

Each SQL query against our database produces at most one thousand records, so Michael wrote a little C routine to sort those records according to social security number. I cannot, for the life of me, figure out why he insisted on doing his own sorting routine, since SQL provides for sorting a query result. In technical terms, Michael was using embedded SQL in a C host language environment, so he wrote the sorting routine in C. The sorting routine he used required the swapping of records produced by the SQL query. He did this by swapping the data field by field. He did a normal swap for all fields except for the infectious disease blood test field. Now, this is really bizarre. Instead of doing a normal swap using a temporary variable, he used a C feature where you can do a swap without a temporary variable using the bitwise exclusive OR operator. Unfortunately, he implemented this incorrectly, with the result that some HIV positive codes were dropping out as negatives.

The fact that this error could have slipped by is evidence that Red Flag has serious deficiencies in terms of software quality assurance and testing. The problem is that we have a mentality where writing new applications is rather routine and testing new applications is done rather informally. In fact, I thought that this application was tested rather exhaustively, but obviously our testing was not adequate. As it turns out, none of the test data suites for Mr. Keating's application included an HIV-positive code of 127, and hence the deadly data was never detected. This was the only input value for which Mr. Keating's code does not work properly. But, I must say, his code is a mess. I can't say that I understand all of its intricate obscurities.

PULITZER: A security system is only as strong as its weakest link, wouldn't you say?

MOORE: Yes.

PULITZER: Isn't the fact that Michael Keating's software could destroy the integrity of your data a security issue? The way I see it, Michael Keating had de facto write access to the sensitive data in the database!

MOORE: No, he did not. He had neither read nor write access to the database.

PULITZER: But his programs could update the database, so it seems to me that he had effective write access to the database. He could compromise the integrity of the database.

MOORE: Okay. I see what you're trying to say.

PULITZER: Someone who writes applications programs for a database of this nature should be treated as a person who could potentially do great harm. Someone should have looked more carefully into his background and his character.

MOORE: I couldn't agree with you more, Ms. Pulitzer.

PULITZER: How many people received tainted blood as a result of this incident?

MOORE: Thirty-one.

PULITZER: Is there any evidence that Michael Keating acted maliciously to subvert the blood donation database?

MOORE: That's not for me to say. That's a police matter.

PULITZER: Did you find Randy Samuels's essay "In Praise of Obscurity" tacked to the wall next to Michael Keating's workstation at Red Flag?

MOORE: Yes, I did.

PULITZER: Would you have allowed Michael Keating onto your team knowing what you now know about his previous history?

MOORE: Absolutely not. I would have told Jane Farnsworth that I don't want this guy working here at Red Flag, and she would have kept him out. At Red Flag, we're dealing with people's lives.

PULITZER: I would like to thank my guests for appearing on "Roundtable" this morning. I hope the viewing audience learned some things about data and data privacy. This is your "Roundtable" reporter, Pam Pulitzer, wishing you a pleasant Sunday. Look for my article on the Red Flag situation this morning in the Sunday *Sentinel-Observer*!

Transcript of "Roundtable," the Sunday Morning Public Affairs Program, Broadcast on Public American Television

Transcript of "Close-Up," Public Television's Premiere Program about Academics and Academe

IS YOUR COMPUTER STEALING FROM YOU?

CLOSE-UP: Good evening. I am Stanley Kahn, reporter for the Silicon Valley *Sentinel-Observer*, and your host for this week's edition of "Close-Up."

This evening's guest is Professor Jacob Lowe-Tignoff, professor of religious studies at Silicon Valley University. Professor Lowe-Tignoff received his Ph.D. from Temple University. He is also ordained as a rabbi in the Conservative branch of Judaism. Last year Professor Lowe-Tignoff's book, *Is Your Computer Stealing from You?*, won several prestigious awards. It was on the *New York Times* Best Sellers list for twelve straight weeks. This evening we

are going to discuss Professor Lowe-Tignoff's thesis that computers can potentially steal away important human capabilities.

Welcome to "Close-Up," Professor Lowe-Tignoff.

LOWE-TIGNOFF: It's a pleasure to be here with you.

CLOSE-UP: Your book states **that** we need to develop a theory of ethical behavior for computers.

LOWE-TIGNOFF: More than a theory-actual guidelines. We need to develop ethical guidelines for computers that are similar to the ones that we already have for people. I am certainly not the only one who is discussing these issues.

CLOSE-UP: When I read what you had to say, my first reaction was that it was . . .

LOWE-TIGNOFF: Ridiculous! Let's be honest. Let's not mince words. People think that ethical constraints are for humans. They don't envision the possibility that we might have to impose ethical constraints on computers as well.

CLOSE-UP: Okay. At first I thought that the idea of ethical guidelines for computers was ridiculous, but by the time I reached the end of your book, I was a true believer. Computers have the potential to do great damage to human beings, as you emphasize.

LOWE-TIGNOFF: In the book I talk a lot about "Human Being," with a capital H and a capital B. "Human Being" refers to our humanity in the broadest possible sense. It refers to all aspects of our being human. I am raising the issue of whether computer systems should be constructed if they have the potential to diminish Human Being.

CLOSE-UP: Your book makes a subtle distinction between the need for ethical computer systems and the need for moral computer systems.

LOWE-TIGNOFF: My book suggests that professional societies develop ethical guidelines for computer systems. That is, using reasonable principles that most computing professionals can agree upon, we need to place constraints on computer systems and also to describe desirable properties for computer systems. I fear that unless the professional societies act first, then legislative bodies may try to regulate computers as a matter of public policy.

These ethical guidelines for computers should be somewhat narrow in focus, like the ethical guidelines that professional societies have already produced for computing professionals. For example, my book attempts to establish the principle that computers should not be allowed to steal as a fundamental principle behind any such set of ethical guidelines.

I also believe, as a religious believer, that computer systems should be moral. If someone tries to develop a computer system that deeply violates my sense of morality, I will try to oppose it. That sense of morality essentially derives from my religious beliefs. Since people have differing religious and philosophical beliefs, there will be conflict over issues of morality. In a

democratic society, all sides have an opportunity to freely express their perspectives.

CLOSE-UP: Would you say, then, that you wrote more about ethics than about morality in your book because you felt that there is a more universal consensus within the computing profession about what is ethical and what is not?

LOWE-TIGNOFF: Yes. *More* universal, but not universal. For example, most computer scientists would agree that computer systems should be easy to use, that they should be reliable and effective, and that they should not steal, lie, or cheat. My book examines computer systems, both real and hypothetical, and asks whether **these** systems behave in an ethical manner based on the general consensus within the computing profession as to what constitutes ethical behavior.

CLOSE-UP: Is there such a consensus within the computing profession?

LOWE-TIGNOFF: I assume that professional codes of ethics, such as the ACM Code of Ethics, represent the broad outlines of such a consensus. These proscribe lying, stealing, cheating, and causing harm. They emphasize that computer systems should be harmless to humans and to the environment. They emphasize the need for systems to be responsive to user needs and not to hurt users physically or psychologically. Thus, my point of departure for discussing ethical computer systems is that such systems should be at least as ethical as the computer professionals who develop them.

Programmers and developers need to adhere to codes of ethics, but we also need a basis for judging whether a computer system is itself behaving ethically. Is a computer stealing, is it damaging the environment, is it spreading rumors and lies, is it promoting unnecessary violence and destruction? When we look at a computer system from an ethical perspective, we must ask questions about its functionality, about its impact on users, society, and the environment. That is the gist of my book.

CLOSE-UP: These concerns seem obvious. Computers should not murder, steal, or lie. But then in reading your book, I realized that your characterization of stealing is more comprehensive than the usual one, and that it subsumes almost every other kind of unethical and immoral behavior.

LOWE-TIGNOFF: Yes, in order to give my discussion of ethics for computer systems greater coherence, I characterize all ethical lapses as acts of stealing.

CLOSE-UP: What constitutes stealing, then, in your framework?

LOWE-TIGNOFF: Stealing means to diminish oneself or another in some way, to diminish oneself or another spiritually, physically, materially, or psychologically. Another way of stating it is that to steal is to do harm to oneself or to another.

When you diminish another person, you always diminish yourself as well. When you diminish yourself, you diminish all others simultaneously.

CLOSE-UP: Suppose I have a headache and somehow you take that away from me. Is that stealing?

LOWE-TIGNOFF: No. Presumably you have given me permission to take the headache away.

CLOSE-UP: But what if I don't want you to take it away?

LOWE-TIGNOFF: Then, if you are an adult person, I should not take your headache away unless there is some other compelling reason to do so. Perhaps you are a pilot, and your having a headache endangers the lives of your passengers.

CLOSE-UP: Before we discuss your concept of stealing in greater depth, I would like to back up a bit and ask you how you got interested in developing ethical guidelines for computers.

LOWE-TIGNOFF: I started thinking along these lines after I read a pair of articles by a computer scientist named Roger Clarke. He tried to use Asimov's laws for robots as a means of deriving principles for the construction of ethical laws for information systems. He ended up with a set of commonsense principles, but I do not believe that Asimov's laws were an essential starting point. Let me read you the text of Asimov's laws of robotics as quoted by Roger Clarke:

First Law: *A robot may not injure a human being, or, through inaction, allow a human being to come to harm.*

Second Law: *A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.*

Third Law: *A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.*

Clarke goes on to discuss the inadequacy of these laws and a set of revised laws that were subsequently developed by Asimov. Finally, Clarke develops a set of general principles for information systems.

I enjoyed Clarke's articles, but I think it is better to start with ethical principles that govern humans as opposed to ethical principles specifically devised for robots and computer systems. As a first approximation, if we view a computer system as being capable of doing whatever a human being can do behaviorally, then certainly the computer system should be bound by the same ethical guidelines that guide the behavior of human beings. For example, since human beings are not allowed to steal, computer systems should not be allowed to do so.

Now, this is just a first approximation because computer systems can potentially do things that humans cannot do, and we need to discuss that possibility as well.

Ethical principles are constraints on the human capacity for harmful behavior. If we study codes of ethics, we see that they also encourage positive contributions to society. These codes are not just proscriptive.

I think that a computer system should be programmed to fulfill the same ethical obligations as a human being. It should not do harm. It should contribute to the good.

However, as my book indicates, there is a subtle danger lurking in computer systems that needs to be addressed.

CLOSE-UP: Namely, the possibility that a computer system can potentially steal human capabilities?

LOWE-TIGNOFF: Precisely. That is my main message. That is my reason for being, to communicate this message and to alert people to that danger.

CLOSE-UP: Before we get to that, let's discuss your general theory about ethics. Your book illuminates the relationship between modern ethical problems and the source materials that provide the spiritual, cultural, legal, and ethical foundations for what you call "traditional civilizations."

LOWE-TIGNOFF: I also call them "self-renewing civilizations." We are talking about Hinduism, Judaism, Buddhism, Christianity, and Islam. These are the civilizations that have survived for at least one millennium and some for as many as five millennia. Thus, they seem to be self-renewing. Some people say that Hinduism has been around for five millennia. Judaism, Christianity, and Islam can trace their origins back to Abraham, who lived almost four thousand years ago. Buddhism is twenty-five hundred years old. By contrast, Nazi Germany lasted twelve years, and the Soviet Union lasted about seventy years.

The basic ethical precepts in all these traditional civilizations are proscriptions against stealing in various forms. The fundamental moral precepts are these:

1. Do not steal.
2. Do not murder.
3. Do not lie.
4. Do not use sexuality to hurt others.
5. Do not destroy your wisdom (with drugs, destructive emotions, or intoxicants).

CLOSE-UP: In your book, you show that these principles are actually proscriptions against stealing.

LOWE-TIGNOFF: Yes. I took the Ten Commandments and showed how each commandment could be viewed as an attempt to prevent people from diminishing themselves and from diminishing some "other" or "others."

CLOSE-UP: So, for example, if you worship idols, you diminish yourself, or you steal from yourself.

LOWE-TIGNOFF: Correct. You also steal from God.

CLOSE-UP: Could you explain the five precepts in terms of the concept of stealing?

LOWE-TIGNOFF: "Do not steal" is a proscription against taking the property of another. Of course, when you steal someone's property, you are diminishing their wealth and perhaps their happiness, but you are also diminishing yourself. You are a thief.

Stealing property is the concrete manifestation of stealing as a cosmic idea. In other words, conventional stealing is a symbol for and a reminder of the other forms of stealing.

"Do not murder" is a proscription against stealing away the life of another. The universe has given that person a life, but you decide you don't like the way the universe has arranged things, so you murder that person to try to get the universe to bend to your will. But you cannot get the universe to bend to your will. Instead, your act of deliberate evil forces the universe to compensate for your action. The universe is all-aware, alive, and intelligent.

In committing a murder, you diminish the other and you diminish yourself. You diminish yourself because you will have to suffer the consequences of your act at some level. You are a murderer, and you know it at all of the various levels of your being, conscious and unconscious.

CLOSE-UP: What about the other precepts?

LOWE-TIGNOFF: When you lie, you steal the truth away from the person who listens to your lies, unless they have the wisdom not to believe you. When you bear false witness against your neighbor, you are stealing away his or her integrity, dignity, and honor. Lying diminishes those who believe the lies by rendering them less knowledgeable. In lying, one diminishes oneself because one lives in an environment of falsehood. Furthermore, one is a liar, and one knows that one is a liar at all of the various levels of one's being, conscious and unconscious.

To gossip is also a sin, because you are stealing from a person's integrity, dignity, and honor. In other words, you are diminishing that person's reputation and the affection due to that person. Of course, the gossip also reduces his or her own stature on many levels.

The proscription against hurting others using sexuality is also about stealing. For example, if a man commits adultery with another man's wife, then he is stealing from his own wife and from that other man. He is also stealing from himself and from the woman he slept with. If you look at it in terms of diminishment, those who committed the act and all of those who are influenced by the act, especially the spouses and the children, are all diminished. At the very least, their happiness and peace of mind are diminished.

The proscription against destroying one's wisdom usually takes the form of a proscription against the use of alcohol, drugs, and intoxicants. However, anger, hatred, excessive fear, and anxiety also destroy wisdom. Clearly, if one destroys one's wisdom, then one has greatly diminished oneself. One's stature as a human being is much less. In addition, by destroying one's own wisdom, one contributes much less to the world, and thus the whole world is diminished.

Hurting the brain with excessive alcohol and drugs is stealing away from oneself in a dramatic and sometimes irreversible manner. The alcoholic not only steals from himself, the alcoholic steals from his family and friends. He steals away their peace of mind and, if he is violent, their physical safety. He can also steal away the prosperity his family might have enjoyed if he had been able to keep gainful employment. The alcoholic or other drug-dependent person diminishes himself and everyone in his environment.

What I am describing is the actual spiritual state of things. By thoughtless and unethical behavior, we greatly diminish ourselves and everyone and everything around us.

CLOSE-UP: I appreciate your perspective that when we act unethically or immorally, we are diminishing ourselves and others.

LOWE-TIGNOFF: Thank you.

CLOSE-UP: Now, taking these ideas and applying them to computers, what you are saying is that computers and computer technology have the potential to diminish human beings, to reduce our stature spiritually and psychologically.

LOWE-TIGNOFF: Absolutely. That is what I wrote, and that is what I believe.

CLOSE-UP: This is an important assertion, so how can you convince us of it?

LOWE-TIGNOFF: Let's analyze the impact that a computer system has when it is placed in its intended environment. Is that computer system stealing in any way? If it is stealing, is it contributing something in return that compensates for the stealing? In other words, we must demand that if a computer system steals, then it must more than compensate for its stealing by adding something of value to our lives and society.

CLOSE-UP: In other words, it must contribute more than it takes away?

LOWE-TIGNOFF: Yes.

CLOSE-UP: For example, you wrote that a computer system with a poor user interface is stealing from the user.

LOW-TIGNOFF: Yes, it could steal in a variety of ways. It could steal the user's eyesight, or it could convert an enjoyable job into a hellish one. A computer with a poor user interface can diminish the user's happiness and physical health.

If a computer system is not reliable, it can steal the life of its operator, as in the case of the killer robot incident. A database system can steal away truth if it does not maintain the integrity of its data.

CLOSE-UP: But your main thrust is that computers might diminish human capacity, in effect, stealing from the human race.

LOWE-TIGNOFF: I am concerned about all of the issues that we have been discussing on this program. I believe we should carefully analyze each computer system from the point of view of how it can steal from users, clients, and what has been called the penumbra of that computer system.

CLOSE-UP: What is the penumbra of a computer system?

LOWE-TIGNOFF: This term was used by Robert Collins and his coauthors in their article on ethical guidelines for computer systems. It refers to all people indirectly influenced by the computer system. You see, a computer system can have a large penumbra, and thus it may diminish the quality of life for many, many people.

CLOSE-UP: Or it may be making their lives better in some way.

LOWE-TIGNOFF: For example, the penumbra for the killer robot includes Mrs. Matthews, the widow of the robot operator, her children, and other relatives. It includes the lawyers who have gotten involved in the case and, most broadly, the public that has gotten so absorbed in the issues raised by this incident.

If computer scientists begin to develop computer systems that steal fundamental human capabilities—such as medical judgments, ethical judgments, artistic and scientific creativity—then the penumbra will include all of humanity. There is a danger that we will all be greatly diminished by such a technological development. It will be a threat to our very humanity.

CLOSE-UP: These systems might augment human capabilities and human stature.

LOWE-TIGNOFF: This is what we need to look at. We shouldn't go into this new era with our heads buried in the sand.

'Here are some hypothetical examples that I used in my book to illustrate this point.

Suppose a computer system is designed to make decisions in criminal trials. The system is motivated by a lack of faith in the jury system. The system requires that trained professionals take the facts in the case, prepared by lawyers, and enter them into a knowledge base. The computer system then reaches a verdict in the case and decides on a sentence. The judge is reduced to a mere administrator of data. One of the judge's responsibilities would be to read the computer results to the defendant and his lawyer.

The researchers who developed the system might test their product using the following experimental design. They choose a set of actual court cases that were decided by a traditional jury and judge. They then ask the computer system to arrive at a verdict and a sentence for each of these cases. The researchers then ask a population of legal scholars, in a blind test, to study the court cases and to indicate which gave the better judgment: the actual judge and jury in the case or the computer system. Suppose that the

researchers discover that the population of legal scholars favor the computer system over the actual judge and jury eighty percent of the time.

I think a scenario like this is very possible in the next fifty to one hundred years, if not sooner. If we allow computer systems to try people and to pass sentence, are we allowing the computer system to steal from us? I would also like to know what it is that the computer is stealing and whether we, as human beings, can afford to allow computers to take that ability away from us.

The hypothetical computer system that I am discussing will be taking away our right to be judged by a jury of our peers.

CLOSE-UP: Someone might respond, But the peers are in the code. That is, the peers are the ones who wrote the program.

LOWE-TIGNOFF: But those are not representative peers.

CLOSE-UP: But what if the computer system is actually better than a human jury? Shouldn't we use the computer system?

LOWE-TIGNOFF: Would you want to be judged by a computer program if you were on trial for a crime you did not commit?

CLOSE-UP: Yes, if the computer program could do a better job than a human jury. Look at the jury pools that we have seen in some recent celebrity trials.

LOWE-TIGNOFF: So we should allow computers to serve as judge and jury because our justice system has so many serious deficiencies? Is that it?

CLOSE-UP: Perhaps. I really need some time to think this over.

LOWE-TIGNOFF: Your willingness to give up that human prerogative to a computer is what I call stealing. The computer would be stealing away a human function that has been central to our humanity for millennia.

CLOSE-UP: Is it stealing if we willingly, as a society, turn over this function to expert systems that are better than the "man on the street"? Stealing involves taking something from a person without that person's permission.

LOWE-TIGNOFF: Is it stealing if you voluntarily give something over to a con artist because you don't understand the value of what you are surrendering?

CLOSE-UP: You're implying that I am not seeing the full magnitude of what we human beings would be losing if we hand certain functions over to the computer.

LOWE-TIGNOFF: Yes. Are we not losing something by having no human element, no human involvement, in the judgment process?

CLOSE-UP: But there is some human involvement. Human beings programmed the computer.

LOWE-TIGNOFF: In the Bible, Solomon displays his wisdom, in part, by adjudicating difficult cases. His brilliance is demonstrated by extraordinarily intuitive but effective leaps of judgment. This is part of our humanity.

CLOSE-UP: I see. A figure like Solomon could not be recognized or acknowledged in the brave new world of the computerized court room. A fundamental human capacity would be reduced . . .

LOWE-TIGNOFF: Would atrophy. It would atrophy from lack of use. In other words, Human Being, with a capital H and a capital B, would be diminished.

Is there something fundamentally human about the wisdom that is required to pass judgment on another human being? Isn't this one of the nexus points where justice, compassion, and mercy meet, are developed, and are exercised? If this nexus point is removed, will justice, compassion, and mercy still have an opportunity to develop and manifest among humans?

Judges have been around since Biblical times. In all just societies, the responsibility of the judge has been based on a keen sensitivity to the rights of the accused and the presumption of innocence until guilt is proven. This is one nexus point where society expresses its capacity for justice, compassion, mercy, and empathy. Can we afford to say "Oh, this is too messy. Let the computer do it"?

Now, what if this happens in medicine and other professional areas as well? More and more, tough decisions are handed over to computer systems. Decisions about diagnosing illness, prescribing drugs. The doctor becomes a mere adjunct to the computer, which has enormous databases filled with knowledge. Engineers, computer scientists, accountants, and so forth all defer decision making more and more to computer systems. Is something being stolen? Is Human Being diminished?

CLOSE-UP: I see. A certain kind of thought process, involving risky thinking, involving tough decisions, is moved more and more onto the computer systems. Then what is the brain being used for? Who needs a brain one hundred or two hundred years from now?

LOWE-TIGNOFF: It's not just a thought process. It's that nexus of responsibility for our actions that makes us ultimately human. That nexus implies that we are conscious beings, that we are not automata. Shouldn't we think about all this *now* before it becomes a *fait accompli*?

Here's another possibility: A computer system that does our moral and ethical reasoning for us. Thinking ethically is too difficult-let's reduce it to a bunch of question sets, rules, and guidelines, and let the computer grind it out. Then we are free of all of those messy, difficult decisions. What if some computer system becomes the dominant means of ethical decision making?

CLOSE-UP: That reminds me of the infamous "killer robot tapes" that were played on the radio a few months back. Maybe Randy Samuels should have used such a system before deciding to tell George Cuzzins what Zelda Riddle-Davis told him about her being a lousy programmer. I face decisions like that all the time, and I wouldn't mind having the help of a computer.

LOWE-TIGNOFF: But what will the computer be stealing if this happens? Isn't there a sense in which your being will be diminished? Isn't there some human value in wrestling with ethical dilemmas and coming to a courageous decision when courage is required? Isn't the ability to decide between right and wrong a fundamental part of human nature? Can we afford to have that stolen from us?

CLOSE-UP: But couldn't human wisdom just be shifting from one domain to another? Solomon displayed his wisdom by judging disputes, but human beings in the future could display their wisdom by judging other issues, for example, issues involving the use of technology.

LOWE-TIGNOFF: I do not think so.

CLOSE-UP: Maybe the future will be a paradise. We will spend our years writing, composing, and doing those things that are challenging and soulful.

LOWE-TIGNOFF: Well, let's find out. Consider this hypothetical situation: A computer system is developed to compose newspapers automatically from compilations of raw facts and data. This is necessary because competing sources of information are all on-line and newspapers must be produced with incredible speed in order to provide relevant information. Human story writers are not up to the challenge, so a computer system now writes and composes newspapers. Reporters are laid off by the thousands. What do you think of that? You're a reporter.

CLOSE-UP: That damn computer just stole *my* job!

LOWE-TIGNOFF: Is that all?

CLOSE-UP: It stole away my passion-I am passionate about writing. Now I have no outlet for my passion.

LOWE-TIGNOFF: Who cares? People dislike the press, so they are happy to see you get yours.

CLOSE-UP: This system is stealing away something from the newspaper, something human, and it's stealing from the readers. They're not going to get that passion.

LOWE-TIGNOFF: Oh, but this computer system can add as much passion as it needs to. On a scale of 0 through 10, how much passion do we want in today's paper? Okay, put that into the machine, and let the machine grind it out.

Now consider the following hypothetical scenario: A computer system is developed to compose and perform music of a particular genre. Its compositions prove demonstrably better than human compositions in that people enjoy the computer-generated music much more. All musicians who compose and perform music in this genre see themselves being pushed aside by the public's enthusiasm for the computer-generated music. Is something being stolen here?

CLOSE-UP: Musicians are losing their jobs, and musicians are passionate about what they are doing. Does this music have feeling?

LOWE-TIGNOFF: This is the future, and computers have become incredibly adept at capturing feeling. The audiences love it. This new music is creative and brings in new elements the human composers had not imagined.

CLOSE-UP: Maybe this is just a better way to compose music. The system must reflect the true musical genius of at least one person.

LOWE-TIGNOFF: But what about the human musicians who are out of a job? No one wants to listen to their music any longer. Have we stolen something from them? If we say “that is progress” as group after group is dispossessed by computer systems, where is this heading?

CLOSE-UP: Artistic capabilities might atrophy as well as technical capabilities.

LOWE-TIGNOFF: So, what is left-computer programming? All human creativity is moved over to computer programming, but then the computer is better than the human even when it comes to computer programming.

In the end, the computers will do nearly everything that is intellectually challenging. Human beings will end up doing things that are too boring for the computer, like hauling manure. Humans will be mere slaves to an all-encompassing network of intelligent computers that will intrude into every aspect of human life. We’ll be slaves, like the Israelites in ancient Egypt. The computer will be new Pharaoh.

CLOSE-UP: So, inevitably, we human beings will have to impose limitations on our computer systems,

LOWE-TIGNOFF: You’re assuming that the man on the street will have the intelligence to see what is happening.

Let me put it bluntly. Inevitably and inexorably, computer technology will become as potentially dangerous as genetic engineering is today. Just as genetic engineering has the potential to introduce something into the environment that is devastating, so computer technology can also threaten human self-confidence, wisdom, and sovereignty.

CLOSE-UP: You said that the computer has the potential to become an intoxicant, like alcohol or drugs.

LOWE-TIGNOFF: Yes. That is already happening. Consider virtual reality entertainment. As the earth is despoiled, as the oceans become ever more polluted, as the air quality diminishes, as species disappear, as nuclear waste and garbage accumulate, as the population explodes and ecological systems are choked by human wastes and garbage, as free spaces disappear, as more and more of the natural world is pushed aside to make room for development and progress, as all these things happen and manifest, people can always resort to their virtual reality helmets and data gloves in order to experience mountains, oceans, rivers, the skies, the long-gone natural order.

You see, this is just the purpose of drugs and intoxicants. Because actual reality is becoming so bleak, we will resort to an intoxicant that will allow us to escape from that reality. So now we are creating a technology that will allow us to smell beautiful flowers and see beautiful scenes, even as we are up to our eyeballs in our own garbage.

CLOSE-UP: So what is being stolen here?

LOWE-TIGNOFF: In this case, our ability to process the feedback that the universe is trying to give us concerning the stupidity of our behavior. Whenever another precious species disappears, whenever another beach is closed because of pollution or oil washing ashore, whenever the earth appears diminished in its beauty and its richness, that is feedback from the universe that we are not being good stewards of the earth.

CLOSE-UP: Virtual reality can shut out the social reality as well: the poverty, the violence, the inequity, the despair.

LOWE-TIGNOFF: Of course.

CLOSE-UP: Do you think we need laws to govern technological development?

LOWE-TIGNOFF: I think laws will be necessary for those sectors where jobs are being lost and human capabilities are being stolen. Our civilization needs to consider whether we want our newspapers to be written by computers, whether we want computers as judge and jury, whether we want computers to diagnose and treat our illnesses, to make ethical and moral decisions for us, to write our novels, to educate our children, to compose and play our music, and to create our motion pictures.

CLOSE-UP: Are you pursuing these issues further? Is there another book that we can look forward to?

LOWE-TIGNOFF: I am working on a new book, or maybe several books. I have some chapters written out. This new book is not so much about computers, but the whole idea of how we diminish ourselves and others. My computer book emphasized stealing as a primordial ethical lapse. By the end of that book, I started to focus on this idea of how we diminish ourselves and others. The “other” could be God, my neighbor, a tree, an owl, the ocean, this entire planet. Whenever I do not admit the profound greatness of the other-of God, of my neighbor, of that tree in front of my house, of that majestic owl looking down from that tree, of the ocean, of this beautiful earth-then I diminish myself. I diminish who I am. You see, the other is not “other” at all. When we diminish another, we are diminishing ourselves. When we diminish ourselves, we diminish all others simultaneously.

CLOSE-UP: Can’t you state it more positively? Shouldn’t our effort be to augment, increase, or glorify everyone else and ourselves? I can’t find the exact word I need here.

LOWE-TIGNOFF: I would use the word *redeem*. We would like to redeem ourselves and the world, to make the system whole and harmonious again. The,

danger is that people who think that they know how to redeem the world often become too aggressive, too imperialistic. It is a very delicate matter to think that you know what is good for another person. Thus, I emphasize respecting others—not diminishing them; not lying, stealing, and cheating; not doing things that will hurt other people. Perhaps if we remove the hurtfulness, redemption will flower by itself.

CLOSE-UP: One difficulty I had with your book was that you seem to insist that the computers are doing the stealing, when in fact it seems more accurate to describe the situation as our surrendering our talents to the computer. Who is really doing the stealing?

LOWE-TIGNOFF: My answer to your question is that we are stealing from ourselves. We have free will. If we give something away that is rightfully ours, we are stealing from ourselves.

CLOSE-UP: In your book you wrote that you were afraid that human beings were surrendering their spiritual power. What is spiritual power?

LOWE-TIGNOFF: I think that it is the power to redeem, heal, and repair.

CLOSE-UP: Did it ever occur to you that what you call “stealing” is just the opposite of love? Diminishing oneself, diminishing others, all of this is just the opposite of what love would do. Love heals, replenishes, and redeems that which has been diminished. Do you think that we are allowing computers to steal from us because we do not know how to love one another?

LOWE-TIGNOFF: That could explain what is going on.

CLOSE-UP: That’s it for this evening’s program. If you have any comments about this evening’s program, please send them to close-up@naptv.org. I would like to thank Professor Lowe-Tignoff for offering his perspective on where computer technology is heading.

An Interview with Professor Jacob Lowe-Tignoff

Silicon Valley University’s
CANDID PROFESSOR

ANNOUNCER: Those are the famous chimes of the Silicon Valley University Clarion Tower. It must be time for another lecture in our “Candid Professor” series. This week’s lecture will be given by Professor Harry Yoder. The

course: Computer Science 412-Ethical Issues in Computing. As you can see, the students are anxiously awaiting the illustrious Professor Yoder, who has written numerous books and articles about the social impact of computing and computer ethics. He is the Samuel Southerland Professor of Computer Technology and Ethics at the university.

As regular viewers of “Candid Professor” are certainly aware, Professor Yoder has no idea that he is going to be on television this morning. Of course, as soon as he sees our lights and cameras, he’ll know that he’s the “Professor of the Week” for this venerable Silicon Valley University tradition.

YODER: Oh, no! Of all weeks, why did you have to choose this week?

ANNOUNCER: We’re taping this even as you speak, Professor. You know the rules. Harry Yoder, you’re the “Candid Professor’ Professor of the Week.”

YODER: How can I lecture with all of those lights in my face? I can’t see my students. That’s better, thank you.

Class, I have your exams, and I will return them to you at the end of the hour.

CLASS: Can’t we get them back now?

YODER: Look, I cannot lecture when you guys are having evil thoughts about me.

Now, today’s assignment was to read and discuss the paper “How Good Is Good Enough?” by Collins, Miller, Spielman, and Wherry published in the *Communications of the ACM* in January 1994. I hope you guys did a good job of reading that paper because, as you know, I was away in Hawaii last week at an important computer ethics conference. On top of that, I had to grade your midterms. On top of that, I rushed to class straight from the airport.

CLASS: Poor Professor Yoder!

YODER: Okay, let’s cut out the sarcasm. It was not all fun and games as you imagine. This was a computer ethics conference, not OOPSLA. So who can give me a brief description of the paper that you were asked to read?

CLARISSA: It’s about developing a framework for deciding whether it is ethical to release a software system.

YODER: An ethical framework as opposed to what?

CLARISSA: As opposed to a framework that is based just on technical issues, on software quality assurance.

YODER: And why is that important? Someone else.

JAMIE: Because the science of software quality assurance is not far enough along in terms of its reliability to allow those techniques to be used to decide whether software gets out the door.

YODER: Can you give some specific examples?

JAMIE: Software testing cannot find all bugs. Some bugs are related to the way in which users interact with the software. They may escape detection during normal software testing.

YODER: Do you buy this idea that technical criteria alone may not be sufficient to decide whether software is ready for release?

CLASS: Yes.

YODER: Any dissenters?

ABIGAIL: I think that the ethical criteria are just as primitive as the technical. The authors themselves state this at the end of their paper. In other words, I think these ideas need to be discussed over a period of time. Perhaps after ten years of intense deliberations we will arrive at a good set of ethical criteria.

NATE: I agree with Abigail. But it may not take ten years. It could take longer. I believe that ethical criteria may be as difficult to develop as technical criteria.

YODER: Let's discuss the criteria developed by Collins and his coauthors. First of all, what were the basic principles behind the development of their ethical guidelines?

NATE: The fundamental idea is that of a social contract, as found in the writings of Rousseau and Locke. For example, the proper behavior for the government and its citizens is based upon a contract between the government and its citizens. For Americans, the Constitution embodies that contract.

JAMIE: Perhaps we need a standard document like the U.S. Constitution that would explain the duties and obligations of the various parties involved in a software project.

ABIGAIL: That's why this paper is so important. I think it is leading us in that direction. I didn't take it as the final word. I think a document for computing professionals based on the notion of a social contract would go a long way toward outlining the duties and obligations of the various parties.

CLARISSA: Yes, the contract concept seems to be cropping up more and more in discussions of software and software engineering. In my software engineering class with Professor Silber, we discussed Bertrand Meyer's notion that classes should be designed around the concept of a "contract" that lists the duties of the various pieces of software as well as the benefits accrued by the use of that software. I would assert, therefore, that the contract idea is fundamental in computing because of the nature of computing and because of its complexity. The contract concept is inherent in the need to divide up responsibilities in the construction of a software system. It may be that software itself will need to be organized using the contract concept, just as society has been organized around the idea of a social contract.

YODER: Anyone else?

NATE: I am also taking software engineering this semester. I like what Clarissa said. Obligations and benefits are the building blocks of contracts, and it

is interesting that these concepts apply at various levels of software construction. Turning back to the ethical and social implications of software systems, I think the message is that we have to carefully consider the parties involved, their duties, and the benefits that they can expect. Benefits are just as important as duties.

YODER: Okay, we will discuss this issue of benefits and duties a bit later. But I would like to return to the framework that Collins and his coauthors developed. Can someone give a summary of how that framework was constructed?

JAMIE: The authors relied heavily on an ethical theory of John Rawls that stressed negotiations between interested parties leading to a contract that would list the obligations of the various parties. The benefits are implicit, although I think they should be explicit.

CLARISSA: There was more to the ethical theory than the idea of contracts. The authors listed fairness, rationality, and a certain kind of ignorance as important for the negotiation of the contractual obligations. The parties in the negotiation would be fair, rational, and also ignorant of their own possible stake in the negotiation. Thus, the parties would naturally try to protect the most vulnerable party, since they might be the most vulnerable party*

YODER: So the people in the negotiation do not know their stake in the negotiation? They just know the kinds of parties involved, and they make an effort to protect the rights of all parties.

CLARISSA: Yes. Each negotiator could actually be the "least advantaged," or most vulnerable, party.

YODER: What parties did the authors see as having a role in the negotiation?

HANK: Software providers, such as Silicon Techchronics. Software buyers, such as CyberWidgets. Software users, such as Bart Matthews. And the software penumbra, such as Bart's wife and his children.

YODER: Could you define the *software penumbra*?

HANK: This is the population that lives in the "shadow" of the system, that can derive benefit from or can be harmed by the system. For example, fly-by-wire software would include in its penumbra all the passengers who fly in airplanes controlled by fly-by-wire software and all the friends, relatives, and acquaintances of these passengers. The penumbra can be quite broad. The fly-by-wire example is actually used in the article we read.

ABIGAIL: When the United States and the former Soviet Union were considering computer-mediated launch-on-warning systems, the penumbra would be the entire human race.

YODER: Can anyone explain what *launch-on-warning* means?

ABIGAIL: It means that a computer would decide to launch a nuclear strike in response to a perceived attack from the other side.

YODER: So a system like that would have an enormously broad penumbra. It's a scandal that the human race ever got to that stage of thinking during the Cold War. What came next in the "How Good Is Good Enough?" paper?

SAM: We have four parties—software providers, buyers, users, and the penumbra—and a negotiation that is going to decide their obligations to one another under a contract. Those involved in the negotiation are fair and rational, and they do not know their own stakes in the negotiation. They may not have any stakes in the negotiation, but they understand the nature of software development, and they understand the roles played by providers, buyers, users, and those in the penumbra. They will naturally act to protect the least advantaged or most vulnerable.

JAMIE: In fact, the authors came up with three principles to guide the negotiators. These are, Don't increase harm to the least advantaged, the most vulnerable. Don't risk increasing harm in an already risky environment. From these they derive the corollary that, and I'm quoting here from page 86, "Software designed for a low-threat context should not be used in a higher-threat context." Finally, they propose a publicity test for difficult cost-benefit trade-offs. For example, and this is my own example, let's say a car company knows that a gas tank explodes when there is a rear-end collision. They can do a cost-benefit trade-off on whether to redesign the car. They might decide that it is worth absorbing ten deaths per year and the lawsuits that would ensue instead of absorbing the costs of redesigning the car. If the public would be outraged if it found out about such a trade-off, then the trade-off would be judged as being unethical according to the publicity test.

YODER: Can anyone give us a summary of the contractual obligations that the authors developed in their paper?

NATE: Each party has obligations to each of the other parties. Providers have obligations to buyers, to users, and to the penumbra.

CLARISSA: And to themselves! They must make a profit!

YODER: But can we have a summary of these obligations so that we can try to apply them to the case of the killer robot?

JAMIE: I made up slides that contain the content of their tables 1 through 4. These might be helpful for those who will be watching this on television.

YODER: You knew this would be on television?

JAMIE: I work at KPAT. You know it's affiliated with the university, and I'm a work-study student.

YODER: And you didn't warn me?

JAMIE: It's strictly against the rules. "Candid Professor" is a sacred university tradition.

YODER: Okay. Let's put up your slides.

JAMIE: They're on disk. It will take a few seconds to set up the computer and the projector. These are verbatim reproductions of tables 1 through 4 of the "How Good Is Good Enough?" paper from pages 88 and 89.

TABLE 1. OBLIGATIONS OF THE SOFTWARE PROVIDER

To the provider:

- make a *reasonable* profit

To the buyer:

- *reasonable* use warranty
- *informed consent* about testing process and potential shortcomings

To the user:

- *clear* operating instructions
- *reasonable* protections from and informative responses to use and abuse
- provide *reasonable* technical support

To the penumbra:

- reasonable protections against physical, emotional, and economic harm from applications
- open about software development process and limits of correctness

(from Collins et al., p. 88)

YODER: Excellent! Now let's take a few minutes to study the slide, and then I'd like to ask you whether these ideas apply to the killer robot case.

Okay. How would you apply these obligations to the software provider in the case of the killer robot? Clarissa?

CLARISSA: Clearly, Silicon Techchronics was not forthright in its communications to CyberWidgets about software testing techniques.

NATE: They failed miserably in their obligations to the user. The robot operators did not receive adequate training. They certainly did not protect the robot operators against potential harm that could be done by the robot.

HANK: It's not clear to me that they had any obligation to communicate to the penumbra about risks involved in the use of the robot. It seems like the

penumbra, including Mrs. Matthews and her children, get involved after the fact.

YODER: Is this always true? Does the penumbra always come in after the accident, after the disaster?

ABIGAIL: No, I think that the authors are making an excellent point. Some systems have a more intimate connection to their penumbras. For example, the entire population was in the penumbra for the proposed launch-on-warning systems in the 1980s. The people should have a say in whether such a system gets deployed or not.

NATE: I wonder how companies like Silicon Techchronics can fulfill the obligation to be “open about software development process.” If they had been open about what was going on, would that have dampened the negative public reaction to the killer robot?

JOHN: That’s why I think this paper is heading in the right direction. The public does not understand how software is developed, and they do not understand the limitations of computer systems, but they need to.

YODER: Okay. Let’s study Jamie’s next slide, and then we’ll discuss it.

TABLE 2. OBLIGATIONS OF THE SOFTWARE BUYER

To the provider:

- negotiate *in good faith*, recognizing the importance of provider’s fair profit
- learn *enough* about the software to make *an informed* decision
- facilitate *adequate* communication to users.

To the buyer:

To the user:

- provide *quality software appropriate* to users’ needs within *reasonable* budget constraints
- *prudent* introduction to automation
- *informed* consent to using software
- *represent* user’s interests with providers

To the penumbra:

- buy software only with *reasonable* safeguards for the public
- *open* about software capabilities and limitations

(from Collins et al., p. 88)

YODER: The software buyer was CyberWidgets. Did they fulfill these obligations?

MIKE: I think they were deficient in the way in which they communicated with Silicon Techchronics. Nothing that I have seen on TV or read in the papers so far indicates that they really tried to learn all they could about the Robbie CX30 project from the point of view of safety and so forth.

HANK: **There** didn’t seem to be any effort on the part of CyberWidgets to facilitate communications between Silicon Techchronics and the robot operators. That would have solved a lot of problems.

ABIGAIL: **But** shouldn’t Silicon Techchronics have been more aggressive in that way? This seems to be lacking in the authors’ list of obligations by the provider. I think the provider must communicate with users to determine their needs *early* in the project, before the thing gets built.

CLARISSA: **Good point.**

JAMIE: This is really interesting. I think we’re discovering significant deficiencies in the way that CyberWidgets behaved. The media has concentrated **its attention** on Silicon Techchronics. CyberWidgets did not fulfill its obligations to users based on this slide. They did not check to see that the new system would be appropriate to the needs of the robot operators. They did not really ensure a quality product, which was their obligation to their employees, the users: I don’t see that CyberWidgets really consulted with users or represented their interests to Silicon Techchronics.

MIKE: I think that CyberWidgets failed in its obligations to the software penumbra. The obvious software penumbra includes the family of Bart Matthews, but **it** includes all of the families of the robot operators. I don’t think CyberWidgets had any awareness of a penumbra.

YODER: Okay. Let’s look at the next slide.

TABLE 3. OBLIGATIONS OF THE SOFTWARE USER

To the provider:

- respect ownership *rights*

To the buyer:

- *active communication* with the buyer
- *good faith effort* to learn and use software responsibly
- *reasonable* requests for computing power

To other users:

- *willing* cooperation in learning and using software

To the penumbra:

- *conscientious* effort to reduce any risks to the public
- encourage *reasonable* expectations about software capabilities and limitations

(from Collins et al., p. 89)

YODER: What do you think? Did the robot operators fulfill their obligations?

CLARISSA: I don't see how they could be faulted.

JAMIE: I think they can be faulted. First of all, if you are a husband and a father with three children, isn't it your responsibility to make sure that the robot you are using is safe? That's your obligation to the penumbra.

CLARISSA: Okay, but we see no evidence that Bart Matthews was irresponsible in the way that he used the software. He did communicate his concern to CyberWidgets that the robot console had a poor user interface. I think he was willing to learn, although he didn't know everything he needed to know in order to save his own life.

NATE: I can see how in certain situations the obligations of the user to the penumbra are very significant. For example, consider the operator of an X-ray machine, such as the one that we studied [the Therac-25]. The operators really have the obligation to inform patients of the nature of the machine and of possible risks.

MIKE: This slide shows that users have more power than we usually give them credit for. They need to understand the limitations of computer systems so that they can make reasonable requests for new systems. I wonder if the workers at CyberWidgets were pushing for the new robots and whether that was a factor in the haste with which the robots were introduced.

YODER: Let's consider the final slide.

TABLE 4. OBLIGATIONS OF THE SOFTWARE PENUMBRA

To the provider:

- become *aware* of the capabilities and limitations of software.
- advocate a *suitable* economic and statutory environment for quality software

To the buyer:

- advocate a *suitable* economic and statutory environment for quality software

To . . . [the] users:

- expect only *reasonable* service from users
- become *aware* of the capabilities and limitations of software

To the penumbra:

- become *aware* of the capabilities and limitations of software
- advocate a *suitable* economic and statutory environment for quality software

(from Collins et al., p. 89)

YODER: Okay. What about the obligations of the penumbra?

CLARISSA: These seem to be rules about good citizenship in the computer age.

YODER: But what about the killer robot case?

MIKE: Mrs. Matthews is the most obvious member of the penumbra. She is advocating a suitable statutory environment in that she wants to see accountability in the death of her husband. She is not going to be satisfied with liability damages against Silicon Techchronics. She wants those who caused the accident to suffer appropriate legal penalties.

YODER: Do you support that kind of accountability?

MIKE: No. I would find that kind of legal environment very intimidating. It might get hard to get anything done.

ABIGAIL: Jane **McMurdock**, the prosecutor, is also part of the penumbra, and she is working for a suitable statutory environment for software development in her own way.

YODER: Is there a broader penumbra in the killer robot case?

JAMIE: Yes. It includes everyone who has read about it, everyone who has seen a television program devoted to it, and everyone who has listened to a radio talk show about it. We are all in the penumbra because this issue is being discussed in the media. I think our obligation is to learn as much as possible so that we can understand the issues.

YODER: Is there anything beyond that?

NATE: Yes, I think there's a missing element in the authors' analysis, although I think it is an excellent paper. It almost seems like a top-down analysis. The apex is the obligation of the provider to make reasonable profits. I would like to see a new obligation added to the penumbra's obligations to itself, and that would be to increase the well-being and happiness of the penumbra. I think the obligation of the penumbra to itself is to protect its

own happiness, health, and well-being. This may place providers at odds with the penumbra. It may require new obligations in all four tables that were given in the paper.

YODER: I think you have a good idea there. So, we have a top-down analysis where business leaders are trying to maximize their profits and a bottom-up analysis where the penumbra is always asserting its right to protect its well-being and happiness.

NATE: Yes.

YODER: Class, that was a good discussion. Please pick up your exams on the way out. And Jamie, don't forget your diskette!

Is that thing turned off?

I can't believe you people chose this week to put me on "Candid Professor." You chose the worst possible week. I'm just returning from the airport. I didn't even have a chance to get home yet.

ANNOUNCER: It was a good class.

YODER: I usually give more of a lecture. This was more like a discussion.

ANNOUNCER: Your students were really on top of the material. They made lots of good points.

YODER: Jamie, you really saved the day with your slides.

JAMIE: Well, this is the second course I've taken from you, and I remember what happened last fall when you came back from a computer ethics conference in the Virgin Islands. Remember when you came to class right from the airport, and you opened your briefcase and your bathing suit fell out?

YODER: How can I forget? Well, at least it was my bathing suit.

Transcript of the Television Program

APPENDICES

A	Real People and Institutions	xv
B	Endnotes and References	xix
C	Discussion Questions	xii