Network" (www.isen.org). This is because the packets that carry the information payload also carry, embedded within them, processing instructions to the network. This instructions/content relationship has analogies to the syntactic/semantic relationship of human language.

[12] In practice, spaces such as America Online's (AOL) chat rooms are open only to AOL members and can exclude known abusers.

bulletin-board system [BBS] that predated the Worldwide Web site), are by invitation only. The more desirable neighborhoods in cyberspace are often so because they are better organized with more entertaining activities and intriguing conversationalists; others are more interesting because they permit certain types of business to be done.

Even, perhaps especially, Islamic (more technically, Salafist-Jihadist) terrorists hang out in their own neighborhoods to transact their "business." In some cases, notably when propagandizing for the masses, seeking recruits, or distributing Web materials, these neighborhoods tend to be public. In other cases, when mooting plots among themselves, Jihadist sites are more private; access is carefully revealed to known individuals. These are not rigid or even rigidly enforced distinctions. Sites have been penetrated by researchers who have figured out how to sound like a potential terrorist.[13]

It turns out that hostile conquest in cyberspace takes place largely at the physical and syntactic layers, while friendly conquest in cyberspace, because it has to do more with the exchange and encoding of knowledge, tends to take place at the syntactic and semantic layers.

1.2 Précis

The contention – that it is hard to control the world by hostile conquest in cyberspace but that the power available through friendly conquest merits attention – is developed in three parts followed by a conclusion.

Part I deals with hostile conquest in cyberspace. It starts with the premise that information systems exist to generate information, that information is used for decisions, but that humans are the agents of knowledge and decision making.[14] In other words, if one is to attack systems and so affect decisions, one must recognize the decisions are ultimately made by natural rather than artificial cognition. People, unlike

[13] Notably the Search for International Terrorist Entities' (SITE) Rita Katz; see Benjamin Wallace-Wells, "Private Jihad: The Woman Who Became a Freelance Spy," *New Yorker*, May 29, 2006, pp. 28–41.

[14] While acknowledging that systems are also used to make decisions and carry out actions automatically, the fundamental choice of where and how to put people into the decision loop is quintessentially a human one as well.