Towards an Ethical Code for Information Security?

Panel Chair/Editor: Steven J. Greenwald

Steven J. Greenwald
Information Security Advisor
2521 NE 135th Street
North Miami, Florida 33181
U.S.A.
sjg6@gate.net

Brian D. Snow
Independent Security Advisor
Clarksville, Maryland 21029
U.S.A.
briansnow@comcast.net

Richard Ford
Research Professor,
Computer Sciences
Florida Institute of Technology
Melbourne, Florida 32901
U.S.A.
rford@cs.fit.edu

Richard Thieme
ThiemeWorks
PO Box 170737
Milwaukee, Wisconsin 53217
U.S.A.
rthieme@thiemeworks.com

ABSTRACT

Most computer scientists reflexively reject the idea of a malicious universe due to its conflict with the dominant scientific paradigm of a non-teleological impartially disinterested universe. While computer scientists might not view the universe as benign, neither do they view the universe as actively hostile. In addition, most scientist would take the view that a teleological universe equals paradigmatic heresy.

To this we say: *feh!*

Outsiders call us "paranoid," but any sensible member of our field knows for a fact that the information security universe *does* act maliciously.

Our universe really *does* try to cause us harm.

Two of us (Snow and Greenwald) have recently given some thought to ethical notions in a somewhat related field; we realized that due to the paradigm-conflicting presence of a malicious universe, we may need a specialized code of ethics for the computer security field. We therefore assembled a group of experts with different viewpoints on this subject for a New Security Paradigms Workshop panel. We felt that NSPW would provide the perfect venue to discuss this radical concept.

We gave the panel the charge of considering the mere *notion* of a specialized code of ethics for the field of cybersecurity. Do we really need or want a specialized code of ethics? We therefore had no interest, at least for the purposes of this panel, with the possible *contents* of such a specialized ethical code.

Our panelist positions run the gamut from "We desperately need a strong code of ethics" to "A specialized ethical code would cause great harm." Along with our positions, we report on the feedback we received from the NSPW process and what we learned.

NSPW'08, September 22–25, 2008, Lake Tahoe, California, USA.

Categories and Subject Descriptors

A.0 [**GENERAL**]: Conference proceedings; J.7 [**COMPUTERS IN OTHER SYSTEMS**]: Command and control; K.4.1 [**Public Policy Issues**]: Abuse and crime involving computers; K.4.1 [**Public Policy Issues**]: Ethics; K.4.1 [**Public Policy Issues**]: Privacy; K.4.1 [**Public Policy Issues**]: Regulation; K.4.1 [**Public Policy Issues**]: Use/abuse of power; K.4.4 [**Electronic Commerce**]: Security; K.5.2 [**Governmental Issues**]: Regulation; K.6.5 [**Security and Protection**]; K.7.4 [**Professional Ethics**]: Codes of ethics; K.7.4 [**Professional Ethics**]: Codes of good practice; K.7.4 [**Professional Ethics**]: Ethical dilemmas

General Terms

ECONOMICS,LEGAL ASPECTS,MANAGEMENT,SECURITY

Keywords

Code of ethics, command and control, computer security, cybersecurity, ethics, ethical code, information security, Kuhn, New Security Paradigms Workshop, NSPW, paradigm, professional ethics, regulation

1. INTRODUCTION

With the ever-increasing ethical challenges to computer security professionals[1] in the information security[2] field, we felt the need to discuss the *notion* of an ethical code specifically designed for our field and that also changes the dominant scientific ethical paradigm. Yes, the universe really wants to "get us" — an idea that most mainstream computer scientists would almost reflexively reject due their use of a different Kuhnian paradigm [15]. Of course, we use the term *universe* to mean the "universe of discourse" that we all use in our paradigms.

[1]By "professionals" we limit our scope to meaning practitioners, designers, and in appropriate cases sponsors in the field of computer security, and unless specifically noted otherwise, we do not intend it, or this paper, to apply to general computer professionals.
[2]We synonymously use the terms "computer security," "information security," and "cybersecurity."

Information security ethics involves a paradigm shift away from the standard ethical paradigm of computer science as found in our canonical professional societies like ACM [1] and IEEE [13]. The dominant paradigm in computer science ethics does not assume a *malicious* universe; instead it assumes the standard scientific paradigm notion of an *impartially disinterested* universe; *nota bene* that the idea of teleology in most scientific paradigms comes as close as science ever does to heresy (the sign of a Kuhnian paradigm shift of course). But in the information security field we certainly must assume a malicious universe not only due to the adversarial nature of the people involved, but also due to the existence of proxies (malware) that run on the behalf of human adversaries (or even taking on a "life of their own").

We therefore challenge the *status quo ante* idea that scientists should assume a disinterested and impartial universe (or even worse, a benign one). In addition, we note the lack of any *valid* ethical code specifically designed for information security. We can certainly understand why we have this lack: given the disinterested universe paradigm it makes no sense at all to have a separate code of ethics for information security that works separately from the standard computer science ethical codes.

1.1 The Panel Charge

We use the term "ethics" in this paper (as distinct from the term "moral") as defined by the third definition given in the *American Heritage Dictionary of the English Language, Fourth Edition* (2000): "The rules or standards governing the conduct of a person or the members of a profession: *medical ethics* [emphasis in original]."

All *bona fide* members of the computer security field have a horrifying awareness of the ethical dimensions of the field. We regularly encounter the results of the malicious actions of miscreants and adversaries and wonder about things such as the appropriate response to these actions. We also sometimes wonder about the behavior of some of our peers and sometimes question their ethical choices. Some of us must choose actions in ethically "gray" areas where even refusing to act becomes an ethical choice.

All of the authors agree that we need to clarify ethical ideas, and among ourselves we have discussed in detail many ethical issues in our field. We realize that, while meta-ethics has its intellectual appeal, if we wish to alter the behavior of people in the computer security field then we should *consider* having a specialized ethical code for our field. To take an example unique to our field and current at the time of writing this: should a research team divulge the hitherto publicly unknown details and techniques they used to exploit security vulnerabilities in a particular widely-used implantable medical device (IMD) [12]? If they do, then they might put a large number of patients at risk. If they don't, then they might fail to provide stimulation in the IMD field for the device manufacturers and other researchers to improve security. Any one of the information security experts invited to attend NSPW 2008 can think of literally hundreds (at least) of similar ethically challenging cases/situations.

We therefore feel that we must consider (at least) the following three fundamental questions if we wish to avoid a naïve approach.

1. Does information security actually need specialized ethics?
2. If we answer "yes" to (1), then what characteristics of specialized ethics do we want/need?
3. If we answer "yes" to (1) and have an answer for (2), then how do we implement the ethics that we want/need?

Our panel addressed these fundamental questions through the following positions of our four panelists.

1. We need ethics and can implement them via an ethical code.
2. We need ethics, but implementation issues may prevent the creation of an ethical code.
3. Computer security already has default ethics.
4. We do not need or even *want* ethics; even if we did we could not implement them.

These positions form several dimensions, so we will address these positions in the next four sections, each written from the individual perspective of each panelist and giving some of their individual flavor. We have the added a section containing some of the comments and interaction from our NSPW "shepherd." We received a lot of feedback from the workshop process and we report on that in the next section, and finally give our conclusions.

1.2 A Note on Cultural Bias

Part of the challenge with any discussion of ethics is that one tends to approach the subject with a strong cultural bias — in our case, a primarily occidental one. In terms of computer security, this can cause a problem as in many cases our allies and adversaries have very different ethical worldviews compared to our own. These cultural biases make any global security ethical framework extremely difficult. For example, different ethical attitudes toward privacy make any agreement on protection of user data difficult. If different ethical standards (and therefore in all likelihood laws) exist in Europe, for example, a profiteer may base his or her business in another jurisdiction that holds different values. However, due to the global nature of the Internet, this may still impact European users.

Fortunately, several factors help mitigate this problem. First, our mores regulate the way we believe we should act in a particular circumstance. Often, these do not get influenced by how others might handle the same situation. Second, no real disadvantage occurs with this discussion taking place across multiple cultures — correctly done, these issues should generate useful insight into the different ways our international colleagues make decisions. Finally, we have noted how the field of computer security has a basis on the concept of a hostile world; as such, our decisions do not assume the "right" course of action on the part of the world in general.

Thus, while we acknowledge that the following discussion reflects a Judeo-Christian (Western) worldview, we believe that the discussion still has relevance as it effects how we structure our own ethical stance. However, the reader should at all times maintain an awareness that we based much of the argument here on certain cultural assumptions, and as such, they might not have global applicability.

2. *BRIAN D. SNOW'S POSITION:* WE NEED ETHICS SO LET'S DEVELOP AN ETHICAL CODE

Do we need ethics for computer security practitioners? I think so. So have others. But I will not be insistent on the point. Ethics can only help mitigate some problems, not solve them.

Norbert Weiner said in 1948 ". . . we are already in a position to construct artificial machines of almost any degree of elaborateness of performance. Long before Nagasaki and the public awareness of the atomic bomb, it had occurred to me that we were here in the presence of another social potentiality of unheard-of importance for good and for evil." [20, pp. 27–28]. In 1950 Weiner published *The Human Use of Human Beings* [21], which, to paraphrase Terrell Ward Bynum [4] "established him as the founder of computer

ethics and built a computer ethics foundation which even today is a basis for analysis. It contains a method for doing applied ethics, discussions of the fundamental questions of computer ethics, and examples of key computer ethics topics."

Computer security practitioners tend to be (at least) engineers, computer scientists, or mathematicians, and each profession already has an established ethical code [13; 1; 2, respectively]. We also have [14] as a code of ethics for one set of security practitioners ((ISC)^2).

So some could argue no further work is needed. I beg to differ.

To protect us from known human frailties, an effective ethical code must be tailored to the specifics of a particular work environment and the pressures within it. They are meant to encourage practitioners to avoid "temptations" that are prevalent in their field that could lead to embarrassment, misfeasance, malfeasance, corruption, or more severe lapses. An Ethics Code should protect the weak from the self-interest of the strong, and its key points are typically shaped by the structure and workflow of the given field. For Security practitioners, the field is weak, and it is the enemy who is strong. Computer security is new enough and weak enough that it needs all of the help a carefully structured code of ethics could provide.

The security field has unique aspects not shared by other fields. The main difference is that they should practice to prevent *malice* from impacting their clients. Architects designing skyscrapers try to counter normal stress, or expected weather or other expected environmental factors. Once a building is up, the architect does not need to alter his design process for his next building in another city to cope with how weather will now behave; the weather does not change significantly based on his prior practice. Such is not the case with a security architect facing malice. The security professional faces an environment that adaptively (and rapidly) changes to nullify his efforts, and once again obtain an advantage against his design.

The early Arpanet designers innocently assumed a benign environment, and we have been living with the consequences ever since. William Cheswick of AT&T observed that current net security designs are "crunchy on the outside and chewy inside" as a continuing reflection on design practice even today. Designs really should be more resilient to malice — offering not only a strong surface when appropriate, but displaying resistance in the interior as well when the enemy progresses. This is a mental adjustment needed both in design methodology and also in ethics. It will temper the concept of what "due diligence" is owed to the client, requiring not only attention to current attack modalities, but also explicit attention *during design* to unintended consequences and next probable avenues of attack. We will never get ahead of the opponent if we patch current problems and wait for the next problem to emerge.

In my experience at NSA, when we designed a tactical radio for soldiers, we provided encryption to keep the opponent from over-hearing plans. But we understood that if the enemy could not read the traffic, they would try to jam communication to deny its contents to the soldier. So we also provided Anti-Jam capability in the radio, and also included low probability of intercept functions, knowing that the opponent would attempt to home in on the signal and destroy the radio (and troops) if they could not gain a useful advantage either by reading or jamming the signal.

A lot of thought went into the first build of a radio to counter as many direct threats and follow-on threats as we could envision, so the radio would have a reasonable useful life-span before being nullified by the opponent.

I will harshly over-simplify today's commercial security market build-and-patch strategy as providing one feature in the initial build, which is quickly countered by opponents. The vendor then provides the minimum patch, possibly at additional cost, which is quickly countered and generates the next patch . . . ad infinitum. The vendor does not really attempt to "game" the opponent's evolving strategy in order to provide robust gear with a reasonable operational life prior to needing patching or replacement due to evolving malice by the opponent.

This is no way to get ahead of the game, and is possibly focused more on protecting a profit stream than adequately protecting the client! This is itself an ethics problem.

Malice comes in two flavors: generic or targeted. Generic is aimed at anyone who happens to be vulnerable (open ports) the way robbery often hits unlocked houses. But targeted malice (whale phishing) is set up to get you in particular, unless you really put in a lot of effort to head it off.

The security practitioner needs to protect not only his client against targeted malice, but must also be prepared to handle malice directed at him in order to get to his client. This requires more ethical focus and advance preparation than other fields might face.

Given that computer security is a young field, almost universally viewed as not yet being fully equipped with adequate tools to do the job even against generic malice, there is also pressure on the practitioner to bow to a client's wishes for "anything cheap" to keep the auditors and compliance types at bay with appearance treatments, rather than paying more for better (but minimally needed) protection. There simply are not enough widely accepted (and valid) best practices or even nominal "customary and usual" business *security* practices to judge the performance of systems (or security consultants or advisors) against. *This will change over time* (see the report of Rueschlikon 2005 [8]) as the insurance industry steps in, but we simply aren't there yet.

We also need to describe the scope of the field, and what practitioners are included and which are not. Security consultants are clearly in need. What of system administrators? System integrators? OS designers? Systems engineers? Can we provide "addendums" to existing ethics codes in some of these fields, or do we need to start from scratch?

Also, we must be modest in our expectation of just how much protection (or constraint) a code of ethics (or best practices, or code of professional behavior) can actually provide. There is literature on codes and their limitations as well; [6] is just one such.

Computer Ethics has not stayed stagnant since Norbert Weiner's time; I again paraphrase Terrell Ward Bynum in the following four paragraphs [4].

In 1995 Górniak [11] predicted that computer ethics, which is currently considered just a branch of applied ethics, will eventually evolve into something much more. It will evolve into a system of global ethics applicable in every culture on earth: a few quotes from her work will show the flavor of her thought.

> Computers do not know borders. Computer networks . . . have a truly global character. Hence, when we are talking about computer ethics, we are talking about the emerging global ethic. [11, p. 186]
>
> . . . the rules of computer ethics, no matter how well thought through, will be ineffective unless respected by the vast majority of or maybe even all computer users. This means that in the future, the rules of computer ethics should be respected by the majority (or all) of the human inhabitants of the Earth. . . . In other words, computer ethics will become universal, it will be a global ethic. [11, p. 187]

According to the Górniak hypothesis, "local" ethical theories

like Europe's Benthamite and Kantian systems and the ethical systems of other cultures in Asia, Africa, the Pacific Islands, *etc*., will eventually be superseded by a global ethics evolving from today's computer ethics. "Computer" ethics, then, will become the "ordinary" ethics of the information age!

So there is plenty of energy in developing computer ethics. What is lacking is a sufficient focus on explicitly countering malice, the role of the cybersecurity practitioner. Attention to malice is indeed a major paradigm shift in the culture at large; let's get on with it!

Other speakers on the panel bemoan the inability to provide *detailed* proscriptive guidance in an ethics code due to the rapid evolution in the field; I agree. But that is a red-herring; I do not ask for a detailed proscriptive code or one that focuses on technology; merely a high-level Aspirational code establishing broad outlines of appropriate practice and behavior.

It is an appropriate starting point for an as yet immature field, and I believe a helpful one. We are certainly nowhere near the point where regulation or laws could safely be pursued

There is no need today to get into detailed wording of an actual code. In fact, any first version of an ethics code should be only the most general guidance stated in short, succinct sentences, not far from statements of Core Values. The more detail that is included, the closer we get to policies, regulation, and law. We are not ready for that yet.

I suggest that we focus our discussions today on developing a set of core concepts and values, and whatever brief wording addressing them that would focus on the security practitioner.

I offer the following topics not as a mature list, but merely as a draft point of departure. I'm sure discussion will lead to additions and deletions.

- Who is the client?
- Accountability.
- Dumb clients and/or pressure to "rush to market" do not excuse lack of proper exercise of available procedures and safeguards.
- Truthfulness — including clarity on the current limits of technology and practice.
- Prioritization of decision parameters.
- Full "due diligence" in planning prior to design work, including explicitly addressing potential unintended consequences of action, whether through success or failure.
- Protect the innocent from your actions (and their actions).
- Know the limits of your authority and scope of action.

Language already exists in current codes that address, at least partially, most of the above points. Let's look at a subset of language in each of three codes;

The $(ISC)^2$ code [14] has the following four canons.

- Protect society, the commonwealth, and the infrastructure.
- Act honorably, honestly, justly, responsibly, and legally.
- Provide diligent and competent service to principals.
- Advance and protect the profession.

It also has some specific guidance to discourage the following.

- Raising unnecessary alarm, fear, uncertainty or doubt.
- Giving unwarranted comfort or reassurance.
- Consenting to bad practice.
- Attaching weak systems to the public network.

The IEEE code [13] states the following.

2. Avoid real or perceived conflicts of interest . . .
4. Reject bribery in all its forms.
5. Improve the understanding of — potential consequences.
9. Avoid injuring others — by false or malicious action.

A *draft* "U.S. Intelligence Community Code of Mission Ethics"[3] says:

3. Expediency is not an excuse for misconduct.
6. We will resolve difficult ethical choices in favor of constitutional requirements, the truth, and our fellow citizens.
7. We will address the potential consequences of our actions in advance, especially the consequences of failure, discovery, and unintended or collateral consequences of success.
8. We will not impose unnecessary risk on innocents.

Note that for all three codes, these objectives are in the human domain, not the technological domain, and should hold up for some time.

In conclusion, I am willing to recommend asking employers to require employees lacking professional certification to abide by the $(ISC)^2$ code as a reasonable preliminary step to a more widely promulgated code that will address all germane issues, including malice.

3. *RICHARD FORD'S POSITION:* WE NEED ETHICS, BUT WE CANNOT IMPLEMENT THEM; FIDDLING WHILE ROME BURNS?

There is little doubt as to the value of ethics; ethical worldviews are implicit in shaping the "look and feel" of different disciplines. While this is certainly true for Information Assurance, there is a strong argument that despite their importance, coming up with a meaningful "ethical code" for security practitioners and users in general is a task that is both Herculean and pointless.

In order to provide a foundation for this belief, it is important to recognize the abysmal state of computer security today. SCADA systems running on Windows NT4 are still controlling highly-critical infrastructure — machines that are so unstable that they can be crashed simply by sending traffic to them! Developing nations are coming on line, bringing with them new victims for well-known (but still successful) Internet scams. The Botnet problem has exploded — some researchers have claimed to have discovered networks of over a million machines in one single network. Home users are using ten year old browsers to access "secure" banking information. By any measure, the security of both cyber and real space is crumbling. All it takes is the right pressure applied to the right place, and the losses will be vast. The state of the virtual nation is not good, no matter how much one might wish it to be.

[3]Not an official document, but a work in progress; ask `briansnow@comcast.net` for a full copy.

In addition, the computing environment is *highly* fluid: today's technological innovation is obsolete by the time it hits the shelves. This rapid progression makes maintaining a code of ethics that is highly detailed time consuming, as new issues continually arise and old issues change.

It is in this environment that computer security professionals have spent considerable time thinking about edge cases in computer security. How far can a defender go when under attack? Is it ethical to attack a system that is attacking your system? Is there ever a reason to write self-replicating code? Can one ethically be identified with the "hacker" culture? Given the preceding "state of the virtual nation" these questions pale into insignificance: computer security ethicists are, essentially, fiddling while Rome burns.

In addition to the environmental issues specifically related to computer security, there are also broader influences related to the shifting ethical stance within the last century. Post-modernism has begun to erode the foundations of classical ethical systems. Post-structuralism has begun to influence the minds of the next generation of computer users, and the inerrancy of fact has given way to a more dynamic idea of personal truths. While it would be easy to overstress the impact of these movements, the popular philosophy of the 20th Century can perhaps by summarized by *Il est interdit d'interdire*[4].

With this in mind, consider the task of constructing an ethical code of conduct for computer security practitioners. At the least granular level, the problem can be approached either as an exercise in applied ethics (that is, a set of rules specific to circumstances) or in a more normative way. While the normative is, at least on the surface, more attractive, neither of these approaches is likely to have a large impact at the level of action.

Beginning with rule-based ethical guidelines, the primary problems are:

1. The rapid rate of change in computing makes usable rules highly specific. Thus, any set of ethical guidelines will need continual updating.

2. Any set of ethical guidelines will be highly culturally-dependent. That is, an action in Company A may be entirely acceptable within the culture of Company A, but *verboten* in Company B. For example, user attitudes on privacy vary dramatically from company to company and country to country. There is no set of rules that will fit everyone.

3. Rule-based ethical systems fly in the face of popular culture, where many users attempt to "find their own truth." As such, these rules will be ignored.

4. Proscriptive ethics will inhibit "out of the box" thinking and innovation. Given the pitiful state of computer security, there is a desperate need to try new solutions, but incumbent players are often highly resistant to ideas that challenge the *status quo*. Many researchers have experienced first hand the resistance encountered when challenging deeply-help ethical beliefs. In security, the complexity of the problem cannot (and should not) be reduced to a set of "if–then" statements: this quashes innovation.

On the flipside, taking a normative approach is also doomed to failure as the process of ethical decision making is not taught well in the college system. Historically, the "jewel in the crown" of a student's education was their course in ethics. This course was frequently taught by the President of the university, and was seen as the completion of the student's education. In contrast, modern computing students typically have a very poor grasp of ethical concepts, and often are incapable of making well-constructed ethical arguments for different positions. Thus, many security practitioners simply lack the training to derive rules from ethical theories. Even if they had these skills, is there any pressing evidence that the "end would not justify the means"?

The outworkings of this show themselves in the incongruous ethical values many users display. For example, based on personal and direct interactions with students, many people see a huge difference between stealing a physical CD and "ripping" a copy from a friend. While both actions are illegal and (usually) against the stated request of the intellectual property owner, the virtual action seems to be societally acceptable — certainly within large subcultures. Compartmentalization of beliefs — especially between the private and the public — is a major complicating factor. This is part and parcel of post-modernism . . . and it is the worldview that many new professionals have.

Finally, despite the lofty goals of the numerous "code of ethics" professionals are asked to subscribe to when joining industry organizations, students receive terribly mixed messages from the "establishment." Sharing of copyrighted information is rampant, yet is not really handled well by the industry (how many of you have seen "fair use" applied to commercial presentations?). Furthermore, the educational establishment cannot agree on whether some basic stances are ethical or not (for example, there was a university course offered where students had to write viruses in an uncontrolled environment). For years, security professionals have said that hacking is no way to enter the security industry but (ahem) most researchers know that is not *necessarily* the case. "Do as I say, not as I do" is a highly counterproductive stance.

Given these challenges, it seems wiser to invest in more general ethical education and leave security-specific codes alone. Dealing with edge cases that will be irrelevant in a year or that will never be encountered in practice is not the best use of the limited educational resources that are available. Instead, time and energy should be invested in training to make users better able to reason ethically, especially with respect to the virtual/real world discontinuity so often encountered. As such, any attempt to lay down either a general set of ethical principles or a "one size fits all" set of rules is not likely to help and distracts researchers from the real problem.

3.1 Panel Chair/Editor's Note

The NSPW Small Discussion Group for this paper (prior to the plenary panel discussion) strongly recommended that we include a reference to Ford and Gordon's important NSPW 2006 paper. Due to Richard Ford's ethical reluctance for self-citation, we (the set of authors modulo him) cite it here for him. Ford and Gordon's paper, *Cent, five cent, ten cent, dollar: hitting botnets where it really hurts,* [9] raises several real-world ethical concerns for security practitioners and we recommend it as a good example of some of the very real ethical issues that some practitioners face.

4. *RICHARD THIEME'S POSITION:* CHANGING CONTEXTS OF SECURITY AND ETHICS: YOU CAN'T HAVE ONE WITHOUT THE OTHER

Because implicit ethical and moral dimensions emerge from new social and cultural structures as a result of technological transformations, any discussion of ethics in relationship to the implementation of new technologies must take into account a heightened

[4]French; literally: "It is forbidden to forbid" — a popular slogan of the May 1968 revolt in France

awareness of those dimensions. Because the philosophical and religious systems that animate society simultaneously undergo transformation, emergent paradigms must find expression in formulations as explicit and precise as possible and the implications of those paradigms correlated to new possibilities for action. Implications of this discussion for human identity at all levels necessarily inform this exploration.

Post World War II, R&D in the many branches of the intelligence community and military services have shared responsibility for creating technological engines that have transformed human identity and therefore the Kuhnian paradigm in which we frame possibilities for action. Action means options, and options mean ethics. I define "ethics" for the purposes of this panel as the options that are most congruent with our core notions of identity, self, integrity, and "the right thing to do."

Because all technological transformation processes cause a fundamental"identity shift," our awareness of options must be referenced to those transformational processes because they also alter religious experience, ideation and organizational structures and the way we frame ethical imperatives. It is, therefore, our first ethical imperative to be accountable to a fuller awareness of what this means for the people we serve by our work.

Definitions of everyday reality — privacy, security, legal guarantees — are being transformed by the technologies of surveillance, information, and communication. To articulate a moral dimension in order to formulate a basis for establishing the core values we bring to the various tasks of information security — attack, intrude, co-opt, subvert on one hand, and defend, preserve, and sustain on the other — we discover that we get that for which we test like a physicist determining whether photons are particles or waves.

"Common sense reality" is a function of the technologies from which our social and psychological lives emerge. Those technologies are invisible frames because we live inside the picture, so if we define ethical issues in the context created by prior technologies then we derive familiar recognizable and comforting concepts as a result, but ones that unfortunately no longer fit the real-life context created by new technologies. Our ethical decisions are, in short, inauthentic. It is not that we deceive others but that we first deceive ourselves. That is the heart of the problem.

We do not share a vocabulary, much less a consensus, for discussing how those technologies inform contemporary cultural structures. Yet the need to have this discussion is itself an implicit consequence of the changes I am describing.

Therefore, even a cursory exploration of ethical issues in computer security must include a meta-ethical dimension, one congruent with the newly emergent forms and structures of our lives, up to and including geopolitical and extraterrestrial structures (*i.e.,* confronting the realities mandated by permanent space colonies, lunar and Martian outposts, and the recontextualization of air and ground war by space war).

"All great truths," said George Bernard Shaw, "begin as blasphemies." [19] Today's blasphemy is tomorrow's "truth." Between times, however, we live in the fog of war. In a world which posits terrorists (*i.e.,* enemies of social and economic order) as the Other, the mind of society is the battlefield. Images and ideas are the primary weapons, and the means by which they come into being and move through human networks is the subtext of all security. The paradigms we use determine the questions we are capable of thinking and asking. The formulation of relevant questions may be more important than the answers.

A full discussion of this subject requires much more space than I want to fill, so let me highlight key concepts:

(1) Information security as one task, both offensive and defensive, of the intelligence community sanctions breaking foreign laws while prohibiting similar activities on American soil. But simple distinctions of "foreign" and "domestic" no longer hold. The convergence of enabling technologies of intrusion, interception, and panoptic reach, combined with a sense of urgency about the counterterror imperative and a clear mandate from our leaders to do everything possible to defeat an amorphous non-state entity defined by behaviors rather than boundaries, borders, or even a clear ideological allegiance, has created an ominous but invisible set of conditions that undermine the previous cornerstones of law, ethics, and even religious traditions.

(2) Identity is a function of boundaries. An "individual self" defined by a boundary around biological processes and the complex of energy and information radiated by those processes is undermined by the erosion of those boundaries by the use of connective technologies. The "individual self" we take for granted emerged a few hundred years ago from a cultural shift and is a social construction of reality. New technologies deconstruct it as we speak.

(3) Security, privacy, and intelligence gathering are corollaries of individual and national identities and how they relate to one another. Ethics is a description of "what works," *i.e.,* what is "right" for those identities at different levels of complexity and according to the ultimate goal, whether defense of a community or integrity of an individual.

(4) Security is a function of boundaries. Boundaries define the "other" that threatens "us" and "us" is a felt experience of clan, tribal, and societal kinship still. Prior to the emergence of writing and the religions it facilitated, the "enemy" was the "Other." Ancient societies defined the enemy as one who was not a member of the tribe. After the emergence of writing, the enemy morphed and became — in Christian scriptures, for example — that *in ourselves* which must be fought, resisted, or transcended. This shift in consciousness was a result of emergent technologies of writing. *This distinction is critical because security ethics exist in the tension created by these conflicting definitions*.

When the enemy is "within" the body politic, defined as an element that threatens societal order and economic well-being, defined no longer as a nation-state that threatens our political existence as a nation state, then the distinction between criminals and terrorists or dissenters and supporters of terrorism blurs. Accordingly the tools considered appropriate to their identification and neutralization will also blur.

We continue to speak of ethical norms in relationship to the cultural past as if it is still the context of our beliefs and actions. We speak of individuals as primary moral agents. We speak of nation states as primary determinants of our collective identities. We speak of the intelligence mission as if "we" who live inside one nation are intercepting or penetrating or subverting the technical processes and social dynamics of others who are also "inside" the boundary of a nation state that defines them.

Those distinctions no longer hold.

(6) Current technologies make speaking of interception obsolete. Our technologies constitute the physical framework, and software and informational contexts, of a pan-global society. Boundaries between elements of the network, between the networks that make up the network, that is, are arbitrary and porous. We live in a world literally without walls. Every attribute of a process or structure that broadcasts or transmits information about itself by any physical or electromagnetic means can be detected, often at the source. Often enough, those who built the system in the first place engineer information to come to them. "Here" and "there" are distinctions without a difference.

(7) What if that technology is reverse engineered and used against

Americans in a way than might be said to violate the Fourth Amendment, for example? The Moebius Strip nature of life in a networked world guarantees that unintended consequences must find their way back to the hands (and minds) that made them. In the same way, the idea of "blowback" from disinformation operations conducted in other countries is obsolete: all stories in all publications flow into the single information waters in which we live.

(8) Identity at a fundamental level is transformed. Digital identities can be appropriated, yes, but more than that, *we can invent them on the fly and determine at the moment of action or execution to which matrix we are related as a node in the network.* Our identities exist as potentialities made actual by our intention at the moment of action. They are the equivalent of quantum states, fixed only when expressed.

Identity in relationship to security then becomes a matter of observation and not assertion. Only multi-level observation penetrates the skin sufficiently to reach the meta-level determined by actions which may support or contradict identity-assertions.

(9) Computer scientist Langdon Winner wrote, "To invent a new technology requires society to invent the kinds of people who will use it, with new practices, relationships and identities supplanting the old." [23] In case after case, the move to computerize and digitize means many preexisting cultural forms have suddenly gone liquid, losing their former shape as they are retailored for computerized expression. As new patterns solidify, both useful artifacts and the texture of human relations that surround them are often much different from what existed previously.

This insight has implications for security and ethics. As boundaries go liquid, the task of defining appropriate behaviors in relationship to moral norms becomes difficult because the phrase "moral norms" is a metaphor for the *context* that is generally invisible to members of a society but not to sophisticated security professionals, an elite sanctioned to manipulate those underlying norms on behalf of ends considered important enough to justify a variety of means to achieve them.

Therefore:

Computer professionals exercise an implicit, *de facto* thought leadership because they create structures that bind and inform society and civilization. They create frames of human behavior that determine how we think about ourselves as possibilities for action. Their real implicit charge is not "to defend and protect a nation" but to stabilize a world.

On whose behalf are they acting? Who do they serve? To what end? On the level of the data themselves, the indeterminate but ultimate destination of the data and how they are aggregated to create an image of reality is lost unless the identity of the data and the people securing them are tracked precisely. In effect, people become instantiations of data because only data are meaningful in this context. Yet ethics posits "individual" human beings as the ultimate value in the universe, even as those "individuals" vanish like the grin of the Cheshire cat in the process.

In short: what's a guy or gal to do?

This process has happened before and will happen again. In the past, however, as Alfred North Whitehead said [5], such processes have often all but wrecked the societies in which they occurred. *The dire possibility of societal disintegration elevates the moral responsibility of the security and intelligence communities to a higher level. Linked in cooperative activity, they are responsible for maintaining social and global order at a level of understanding far beyond that formulated in the past by any one nation. These communities in the aggregate constitute a global community of practitioners who share an ethos and modalities of operation not available to ordinary citizens; they have thereby created for themselves an intrinsic vocation or calling to maintain global order in a way that is consistent with the ethical norms and moral order articulated by the great cultural traditions even as those traditions are also transformed by diverse technologies—and even though they and we recognize that in practice that moral order and those ethical norms are often violated as a matter of practice.*

Managing these concerns is quite a challenge. As Machiavelli said in *The Prince* during an equally transformational era:

> ". . .there is nothing more difficult to take in hand, more perilous to conduct, or more uncertain in its success, than to take the lead in the introduction of a new order of things." [16]

5. *STEVEN J. GREENWALD'S POSITION:* WE DON'T NEED/WANT ETHICS; ANYWAY, IMPLEMENTATION: IMPOSSIBLE!

"Un flic dort en chacun de nous, il faut le tuer."[6] — Slogan from the 1968 French Revolution.

I believe that if we reflect on the situation, we do not want or need specialized ethics for computer security. In other words, we should not go beyond the canonical professional society codes of ethics such as ACM [1] and IEEE [13]. Worse yet (from the standpoint of investigators into specialized ethics for our field): even if we wanted them we could not create them.

At a previous NSPW panel [10, § 5, pp. 40-42] I claimed that we had better act extremely carefully in regard to how we view the user non-acceptance issue or we will well and truly hurt not only ourselves, but our clients, customers, funders, innocents, society, and who knows what else? In that panel, I invoked Celine's *Laws of Chaos, Discord, and Confusion* [5] to show that in certain areas our actions will *always* backfire on us by causing unintended consequences for very specific reasons. Celine's Laws follow and I then show how they apply to our topic[7].

1. National security is the chief cause of national insecurity.

2. Accurate communication is only possible in a nonpunishing situation.

3. An honest politician is a national calamity.

5.1 We Should Not Want Specialized Ethics

In the entire history of the human race no effort of this sort has ever paid off in the way the originators intended (and due to the

[5]"It is the first step in sociological wisdom, to recognize that the major advances in civilization are processes which all but wreck the societies in which they occur:like unto an arrow in the hand of a child. The art of free society consists first in the maintenance of the symbolic code; and secondly in fearlessness of revision, to secure that the code serves those purposes which satisfy an enlightened reason. Those societies which cannot combine reverence to their symbols with freedom of revision, must ultimately decay either from anarchy, or from the slow atrophy of a life stifled by useless shadows." [22, p. 88]

[6]Translation from the French: "A cop sleeps inside every one of us, we must kill him."

[7]I do not jest and believe that Celine's laws, incredible as they might seem at first glance, apply with (quite literally) deadly effect

huge number of examples of the law of unintended consequences in our field we have little room for optimism).

We can take a very good guess at some of these unintended consequences based on what has happened in other areas.

1. **Celine's First Law**. Given the nature of our field, more ethics would quickly find a justification due to national security reasons. But Celine's first law tells us that this would lead to a great deal of additional insecurity.

   For example, a practitioner blissfully unaware of certain specific ethical issues would suddenly and forcefully have to consider those issues due to getting exposed to an explicit ethical code. This would result in the practitioner gaining new knowledge about how to ethically violate our systems. We could use, for this example, the case of a practitioner who had never before considered the $(ISC)^2$ ethical code and who therefore would come across the new (to this practitioner) ethical guideline of not giving "unwarranted comfort" — thereby immediately learning a new and perhaps profitable method for the unethical exploitation of people.

   Worse yet, as time goes on any ethical code we create would accumulate more and more examples, case studies, *etc*. until it became a veritable cornucopia of unethical strategies and tactics. Do we really want to create ethical codes that also function as guides in how to unethically exploit? Of course not.

2. **Celine's Second Law**. An ethical code presumes the existence of some authority that can provide (in at least some cases) ethical guidance.

   Certainly some ethical situations can only be resolved by the person in the sticky situation. But most sticky situations can benefit from the guidance of others. But what others? If we had a specialized ethical code, then due to increased demand for authorities to give ethical guidance we would soon have professional ethicists in our field, not to mention lawyers (even though law and ethics are different), ethical boards, and so forth. And of course, due to the presumed authority of these boards most people seeking their guidance would feel inferior in ethical decision making skills. This means most people would have no way to judge the value of this "guidance." Worse yet, if a person's conscience told them that they had received bad "guidance" then this would put the person in a worse situation than before because the person would feel threatened by the (ethically) nominally superior authorities. Unless we are willing to foster "ethical anarchism" (by which I mean that we would have no superior ethical authorities) then this situation must surely develop as it has in every other professional field that has adopted a specialized ethical code.

3. **Celine's Third Law** tells us that in the political realm well-intentioned do-gooders rarely do anything other than create the existence of a new criminal class (because every law they pass makes something new illegal; consider how new laws against illegal drugs create a new criminal class from previously law-abiding people, how intellectual property rights create a new criminal class out of ordinary law-abiding people who do things like exchange certain music files, or even how a law as seemingly benign as banning indoor smoking in public places creates a new bunch of law-breakers[8]).

   The same sort of creation of a new class of law-breakers would certainly happen if we developed a specialized ethical code. The more ethical "case studies" and "examples" and such that we create (or worse yet, actually codify as ethical principles), then the more ethical misbehaviors will also get defined/created by default. This will therefore lead to the creation (by definition/fiat) of a larger set of misbehaving people. You might feel tempted to view this reasoning as some sort of sophistry because labeling actions that have already occurred certainly can not actually cause those actions. Quite true; but please consider the following.

   Right now, a good person might unwittingly, and therefore innocently, break a minor ethical principle that has little or no consequence and, rightly so, not worry overmuch about it. But if that same person felt seriously bound by an ethical code that made it clear the person had broken a codified ethical principle, then that person might very well create feel guilty or culpable and thus consider himself a wrong-doer. This would certainly lead to terrible consequences among many people. For example, some people would drop out of the field. Other people would feel that so long as they were guilty of bad actions and felt self-labeled as "bad" people, then they might as well get damned for some really profitable or interesting behavior, thus leading them to act increasingly worse in a vicious cycle of ethical corruption. If you have difficulty swallowing this argument, then please note that the Marquis de Sade demonstrated that making something wrong, immoral, or criminal often encourages people to perform the act for the thrill of the misbehavior, or even worse can liberate them from refraining from performing worse wrongs [17].

[8] I could give literally thousands of examples of these, but I restrain myself.

Considering that we run the very real risks of educating people about how to behave badly, stopping many people from trusting their own ethical judgment in favor of questionable ethical authorities, and also creating (by fiat) more misbehaving people, it seems to me that the very *last* thing we *want* is a specialized ethical code!

5.2 We Cannot Accurately Implement Specialized Ethics

Even if we wanted a specialized code of ethics for information security, I submit that we could not accurately implement such a code for several reasons.

5.2.1 Lack of Maturity as a Field

As a new discipline computer security simply is too weak and underpowered (in terms of organization) in order to do any proper job regarding ethics. Why? Because in order to have a meaningful payoff ethics must get adopted by (mature) fields that can actually enforce those ethics (for example, politics, medicine, engineering). Our field has not yet gotten to that point; therefore attempting to come up with an ethical code would simply waste time. Some will no doubt disagree with this, and while I would not characterize this as an especially strong argument, fairness compels me to point out that even if true we should not necessarily abandon any attempts (at least for this reason).

5.2.2 The Problem of Rational Behavior

We have a much stronger issue: what economists term "rational behavior"[7, pp. 427–431]. Related to the well-known problem of the "tragedy of the commons," the rational behavior of an individual (or certain individuals) serves to maximize some objective function (usually a utility function) that the individual uses to mea-

sure economic value. This behavior often goes against the strategy that works best not only for the group, but often for the long-term (or sometimes even short-term) benefit of the individual. Consider things like the *Prisoner's Dilemma* in game theory: in many cases the behavior that most benefits an individual involves a type of cooperation that also winds up most benefiting the group. But due to issues with trust (as well as other things) individuals often choose a course of action that they think may benefit them at the expense of the group even when everyone agrees that the optimum behavior would benefit both.

I believe this would apply to ethics in our field. We can view the current internetworked environment as an ethical "security commons" in which no one owns or even controls many of the critical resources. This means that if we wish people in the security community to practice a particular ethical code then we can foster this behavior in only a few ways (as determined by economists who study these issues).

1. Government. I do not find this likely given the nature of things; even if government did get involved, it would likely concern itself with laws and regulations, and not ethics.
2. Outside force. Also not likely given the nature of our resources.
3. Privatization. Possible, but currently very hard to imagine. Could we really privatize enough of the resources that fall under our ethical concerns to make a difference in terms of ethics?
4. Cooperation of everyone through enlightened self-interest. I don't think anyone would buy this. While this does sometimes happen, in our field the many necessary conditions for this do not exist (for one thing, we lack homogeneity as a group).

5.2.3 Ethical "Mutation" and Restrictive Evolution

Even if we could somehow get enough of the right conditions to create and actually implement an ethical code, what would then happen? Well, given the rapid evolution and short generation period of the technology in which we work, I believe an ethical code would start to immediately mutate and therefore evolve; very soon it would bear little relation to ethics. In fact, laws (secular and religious) tend to evolve this way. We can view it as a sort of ethical strictness function where the strictness monotonically increases over time. In other words, ethical codes evolve into more stricter forms, never (absent revolutionary events) less restrictive.

Many religions have evolved in this direction. For example, simple religious laws in Judaism have evolved over the millenia into extremely restrictive versions that bear little resemblance to the original laws (or, presumably, the religious intent). A good example of this involves the laws prohibiting Jews from engaging in a particular Canaanite fertility ritual: boiling a baby goat in its mother's milk (Exodus 23:19, 34:26; Deuteronomy 13:21). The law has evolved drastically [18] into a prohibition against eating any meat and milk together. Or any meat and milk products. Or eating or preparing meat (or milk) on (even spotlessly clean but ritually unclean) dishes and utensils used for the other. And having to wait specified amounts of time before eating meat after consuming milk products (or vice-versa, except the times have radical differences). And so forth. This might serve some modern religious purposes, but it clearly goes far beyond the intent of the original law (and dare I point out the irony that we moderns would not even know about the existence of this ancient Canaanite practice if not for the existence of the law itself?)

The reader might think that Jewish legal authorities had no awareness of the ethical problems related to these types of codification. But we have just the opposite situation: about 2,500 years ago, during the harsh Babylonian Exile, Jewish sages agonized over whether to codify the oral law tradition. They faced a cruel dilemma: codifying the oral tradition would tend to eliminate future flexibility, and yet not codifying it would almost certainly lead to its loss due to the particularly nasty conditions of those times. Their choice of codification led to a doctrine of putting a "fence" around the original laws (in other words, making them more restrictive to prevent people from breaking them).

A code of ethics for security would soon go in this direction. Even the simplest ethical principle would accumulate "ethical encrustations" that would render the code too restrictive, outside the scope of the spirit of the originators, complex, and unwieldy. Some might argue that if we had an ethical authoritative body (analogous, perhaps, to the U.S. Supreme Court) then this need not happen. But consider: if such a body were created it would soon turn ethical principles into specific laws. Because I view law as the enemy of ethics I would view that as a disastrous development

Because of this, and a lot more, I do not view the pursuit of a specialized ethical code for our field as desirable or possible.

6. NSPW SHEPHERD'S COMMENTS

(*Panel chair/editor's note: The NSPW Program Committee appointed us a "shepherd" as part of the workshop process; we so highly valued our shepherd's comments that we decided to include them in this special section.*)

It seems to me that a non-intuitive part in considering a code of ethics for the security realm would be to consider that nature abhors a vacuum, and to examine what existing structures might be filling the role of such an ethics code as part of the *status quo*. I propose that we examine the items in the following outline, and consider what fills each of these voids now.

6.1 Source

When making a decision, there are the explicitly stated rules, and there are situations where one has the notion that breaking the rules may be the right thing to do. What or who is the source for the "higher principle" that justifies that choice? Note that these sources might take many forms — organizational, religious, philosophical, pragmatic (whatever "works" best), monetary/economic, altruistic, *etc*.

6.2 Duty/Responsibility

This covers the question: Whose welfare is to be considered of highest priority? If multiple parties might be harmed by some action, what is the order of preference? To whom are one's loyalties properly bound? This may be a complicated question in the security realm, because practitioners may have a real or perceived duty to multiple parties, whose interests may be in conflict. What's a security wonk to do?

6.3 Benefits/Consequences

What is the perceived outcome of the various courses of action under consideration? Why is one course deemed more beneficial than another is? What are the metrics?

6.4 Calculus of Choice

Given multiple courses of action, what is the method for doing the trade-offs? Is it rigid or situational? Factoring in the previous items, one can see that even at its clearest, this will turn into an

exercise in equations of risk, benefit and consequence in multiple variables and dimensions.

6.5 Roles

In the explicit or implicit model of proper behavior guiding one's actions, there may be an issue of roles. What is proper behavior in one role may be anathema in another. How does the individual in an ethical quandary identify his or her role, and reason about what constitutes proper behavior for that role? As noted earlier, what is the source for the definition and constraints regarding one's role? Also, an individual may also see himself or herself in multiple roles — employee of XYZ Corp, adherent of ABC religion, member of a given community/race/ethnic group, and so forth.

6.6 Exceptions

It seems like every system of reasoning has some "gotcha" point. What are the exceptions or undefined states in whatever structure or rule set the individual is using as a basis for reasoning about the situation and the potential courses of action? Note that the existence of exception conditions in one's primary source of guidance is likely to be at least part of the source of the ethical question to begin with. If the situation were neat and tidy, there would likely be no ethical quandary for the individual to ponder.

6.7 Alternative Models?

Since there is no code of ethics, the gap may be filled by multiple different structures or drivers, each covering all or part of the territory that would ideally be covered by a complete and consistent ethical framework. An individual may consider himself or herself to be a faithful follower of several sets of rules, in some combination. The deconfliction process may be conscious and rigorous or it may be intuitive and internally inconsistent. How this deconfliction currently occurs will be of value in assessing how things would be different if an ethical code or framework were to be established. Before considering the introduction of yet another source of guidance, one must consider the inter-relationships between all the explicit and implicit structures that already exist.

Once one has assessed the existing state, one can then begin to reason about the potential impact of introducing a new ethical code or framework, of whatever degree of formality.

6.8 What Does Not Work?

This is where one examines the failings of the *status quo*, the probability of those situations, and the consequences of them.

6.9 How Would Things Improve?

Would the introduction of an ethical framework of whatever sort address those gaps? Would it make things better, by whatever measure?

6.10 Cost/Benefit

What would be involved in introducing and maintaining the proposed ethical code, model, or framework? Is it worth it?

6.11 Post Workshop Comments

At the risk of seeming fixated on my own ideas, while at the workshop I wanted to hear opinions from the panelists on the general thrust of my "shepherd's comments." I think they can be boiled down to the following three questions.

1. In the absence of a code of ethics, what existing structures fill whatever void there may be, either in whole or in part, and what are the inter-relationships, biases, gaps, *etc*. of those structures?

2. What problems are there with the existing structures, and how do we anticipate that an ethical code would affect the status quo represented by those structures, for good or ill?

3. What is the cost/benefit for whatever ethical code or structure that we might propose — is it "worth it?"

I'd also like to further explore the assertion, left largely unquestioned at the workshop, that ethics in our context is "different" due to malice. I'm not sure that the ethical constraints on the practitioner's behavior change much, if at all, other than needing to consider malice as part of the threats to be countered, both by the systems designed and the practitioner himself.

I personally think that "malice" is a bit of a red herring. I think that the ethical code is made more complex by the range of potentially conflicting allegiances (to one's employer, to the users, to the security community as a whole, to the owner of a system in which one finds a vulnerability, to the developer/vendor of the system, to the people whose information is being protected (who may or may not be the "users,") and so on) than by the presence of malice. Compare this to other "professions," where there are usually at most two allegiances specified - first to the client/patient, and then far more vaguely to "the profession."

7. WHAT HAPPENED AT THE WORKSHOP

Anil Somayaji, one of NSPW's two panel chairs, proposed a slightly new workshop format in order to improve the quality of the workshop interaction (ironically, due to issues related to his wife's parturition, he could not attend and therefore missed experiencing first-hand the success of his proposal). The *Somayaji Process* creates small groups for each paper that meet with the authors/presenters before the (up to that time) normal plenary workshop. The process has the goal of improving the already highly interactive NSPW method. We have split this section into two subsections to correspond to the *Somayaji Process* and therefore give the reader a glimpse as to how things proceeded *in situ*. We only note those comments that we have not directly incorporated into other parts of this paper.

7.1 The Small Group Discussion

For our Small Group Discussion participants, along with all of our panelist/authors, we also had Angelos Keromytis, John McDermott, Sean Peisert, and Cristina Serban (all old-hands regarding NSPW). We authors felt that two issues that we discussed should get addressed as separate topics.

- Our generalized discussion of the panel paper led to a small-scale discussion of relativism in the context of information security ethics. Steve Greenwald asserted that he viewed the term "evil" (and related terms) for the purposes of this topic as "behavior that acts as contra-survival for the group." The rest of the small group participants decided to give this some thought. Steve further asserted that such a definition might have potential for leading to a more objective ethical code for our field. However, we all agreed that the issue falls outside the limited scope of the panel.

- One of the non-panelists made a comment that caused Richard Thieme and Steve Greenwald to observe that a lot of people in our field have a mistaken impression of the Nietzschean *übermench* concept. Steve noted that the "establishment" (for want of a better term) often punishes or demonizes such people in order to maintain the *status quo ante*. In other

words, the "system" must suppress the subversive (new paradigm) ideas of such people in order to maintain the old paradigm and even, in some cases, to survive. Steve also noted that he thought this related to Richard Thieme's idea of how *meta-experts* operate, as the *übermench* concept often fits the bill regarding what many perceive as the role of meta-experts, *viz.*: someone who knows how to appropriately break (or make anew) the rules of the system in order to realize the spirit of those rules. We felt the implications of the notion of such meta-experts and the way they often get mischaracterized has implications for a code of ethics in our field.

7.2 The Plenary Workshop Panel and Ensuing Discussion

After a brief panel introduction, each of the four panelists had five minutes to summarize their position (we presumed, as usual, that each attendee had read the panel pre-proceedings paper). We then opened it up for the usual highly interactive discussion so typical of NSPW and did not get disappointed. Virtually all of the attendees participated and we received outstanding handwritten notes from our tireless NSPW scribe, Bob Blakley, who, as always, did an excellent job or recording the interaction. The following list comes from Bob's notes as well as our recollections based on those notes.[9]

1. As part of his presentation, Richard Ford asked the question: "Where do we learn ethics?" Almost immediately, NSPW's founder, Holly Hosmer, answered, "We learn ethics starting in kindergarten." Of course, Holly did not mean that literally in all cases, but she meant that in some sense we imbibe ethics along with our "mother's milk."

2. Bob Blakley asked an excellent question that elicited many responses: "We have asked society to entrust us, as a profession, with their protection. What do we promise to do if they say yes?"

   (a) Steve Greenwald felt Bob's question had tangential aspects to the user non-acceptance paradigm discussed in [10, § 5, pp. 40-42]; he then answered, "Then I tell them that I promise to do the best I can."

   (b) Brian Snow said, "We can't protect people without their participation, so we shouldn't make such a promise."

   (c) Matt Bishop stated that he didn't feel convinced that we get asked to protect things — it just happens; he wondered if that made the ethical issues more difficult in this case.

   (d) Angela Sasse said she thought that one does not need to observe ethics all the time in order for ethics to improve the situation.

   (e) Richard Thieme responded to Bob's first point by noting that "enemy" used to mean "other"; Jesus and Hillel redefined "enemy" as "that within us that resists integrity, goodness, and wholeness."

3. Bob stated that he thought that Steve made the logical fallacy of generalizing from particulars in his panel statement, thus: Steve generalized "Specialized Ethics is bad" from particulars such as "The ACM code is bad." Steve denied this in the sense that he did not claim that specialized ethical codes must always lead to bad things; only the ones he has seen in terms of specialization in our field or related fields, and that we should therefore act extremely carefully before we get in over our heads.

4. Several participants noted that some ethical codes, like those found in the U.S. military encourage compliance with easy rules (*e.g.*, do not steal) but encourage "breaking" the big rules (*e.g.*, do not kill).

5. Steve mentioned that he thought that a specialized code of ethics will cause literally mindless compliance at the expense of the true thoughtful professionals who have not really needed to internalize ethical principles. Richard Ford agreed, and replied that this highlighted one reason why he believes teaching ethical reasoning is critical.

6. Abe Singer pointed out that in some fields people will decline to perform certain work if they do not have the expertise to do it. Bob thought that Abe made the following inference: because of a lack of ethical expertise we do not do ethics.

7. Richard Thieme pointed out that the American Psychological Association currently lobbies to relax rules regarding research on human subjects to allow less review. Does this mean that ethical review actually inhibits work, or that practitioners just do not want ethics?

8. Richard Ford notes that ethically regulated professions have the structure of a group which confers advantages. We don't have a group structure, mainly due to the newness of our field. Bob, simplifying Richard Ford's statement, said that he did not view the core problem as ethics — rather, as a profession we have no standards at all for anything.

9. Klaus Julisch noted that IBM introduced new core values a couple of years ago; a lot of people were skeptical at first but they now serve a good purpose: to provide basic guidance in a complex world. However, Steve distinguishes core values from ethics; he thinks ethics in such a setting would turn into an organizational committee of ethical authorities/experts where a person who wished to succeed in the organization would not dare to disagree with those ethical authorities/experts even if it conflicted with that person's ethical beliefs. Richard Thieme concurred by noting that in many cases people know the right thing to do.

10. Abe said he wanted to have an ethical code so that he could cite it to his boss as a justification when necessary. Bob agreed. Steve disagreed; he thinks that such a code will eventually evolve so as to prevent people from doing the right thing in the sense that it will constrain truly ethical people into an organizational "one size fits all" mold.

11. Abe said he has an issue with corporate ethics because he thinks that these codes only operate until shareholders decide that they get in the way of profits. Steve disagreed: not "shareholders" (who often have little power or influence, *per se*) but rather the board of directors who report to shareholders.

12. Jon Solworth asked if we are closer to lawyers or doctors in terms of our ethical obligations?

[9] Please note that while we prefer to use a participant's first name after first using their full name, because we have two panelists who share the same first name "Richard," we have forgone that practice with them.

13. Holly asked if we can change the paradigm if we *are* the paradigm. Richard Thieme mentioned one way to do that: walk out of the room. Literally. (No one did, by the way.)

14. Richard Thieme asked that everyone in attendance who had experienced an ethical problem in the field of computer security to raise their hand. Steve felt quite surprised that only seven people (out of about 32) raised their hands. After the panel, later investigation showed that only these seven had actually practiced computer security in a way where ethics got involved. Steve felt very surprised at this as he felt it indicated that only those seven actually practices computer security in any substantive way.

8. CONCLUDING REMARKS

Several important points emerged from this effort.

- We cannot stress strongly enough the Kuhnian paradigm shift aspects of this topic. Any code of ethics must consider the larger potential role of *malice* in the proper execution of the discipline as opposed to other fields. For example, physicians treat patients expecting certain patterns of outbreak of diseases; they do not expect bacteria tailored for lethality for just one particular targeted patient. Engineers build bridges expecting environmental impacts already well-modeled and understood (floods, winds, lightning, traffic load, *etc*.) They have only recently even begun contemplating building bridges designed to mitigate terrorist bombs planted at rush hour. The paradigmatic aspects of this have very deep implications.

  Societally, we generally structure ourselves to accept that malice exists in the world, but it does not function as the dominating driver. As computer scientists, we have approached our ethical code with that assumption. However, in security, we have as a major difference the fact that malice becomes *the driving factor*. In other words the "society" that revolves around INFOSEC has an entire commitment to the worldview that malice, and not collaboration, primarily drives the world. This does not mean that we should view most people as malicious; most people do not fall into the malicious category. But the actions of those few malicious folks certainly drives the system.

  The idea of malice really makes its impact in this area. We do not assume that malice does not exist, or that we need malice-neutral ethical codes, but that in our section of the industry malice works as a kind of driving force. In such a world we must have more complex ethical rules to live by, as we have to trade off action against action. Ultimately, our ethical worldview does not change; but the acceptance of malice as a *driving force* in the world may change our view of exactly what right action we must perform at any particular time.

- Security systems require control systems that maintain control of the system in properly authorized (and audited) hands, even if malicious attackers seek to wrest control for their own benefit or to harm the true system owner or the system's clients. One of us (Snow) has characterized this as *robust control*. Most other design disciplines do not face this concern (but more of them should). A recent pointed example involves control surveillance systems doing data mining on many good folk, seeking the bad among them [3]. What happens if bad guys gain control?

- Too many folk in the U.S. and other countries seem willing to trade privacy or other rights for security or safety. They wrongly feel that one cannot gain security or safety without a corresponding loss of privacy or rights. We view that as simply wrong. We would hope that an ethical practitioner would seek solutions first in the technical domain without first seeking to pluck the possibly easier low-lying fruit available by doing this trade-off. Do not take from your fellow citizens (even if willing) if you can find a viable solution elsewhere, ideally in the technical domain.

- While many workshop participants expressed a strong desire for a specialized code of ethics, we found that many (if not most) have concerns about the immaturity of the field acting as a severe obstacle to forming a true code of ethics. Our field does not have the maturity of other disciplines such as engineering, law, or medicine.

- All four of us panelists now believe that no matter what disagreements we may have regarding a specialized ethical code, that the field of INFOSEC/Computer-security/Cybersecurity (call it what you will) would benefit enormously from a set of *core values*. We leave the creation of the *exact* set of core values for future work.

Where do we go from here? Clearly we need to start work on a set of *core values* that the vast majority of us can accept.

Acknowledgments

We gratefully acknowledge the superb help and contributions that our NSPW shepherd, Ken Olthoff, gave us. We received a remarkable total of 18 separate reviews from the NSPW program committee; we found them thoughtful and invaluable and we appreciate the great amount of work that went into the reviews by the program committee, and especially the program chairs, Angelos Keromytis and Anil Somayaji. Special thanks go to our "small group" attendees: Angelos Keromytis, John McDermott, Sean Peisert, and Cristina Serban; among other things, they rendered great assistance to us and certainly improved our panel presentation and helped improve the quality of the interaction. We received a huge amount of comments during the workshop panel itself that the tireless NSPW scribe, Bob Blakley, rendered for us, thus making our work much more accurate and easier. We also thank the entire set of NSPW participants for their outstanding input. Three of the four panel participants would like to thank the NSPW Financial Aid system and its sponsors for providing the funds necessary for them to attend. We would also like to thank Laura Corriss for her help in reading several drafts of this paper and making valuable corrections. We also thank Edward L Reid, II, M.D. and Percy Aitken, M.D. of the Miami Medical Forum for discussions about the medical ethics paradigm vs. the information security paradigm, particular the assumption of good will in the former vs. malice in the latter. As editor, any mistakes appearing in this paper happened solely due to Steve Greenwald.

9. REFERENCES

[1] ACM Council. ACM code of ethics and professional conduct, October 1992. `http://www.acm.org/about/code-of-ethics`.

[2] AMS Council's Special Advisory Committee on Professional Ethics. Ethical guidelines of the american mathematical society, January 2005. `http://www.ams.org/secretary/ethics.html`.

[3] S. M. Bellovin, M. Blaze, W. Diffie, S. Landau, P. G. Neumann, and J. Rexford. Risking communications security:

Potential hazards of the protect america act. *IEEE Security & Privacy*, 6(1):24–33, January/February 2008. Available at http://www.crypto.com/papers/paa-ieee.pdf.

[4] T. Bynum. A very short history of computer ethics. In J. Dorbolo, editor, *American Philosophical Association's Newsletter on Philosophy and Computing*, volume 99, pages 163–165, Newark, Delaware, USA, Spring 2000. The American Philosophical Association, University of Delaware. Available at `http://www.apaonline.org/publications/newsletters/v99n2_Computers_05.aspx`.

[5] H. Celine. Celine's laws of chaos, discord, and confusion. In R. Wilson, editor, *The Illuminati Papers*, pages 118–125, Berkeley, California, 1997.

[6] P. Davis. Codes of ethics and their limitations. *SIAM News*, 40(9), November 2007. Book Review of *Computer Ethics: A Global Perspective* by Stamatellos, G.

[7] J. Diamond. *Collapse : How Societies Choose to Fail or Succeed*. Penguin, December 2005.

[8] K. Dukier. Ensuring (and insuring?) critical information infrastructure protection. In *2005 Rueschlikon Conference on Information Policy*. The Rueschlikon Conference, September 2005.

[9] R. Ford and S. Gordon. Cent, five cent, ten cent, dollar: hitting botnets where it really hurts. In *NSPW '06: Proceedings of the 2006 workshop on New security paradigms*, pages 3–10, New York, NY, USA, 2007. ACM.

[10] S. Greenwald, K. Olthoff, V. Raskin, and W. Ruch. The user non-acceptance paradigm: Infosec's dirty little secret. In *Proceedings of the 2004 New Security Paradigms Workshop*, pages 35–43, Nova Scotia, Canada, September 2004.

[11] K. Górniak-Kocikowska. The computer revolution and the problem of global ethics. *Science and Engineering Ethics*, 2(2):177–190, June 1996.

[12] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Proceedings of the 29th Annual IEEE Symposium on Security and Privacy*, May 2008. Outstanding Paper Award.

[13] IEEE Board of Directors. IEEE Code of Ethics, February 2006. `http://www.ieee.org/portal/pages/iportals/aboutus/ethics/code.html`.

[14] $(ISC)^2$. Code of Ethics, 1998–2008. `https://www.isc2.org/cgi-bin/content.cgi?category=12l`.

[15] T. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 1970.

[16] N. Machiavelli. *The Prince*. , 1515.

[17] Marquis de Sade. *Justine (Les infortunes de la vertu)*. , 1791.

[18] Rabbi Yosef Karo of Safed. *Shulchan Aruch*. Yoreh De'ah, Chapter 6, Mixing Meat With Milk, ca. 1560. Available from http://www.torah.org/advanced/shulchan-aruch/classes/chapter6.html.

[19] G. Shaw. *Annajanska, the Bolshevik Empress*. , 1917.

[20] N. Weiner. *Cybernetics: or Control and Communication in the Animal and the Machine*. Technology Press, New York, New York, 1948.

[21] N. Weiner. *The Human Use of Human Beings: Cybernetics and Society, Second Edition Revised*. Doubleday Anchor, Garden City, New York, 1954. This later edition appears better and more complete from a computer ethics point of view.

[22] A. Whitehead. *Symbolism: Its Meaning and Effect*. Fordham University Press, 1927.

[23] L. Winner. Who will we be in cyberspace? *The Network Observer*, 2(9), September 1995.