

portable anti-aircraft system. The system which the contractor developed destroys aircraft effectively, but it also occasionally kills the person who launched the missile. The company has declared that this is not a problem because it "[is] in full compliance with the specifications given to [it] by the user." Being a professional involves using one's special skills to give careful and constant consideration to the impact of one's service on others. This consideration is guided by a set of ethical principles.

We have mistakenly understood computer ethics as different from other professional ethics. When we look at medical ethics, legal ethics, journalistic ethics, we have to distinguish the practitioners of those ethics from the ethical principles they affirm. The three professionals work in different contexts: medicine, law, and journalism. However, when we talk of their professional ethics we do not consider them three different kinds. The distinguishing characteristic among professional ethics is the context in which they are applied. Because there are three contexts, it does not follow that there are three distinct sets of ethical rules or three different kinds of moral reasoning. Nor does it follow that computer ethics is another unique set of ethical principles which is yet to be discovered.

By analogy with other "ethics," "computer ethics" can be divided into two spheres. The first embraces a set of ethical problems that can be reasoned about by analogy with most other traditional ethical abuses such as fraud, theft, and trespassing. But this should not even be called *computer* ethics. Ethics for computing professionals is not just another kind of ethics, but it is ethical values, rules, and judgments applied in a computing context based on professional standards and a concern for the user of the computing artifact. It is this sense of computer ethics that has received very little attention. Most of the attention has been directed at the results of the failures of professional ethics or abuses involving a computer. . . .

Computer ethics, as presented here, is modeled on other professional ethics. It can use moral-reasoning models that are similar to those in other professional ethics. The theory of computer ethics we have presented does not rule out the examination of such critical concerns as the impact of technology on the nature of work or computer fraud. The theory puts these concerns in other ethical categories. The former is a concern of sociology, and the latter is a concern of property rights.

Computer practitioners do not have a single representative organization that can control membership in the profession; there is no organization to impose sanctions for violations of professional behavior. But the absence of a single organization does not impede the development of professional ethical standards. The focus of this approach to computer ethics is on the individual professional's responsibility in the practice of his or her craft. As the standards of this craft are being developed, so are the standards of professional computer ethics. The judgments about these standards will be guided by the values of the professional.

Computer ethics as presented here attempts a clear description of the relation of values to the work of the computer professional and sets forth criteria for making ethical decisions in that process. Focusing on stories about the failures of the product has misdirected us. They may be interesting stories to listen to, but they convey little information about computer ethics. I maintain that a focus on the process will resolve many of the problems discussed at the beginning of this paper. This approach will lead to the development of better computing artifacts.

▷ COMPUTERS, ETHICS, AND SOCIAL RESPONSIBILITY

Terry Winograd

INTRODUCTION

What Can a Computer Scientist Say?

Let me begin by admitting that it wasn't exactly clear to me just what the content should be for a "computer science keynote." The path seems clear for my colleagues who will present keynotes on philosophy and sociology, since those fields include ethics and values in their core subject matter. As a philosopher, one can develop theories of ethics. As a sociologist one can study the ways that people learn, change, and exhibit values. But as a computer scientist I don't study ethics and values; I study computers and computing. As we are all well aware, "ethics" and "values" aren't the kind of things to be addressed with the theories and techniques of computer science. For the computer scientist, they are not an object of study, but a domain in which we interpret and assess our actions as professionals.

My role, then, is not as an observer, but as a subject. As a computer professional and a teacher of future computer professionals, my concern is very personal. The questions aren't academic, but practical: not "What is done?" but "What should we do?"

So I will enjoy the liberty in this paper of not having to precisely define the difficult concepts we speak about or of having to argue the logical merits of a particular theory. Instead I will talk about how issues of computing and values show up in the work of our profession. You might think of it as being ichthyology from the point of view of the fish.

Also, in talking about these issues I will not try to draw a careful line between terms such as "ethics," "morals," "values," and "social responsibility." These distinctions can be important for some purposes, but I will interchange them freely here with more of a concern for the ring of the sentence than for the precise differentiation of the concepts.

The Personal Connection

When I speak of my own work, I include more than the narrow pursuit of research and development in computer science. For almost ten years I have been a participant in the work of Computer Professionals for Social Responsibility (CPSR), an organization that has brought together people from around the country (in fact, around the world) to share understandings and to act collectively in many of the areas that are being discussed in this text. That activity is not a diversion but a critical part of the work of a computer professional. One of the things I want to highlight is the way in which organizations like CPSR and the National Conference for Computing and Values (NCCV) play a central role in ethical conduct for computer professionals.

In addition, during the past three years, Helen Nissenbaum (now at Princeton)

and I have developed and taught a course on "Computers, Ethics and Social Responsibility" for undergraduate computer science majors at Stanford University. As all of us in academia know well, there is no better way to expand your own understanding than to throw yourself into a room full of bright undergraduates who want to master a difficult topic and expect you to help. Much of my understanding has grown from the generative interaction that comes in teaching, and that too is a central part of my work as a computer scientist. It has forced me into some hard and productive thinking about the questions being raised at a conference on Computing and Values.

What I Will Say

In this paper I will present and contrast some common views of how ethics and values are related to computing and see what these views imply for the activities we can undertake to promote ethical behavior and social responsibility. My emphasis is on the fundamentally social nature of ethical concerns: with looking beyond the role of the individual to the larger context of discourse and action that generates the world in which individuals make choices and to act. Rather than focusing on the isolated individual faced with an ethical dilemma, I want to direct our gaze to the larger swirl of human discourse, which is the source of the interpretations, values, and possibilities that make ethical choice meaningful.

The announcement for the NCCV conference declared a vision:

To integrate computer technology and human values in such a way that the technology advances and protects those values rather than doing damage to them.

This will require acts of individual moral courage, and it will be based on a lot more. We need to create an environment in which the consideration of human values in our technological work is not a brave act, but a professional norm. We need to produce a background of understanding in which it is simply taken for granted by all computer professionals that value considerations are foremost. We need to forge everyday practices and ways of teaching that reinforce that understanding.

In that spirit, I will argue that the kind of inquiry and discussion that motivate the conference, and that have been at the heart of CPSR's ten years of work, are a primary form of ethical behavior.

BEING A "GOOD" COMPUTER PROFESSIONAL

First, let us go back to the basic question of what values, ethics, and social responsibility have to do with computing. I said that ethics and values constitute "a domain in which we, and others, view and assess our actions as computer professionals." What do we mean by "assess our actions"?

Assessments and Competence

In every area of purposeful endeavor, there exist communities of assessment within which it is possible to meaningfully describe, compare, and evaluate action. As a computer scientist I am part of a scientific community with standards of practice, and

practices of assessment. There may be no straightforward quantitative measure of whether I am a "good computer scientist," but there are ways in which all of us measure the achievements of others and of ourselves. In the academic world these include publication records, peer review, awards, election to various professional and honorary societies, and the like. They also include less tangible but still consensual domains of reputation, status, and in the longer run your "place in the history of the field." I identify myself as belonging to a scientific community and I participate in its consensual processes for assessment. For example, my kids may think I'm a fantastic computer scientist because I could get Tetris running on our Macintosh, but I don't value this in the same dimension as the judgment of colleagues whom I consider part of the community.

When we look to the computer science community in general we see a notable lack of concern for many of the values addressed at the NCCV conference. There is an implicit definition of "good computer scientist" that dismisses people like Joe Weizenbaum as bothersome troublemakers, and accepts without qualms people who are oblivious of the value consequences of their actions. One of my colleagues, in a note rejecting my questioning the sources of research funding said he feared I would describe him, as "Having the moral fiber of a styrofoam cup." In judging whether people are "good computer scientists" the professional norms are strongly attuned to particular concerns and kinds of action and not to others.

But in our common sense assignment of "good" and "bad," we take a broader view. A "good baseball player" isn't just one who hits home runs, but one who contributes to the efforts of the team as a whole. He may be a great source of spirit and enthusiasm, a kind mentor to younger players, and a contributor in many other ways to the success of the team.

But in talking of the success of the team we're back to assessments again: What constitutes "success" for the computer science team? Again there is a gap between what we see in many of our work settings and what we are striving for here. Our measure of success needs to be the one quoted above:

To integrate computer technology and human values in such a way that the technology advances and protects those values rather than doing damage to them.

With this as our measure, we are ready to look for "good computer science."

What is the Domain of Ethical Action?

If our goal is to "advance and protect human values," then what kind of actions will further it? Ethics isn't an immediately obvious domain of actions. If you ask what competence is being developed in a cooking class, it is evidently "cooking." We can identify people at specific times as being engaged in cooking. But we are never "ethicking" in that simple sense. We may be performing an engineering job, making a living, doing scientific research, (or, for that matter, cooking) and find ourselves in situations where our actions raise some kind of ethical question. How do we identify those situations?

In some sense this is an "academic question." We all grow up with a tacit understanding that there can be things we do that are "right" and others that are "wrong," and that as autonomous individuals we have responsibility for choosing between them. We all have a sense that we *should* do what is right, even though that

isn't always what we end up doing. We also grow up in today's global pluralistic society with an awareness that although everyone has a sense that there is a difference between right and wrong, there is no agreement on just what actions should fall under which category. There are tremendous disputes between different cultural, religious and political groups, which have led to arguments, wars, and disagreements at all levels of society throughout history. But, on the other hand, there seem to be commonalities. Nearly everyone would agree that it is wrong to simply walk up to someone on a whim and inflict pain, and that it is right to help others in need.

People have debated for thousands of years what moral and ethical standards should be. Is there a universal ethics that applies to all people in all ages? Or is ethics a purely relative matter in that what is considered a fundamental moral principle by one people at one time may be equally validly rejected in another culture?

Now if I were a philosopher I would feel compelled to try to make sense of all this: to come up with a coherent moral philosophy that could serve as a basis for understanding what we see in the historical discourse about morality and for making decisions about our own actions. But, as I said at the beginning, I am taking the easy way out. I will leave the philosophical analysis for our colleagues who are much more skilled and knowledgeable, and will appeal to a rather commonsensical basis of agreement. I think we can all accept that in at least some interesting range of cases it makes sense to talk about doing "good" and "bad," and furthermore that we all, to some degree, accept the value of "doing good."

Further, we seem to have some general understanding of what kind of "doing good" constitutes an ethical or moral act. If you take a course on programming languages, you may learn that it is "good" to have a grammar that can be parsed by an LR(1) parser, and "bad" to have ambiguous constructs. But this domain of assessment, which is proper to the computer professional, doesn't seem to have much to do with the kind of human values we are discussing here. Something is missing in the equation.

Taking the naive view again, it seems obvious that the missing element has to do with a regard for the interests of others. There is a popular refrain about acts that are "illegal, immoral or fattening." The distinction between law and morality is an important one we will not go into here. But it is clear that there is something different about "fattening." It may be stupid or unhealthy or unwise to fill myself up on chocolate bars and potato chips, but few of us would consider it unethical or immoral. In general we take moral questions to involve a potential for conflict of interests. In the case of religious morality, the "other" may be a deity. For secular ethics, it is among people (and perhaps other life forms or embodiments of intelligence).

To be fair, this is a very complicated issue, but again in the spirit of simplification, we can accept that for most of the issues that attract our interest, our actions have consequences of value to others. Consider, for example the four clusters of values that are the focus of this conference: Privacy, security, ownership and fair access. In each case it is easy to identify the different parties and potential conflicts of interest, and we do so as a matter of course in teaching about these topics.

Finally in completing this background discussion, I want to make a key point about the role of intentions. Putting it generally, the domain in which an action is assessed is not necessarily the same as the domain in which the actor interprets it. If I ask what some person is doing, you may say that she is busy "establishing an image of authority" even though she is not consciously acting in that realm. Someone can

be assessed as a "great teacher" when what she sees herself doing is having an argument or commenting on a talk she is listening to.

Similarly, acts can be observed in the domain of ethics with respect to standard practices, independent of whether the person characterizes them that way. The fact that someone didn't think about the consequences of an act doesn't remove them from being subjected to moral judgment about it. In fact, we can take ethical oblivion as a key sign of "bad" behavior.

But thinking a moment further, it also doesn't seem appropriate to assess an act as wrong if there was no background of understanding in which it could show up as such. If we now see harmful consequences of the farming practices of primitive tribes, it doesn't mean that they were acting unethically to do them.

Again, we must look to the social context. A person does not exist in a vacuum, but as part of one or more social collectivities, with their shared interpretations of actions, values and assessments. A person cannot be held responsible for considerations that lie completely outside the range of vision opened up within these backgrounds. There may often be cases where an individual rejects the current consensus of society and appeals to a larger context of human meanings and values. But in doing this he or she is responsible for participating in the social discourse and not simply ignoring the concerns of his or her co-denizens. This means that a key component of moral action is the development of understanding within a social background, which is what provides the relevant field of choice for individuals.

THREE CARTOONS FOR HOW TO "ETHIC" WELL

So far we have been taking the view of the observer: one who interprets and assesses the acts that have been done by someone (who may be him/herself) in the past. Let us shift to the view of the doer: the person who is engaged in action that can have consequences in areas of values and morals. Faced with a particular range of possibilities, how does one "ethic" well.

One of the things that becomes clear in teaching this material to students is that people come to this question with a variety of tacit pre-understandings of what we are trying to do. They draw on images that are deeply embedded in our culture, and I want to present some of them in the form of cartoons that exaggerate, but also point out some of the key features. For each of the images, we need to ask several questions:

What are the assumptions that lie behind its perspective?

What problems does it raise?

How from that perspective do we develop people's competence to act?

The Angel/Devil Debate

The cartoon shown in Figure 1 is the familiar angel/devil debate you've all seen on Saturday morning TV. A character is faced with an ethical choice and is obviously having trouble deciding what to do. Sitting on one of his shoulders is a little pointy-tailed demon whispering into his ear "C'mon, take it, he'll never know you did." And on the other shoulder a haloed cherub, sweetly whispering "You know you shouldn't



FIGURE 1: The angel/devil debate

steal." In the end one of them is brushed away with a flick of the fingers and the other dances gleefully in victory.

There are several assumptions implicit in this view of morality:

1. You know what is right or wrong in the particular case.
2. Some part of you wants to do the thing that is wrong.
3. You need to exert moral strength to overcome this impulse and do what is right.

If, in fact, this view were the whole story, then the teaching of computer ethics would be a very different matter from what we see here today. Education directed to this kind of ethical competence might include sermons, examples (stories of sinners sizzling in Hell), and practices such as self-denial. In fact much of the resistance to the teaching of computer ethics within computer science departments comes from the impression that this is what it will consist of, and a skepticism as to whether such moralizing has any positive effects.

The fallacy in the angel/devil view is obvious if we look at Bynum's characterization of the goals for teaching computer ethics. (Bynum, 1992)

1. To sensitize students to computer ethics issues
2. To provide tools and methods for analyzing cases
3. To provide practice in applying the tools and methods to actual or realistic cases
4. To develop in the student "good judgment" and "helpful intuitions" for spur-of-the-moment decision making as computer professionals and computer users

Faced with an ethically problematic situation we must first recognize it as such. It doesn't appear with angels and devils drawn in the corners, but must be seen

through a background of interpretation in which ethical issues have been distinguished and made a part of our everyday discourse.

Of course I know I shouldn't kill other people. I walk around every day doing the right thing hundreds of times by not killing someone. But it isn't an ethical issue for me, it's part of the taken-for-granted background. But when I need to decide whether to build a worker-monitoring system or take research funding from a military agency, I am in the situation of debating what is in fact right and what the underlying issues are. It isn't a simple matter of steeling myself to be righteous. We must be able to recognize our specific situation as it fits into the context of issues and other cases that has shown up historically. These activities are skills to be learned and developed, not character traits like "moral will."

The Morality Computer

Having shifted the question from moral character to understanding and analysis, we find ourselves closer to the cartoon shown in Figure 2, the "morality computer."

Faced with a decision, you don't know which is the right action to take. Should you steal a horse to chase the bandit? Should you cut off life support to relieve suffering? Should you work on nuclear physics knowing the results may produce economic prosperity and also may lead to a weapon of mass destruction?

So you type the information into your morality computer, which has been programmed with the correct moral rules. It sifts the facts, weighs alternatives, makes judgments and pops out after a few microseconds with "Here's the right thing to do. . . ."

What are the assumptions lying behind this picture?

1. If you can figure out the right thing to do, then you will do it.

FIGURE 2: The morality computer





FIGURE 1: The angel/devil debate

steal." In the end one of them is brushed away with a flick of the fingers and the other dances gleefully in victory.

There are several assumptions implicit in this view of morality:

1. You know what is right or wrong in the particular case.
2. Some part of you wants to do the thing that is wrong.
3. You need to exert moral strength to overcome this impulse and do what is right.

If, in fact, this view were the whole story, then the teaching of computer ethics would be a very different matter from what we see here today. Education directed to this kind of ethical competence might include sermons, examples (stories of sinners sizzling in Hell), and practices such as self-denial. In fact much of the resistance to the teaching of computer ethics within computer science departments comes from the impression that this is what it will consist of, and a skepticism as to whether such moralizing has any positive effects.

The fallacy in the angel/devil view is obvious if we look at Bynum's characterization of the goals for teaching computer ethics. (Bynum, 1992)

1. To sensitize students to computer ethics issues
2. To provide tools and methods for analyzing cases
3. To provide practice in applying the tools and methods to actual or realistic cases
4. To develop in the student "good judgment" and "helpful intuitions" for spur-of-the-moment decision making as computer professionals and computer users

Faced with an ethically problematic situation we must first recognize it as such. It doesn't appear with angels and devils drawn in the corners, but must be seen

through a background of interpretation in which ethical issues have been distinguished and made a part of our everyday discourse.

Of course I know I shouldn't kill other people. I walk around every day doing the right thing hundreds of times by not killing someone. But it isn't an ethical issue for me, it's part of the taken-for-granted background. But when I need to decide whether to build a worker-monitoring system or take research funding from a military agency, I am in the situation of debating what is in fact right and what the underlying issues are. It isn't a simple matter of steeling myself to be righteous. We must be able to recognize our specific situation as it fits into the context of issues and other cases that has shown up historically. These activities are skills to be learned and developed, not character traits like "moral will."

The Morality Computer

Having shifted the question from moral character to understanding and analysis, we find ourselves closer to the cartoon shown in Figure 2, the "morality computer."

Faced with a decision, you don't know which is the right action to take. Should you steal a horse to chase the bandit? Should you cut off life support to relieve suffering? Should you work on nuclear physics knowing the results may produce economic prosperity and also may lead to a weapon of mass destruction?

So you type the information into your morality computer, which has been programmed with the correct moral rules. It sifts the facts, weighs alternatives, makes judgments and pops out after a few microseconds with "Here's the right thing to do. . . ."

What are the assumptions lying behind this picture?

1. If you can figure out the right thing to do, then you will do it.

FIGURE 2: The morality computer



2. There is a basic set of moral rules from which to deduce the rightness of action in any given case.
3. You may know the rules but not know how to apply them in this case, and more data, knowledge, or computation is needed.

In the "morality computer" view, the problem is determining what is the ethical action. Competence consists in knowing the moral axioms and having deductive skill in applying them to cases where the question "What is the right thing to do?" comes up. This approach appeals strongly to people with a background in science and engineering. When they encounter ambivalence and ambiguity they see it as a symptom that the problem has not been well formulated, or that we do not have enough knowledge. The fix is the kind of fix that works in technical domains: get the rules right, find the correct methods of applying them, and the right answers will come out. It offers the possibility of a "technical fix to the ethics problem."

Education in this perspective, as in other forms of education in science and engineering, is a matter of giving students the right principles and giving them practice in applying them to cases. At times we hear frustration from some of the students who take our course because we are not providing them with this kind of structure: we aren't able to give them the precise rules and methods, so they can learn to plug in the data and come up with answers.

But, of course, it doesn't take sophisticated philosophical reading to recognize that despite millennia of debate, humankind doesn't seem close to reaching agreement as to the general grounding for moral reasoning at all, much less the specific rules. Within any moderately diverse group of people you will find a wide range of beliefs: Some will believe that morality is grounded in some form of divine intention, others that it is a feature of human psychology, and others that it is some kind of "optimization" principle concerned with the welfare of the species. Some will base their moral reasoning on a structure of absolute principles—do's and don'ts—while others see it as some kind of calculus of costs and benefits.

Regardless of which approach you take, no matter how certain you are about the basic principles and rules, you find yourself puzzled by individual cases. One of the things we have become painfully aware of in nearly a half century of work on artificial intelligence is that there is a huge distance from abstract rules to real situations. Before rules can be applied, there must be interpretation as to how the terms in them actually fit the situation, and in doing so there is a wide-open field of human judgment and implicit understanding that has not yielded to logical analysis.

This all may be painfully obvious to those who have been working in the field of computer ethics, but let me give an example to clarify what I am pointing at.

Assume that you accept some form of the rule "Thou shalt not steal." Then in order to apply it, you need to know when an act is "stealing." We may define it as something like "taking property that belongs to someone else, without their consent." That's a good start, but what is "property?" There are clear examples such as someone's wallet or car, but what about their "ideas." Are those the kind of things that can be property at all? Now we are in a complicated realm of definitions, which has occupied philosophers (and lawyers) for centuries. The apparently simple notion of "property" has different interpretations in different cultures, legal systems, and traditions. Further, what do we mean by someone's "consent." What kind of consent is implied by opening a box that has a label on it saying "By opening this box you hereby agree to . . .?"

In reading the literature on computer ethics' we encounter many more such

examples and become painfully aware of how difficult it is to come up with consistent principles and standards for applying them to cases we encounter. The point should be obvious: we are not able to provide the kind of rules and methods that work in normal science and engineering, to come up with answers to problems. The morality computer is an idealized fantasy, and can mislead us if it shifts attention to a quest for the "right answer," away from the questioning activity that is required of each of us and involves us in a never ending dialog with others.

A Troupe of Jugglers

And this leads us to the cartoon of Figure 3, a troupe of jugglers. It may seem that juggling is too frivolous an activity to be a relevant analogy, but let us look more closely at several key features of the situation.

1. Engaged Activity First we note that the jugglers are constantly engaged in action. The first two cartoons directed our attention to conscious moments of decision, and put the locus of ethical action in determining the outcome. This cartoon suggests that we are always "thrown" into acting and that the assessments of ethics apply to these actions, not just those where we stop to ponder. This is suggested by the fourth of Bynum's descriptions of what we are doing in teaching:

To develop in the student "good judgment" and "helpful intuitions" for spur-of-the-moment decision making as computer professionals and computer users.

This "spur of the moment decision making" is the basic condition of acting in the world. In fact, it often does not show up to the actor as decision making at all. We all remember the interviews with someone who has jumped into a river to save a drowning child, when the interviewer asks "What made you decide to do it?" and the hero or heroine says "Decide? I didn't decide, I just jumped in?" In order to be

FIGURE 3: A troupe of jugglers



skillful at "ethicking" we need to develop the kind of continuing judgment in action that a juggler exhibits, not just the kind of careful argument that a logician applies in constructing a proof.

2. Social Context Second, the focus is not on the isolated actions of an individual, but on the coordinated actions of the troupe as a whole. What I do makes sense—is "right" or "wrong"—in the context of what others are doing. When I look to alternatives, I need to consider not just what else I might do, but what we all might do through some kind of agreement and coordination. As I suggested earlier, this is a key feature of ethical action. If we wait until someone is put into a true moral dilemma, we may get exciting drama, but we will not further the overall pursuit of values as much as if we develop standard practices that make it natural rather than heroic to do the "right thing."

One of the most powerful ethical acts we can each do is to participate in creating a social context in which the future actions of ourselves and others are consistently in line with our values. This includes educating our colleagues and students, working to develop professional standards, exploring new technologies and identifying their consequences for values. Even though we may not face hard individual ethical decisions as part of that work, we are actively engaged in the juggling process.

3. Evolving Understandings, Practices, and Standards Finally, we recognize in a juggling troupe the eternal need to learn and change. There is no ultimate "right" way to juggle. Clearly, any form of juggling will have to conform to the laws of gravity and physical motion. Less obviously, but plausibly, there may be perceptions of what is "good" that are grounded in the nature of the human animal and will be true across cultures and times. But within this, the community evolves practices and standards in which its members are trained, and by which their actions are assessed. Part of the skill we recognize in a community of jugglers is their ability over time to recognize new possibilities, develop skills in areas that hadn't been previously explored, and be sensitive to the changing environment in which they perform.

In some cases, this may require focus on detail: evolving a new concept of just what constitutes property and ownership in a new domain such as software, interfaces and algorithms. At other times, we can make major leaps. When Gandhi proposed nonviolent civil disobedience as a way of furthering the human values he cared about, he created a new "clearing"—a new way of looking at possibilities and taking actions, which could have meaning and power in the world of the late 20th century.

Now it should be clear that I favor this third cartoon, and to be fair, we should apply the same questions as we did to the other two. First, what are the assumptions?

1. There is a social activity in which we are engaged, in which characterizations and assessments in an ethical domain can be made.
2. There is no formal system that determines what is right, but there is an ongoing structure of discourse within a community, in which rightness is the issue, and in which there are stable areas of agreement.
3. An individual is never fully aware of what is possible to do, what effects an act will have, or how it will be assessed, and nevertheless will continue to act.

In a way this is comforting and in a way it is challenging. It is comforting because it does not posit some unachievable ideal: either the ideal of always having the strength to do the right thing, or the ideal of being able to determine just what is the right thing. Instead it puts the emphasis on being committed to entering into discussion with others and to taking seriously their concerns and understandings.

At the same time, for the same reason, it is challenging. It says that we will never have the satisfaction of knowing exactly when to apply our social and ethical concerns, or being confident that "Now we have it figured out." We are always being thrown into activity which may, in unanticipated ways, have implications for values, and we are part of a community that is always responsible for evolving new understandings and ways of "juggling" to maintain those values. It's exciting, but at times can be a little shaky.

WHAT DOES IT MEAN TO DO ETHICS AND SOCIAL RESPONSIBILITY?

In the final section, I want to look at what all this implies for the kinds of activities engaged in by individuals and organizations committed to making connections between computing and values.

All along I have been emphasizing the "doing" side of ethics: the way in which our actions more than our deliberations speak to our values. There are three key components in "doing" ethics and social responsibility:

1. Identifying social/ethical issues
2. Entering into serious discourse about the possibilities, with yourself and with others
3. Taking actions

Each of these has both an individual and a social component. There are cases where one person alone is faced with recognizing a problem, considering what to do, and doing it. Many of the most powerful pieces of literature in our culture grip us because they let us feel what it is like to wrestle with this ultimate responsibility of the individual.

But in this paper I want to focus more on the ways in which each of these components is situated in the actions of larger groups, and in particular the kinds of organizations represented by NCCV and by CPSR.

The activities with which I am most familiar are the work over the last ten years of Computer Professionals for Social Responsibility. Initially we were motivated by what appeared to be a mad rush towards nuclear war on the part of our government. As with many groups that emerged in the early 1980s with the words "social responsibility" in their names, we felt that the only responsible thing to do in that climate was to work wherever we had the most possibility of influence, in order to avert catastrophe.

In fact, there were many clear connections between computing and nuclear war, and in particular we came to focus heavily on the proposed Strategic Defense Initiative, or Stars Wars system. As computer professionals, we were sensitized to the prob-

lems of reliability and complexity that made the plans unrealistic, and could lead both to tremendous wasted resources and to a false sense of security that could dangerously destabilize the nuclear situation.

We approached this problem in a number of ways. CPSR members, individually and together, wrote papers analyzing the problems and bringing them to the attention of policy makers, both directly and through publication in newspapers and magazines. A number of us were active in the movement to have scientists pledge not to take research money from the SDI office. It was clear that the promise of research funds was being used as a lure to get tacit approval from the scientific community for the project. One person in particular who recognized this was David Parnas, who resigned from the panel that was convened to develop an analysis of the computing requirements, and later made public his analysis of the problems and of the ways in which the Pentagon was trying to influence scientists' assessments.

I mention all this not just because it is an important piece of CPSR's history, but to illustrate the range of activities that constituted "ethicking" for people involved with the issue. Some of the actions, like Parnas' resignation, can be viewed as explicit and difficult ethical choices made by individuals, and are noteworthy as such. But the impact they had was magnified by the fact that they were part of a coordinated campaign, in the context of organizations that could bring them to public attention and connect them the basic issues being fought. If one scientist quietly decides not to do SDI research, the impact is on that person and his or her work. If a whole community is involved, not only is the political impact greater, but the thinking of the entire profession is moved. The ability to recognize the potential of issues with consequences for values and ethics is increased for all those who become party to the discussion, even if they don't take direct action at the moment.

As the imminent danger of nuclear war appeared to subside, CPSR was able to take into fuller consideration the recognition that "social responsibility" really does cover more than just preventing annihilation (even though that is certainly a good place to start). Our concerns have overlapped strongly with those that form the core of the NCCV conference: Privacy, security, ownership and fair access. For example, we recently took part in a successful nationwide campaign to block the sale of a product proposed by Lotus that would have made it possible to find out private information about millions of consumers. In that case computers played a helping role as well, with much of the education and awareness about the issue being transmitted by electronic mail.

There isn't sufficient space to go into detail here on all of the CPSR program areas, but I have recounted this history to serve as an example as we look back at the three components of "ethicking" previously mentioned:

1. Identify Social/Ethical Issues The first is to identify social and ethical issues to which computers and computing are relevant. Each of us needs to ask how our actions as a computer professional might have ethical and social consequences, and there are a variety of answers. In some cases, as with every profession, the consequences grow directly from specifics of conduct. If I am asked to write a piece of life-critical software then I must proceed in a way that is responsible: taking care as best I am able given current software practices, and being honest with the clients about the risks and limitations. If I do less, I am cheating.

In some cases, harmful consequences come not from the quality of the work itself, but from the uses to which it will be put. If I am asked to build a program that makes it possible for employers to invisibly monitor the details of a worker's activi-

ties, I need to be aware of the consequences such programs can have in the workplace. Often, there is no clear boundary between beneficial and harmful technologies. A data base system used by the FBI to track drug dealers may have a positive effect on reducing the drug traffic, but can also be used to keep track of people with unpopular political beliefs. Even then, if I see the danger as great enough (what if those beliefs are racist and violent?) it may be a net gain to society if I produce such a system. But where is the boundary?

To go a step farther, there will be cases where the work itself is positive, but there is a larger context in which it can play a harmful role. Much of the research sponsored by the Strategic Defense Initiative is of a general kind that most of us would assess as having positive applications. The developing of computer networking, although originally sponsored by military agencies, has had a tremendous affect on our ability to function as computer professionals, and is rapidly becoming available to the entire population.

But what about the larger context? When a General in the U.S. Army testifies before congress that the scientific community is in favor of the SDI plans, as evidenced by the number of them who are actively working on research for it, what consequences has our research had? When a tremendous proportion of the research in computing in general is directed by the military, what long-term effect does this have on the kind of problems that are posed, or on the role of military thinking on the direction of our national economy?

There are rarely easy answers to such questions. In order to make responsible decisions about values, an individual needs a broad understanding of the consequences his or her actions might have in this overall situation. Such an understanding develops only through extended open discussion that brings in people from outside the computer profession as well as within it. It also extends beyond those who engage directly in it. The "styrofoam professor" I mentioned earlier has become conscious of issues of research funding through having interacted with me about my own rejection of military sponsorship. Even if he disagrees, the fact that the discussion exists (and has engaged his students as well) gives it a new standing in his "moral calculation."

2. Enter Into Serious Discourse about the Possibilities The second step in "ethicking" was to enter into serious discourse about the possibilities. I use the word "discourse" here instead of "thinking" to emphasize the social construction that is at the heart of decision-making even when a person does not directly enter into conversation with others. In a real, if extended sense, I am in discourse not only with the people I speak with but with those who have written the things that have influenced me, and those I have talked with, and in turn those in the future who will be influenced by what I say and write.

This includes people within the computer profession, and also in the larger society within which we work. The job of "public education" is a key part of creating the background of expectations that constitutes the fabric of ethical and social responsibility.

It should be obvious in looking at this conference, both at the participants and the materials that have been prepared and will be produced, that the weaving of the discourse is a function of groups of people who gather together (literally or through communications media) to think things through (or should we say "talk things through"?). This is a key role played by institutions and organizations devoted to issues of computers and ethics.

3. **Take Actions** Finally, the bottom line is the actions we take, both individually and collectively. It would be futile to try to catalog all of the different kinds of actions that have ethical implications. There are obvious individual acts such as whistleblowing, in which a clear value statement is being made in spite of some potential damage or loss to the actor. There are many other acts, such as choosing whether or not to work on a particular project or to take a job with a company that pursues projects of social concern, in which the decision is more subtle and the ethical factor may be one of many, which cannot be untangled in looking back at why the decision was made.

As members of a profession and its professional organizations, we also take acts that are intended to affect the direction and activity of the profession as a whole. They often don't have the visceral quality of whistleblowing or rejection of funding, but they contribute to creating the atmosphere in which those acts can be given sense. These include teaching, public education, working with professional organizations, developing standards and many other forms of everyday "ethicking."

Some of these activities have an overt political objective, such as lobbying for legislation or providing expertise to lawmaking and judicial bodies. Others operate at a broader cultural level, helping people both in the profession and in the public learn to see the issues, understand their consequences and apply human values to technical decisions. Although it is easier to get an individual to consider "Is it ethical (or socially responsible) to work on a bomb project?" It is equally important and more frequently relevant to ask "Is it ethical NOT to contribute my part to being responsible for how the public and the profession guide the ways in which computers will be used?" At the CPSR annual meeting last year, the slogan on the posters was "Technology is driving the future. It's time to find out who's steering." In the end, we all have a hand on the wheel.

CONCLUSION

Looking back at the three cartoons we might come up with different views of what we are engaged in when we participate in a conference on Computing and Values. From the angel/devil perspective, we could see it as a "revival meeting" at which we encourage each other to act in accord with our values, and tell stories that will help us to be resolute and remain steadfast. From the morality computer perspective, the conference is a think tank: it is our job to come up with the right rules and descriptions that will form the knowledge base that computer professionals can use to figure out what they should do.

These both have some particle of truth, but I much prefer to see our activity as a working session in which we are engaged in juggling the issues, ideas and discussions that generate the world of possibilities in which we and our colleagues live and work. We are creating those possibilities, increasing our own understanding and commitment to their value, and building a community that can continue to create and learn in the future.

NOTES

1. For a wide variety of cases, see the papers in Ermann and Williams, Dunlop and Kling, 1991, Parker, Swope and Baker, 1990, and Johnson and Snapper, 1985. For a

list of syllabi covering a broad range of topics related to computing and ethics, see Friedman and Winograd, 1989.

2. For further discussion of the issues of military funding, see Winograd, 1989 and the papers in Mitcham and Siekevitz, 1989.

REFERENCES

1. Bynum, Terrell. 1992. "Human values and the computer science curriculum." In *Computing and Human Values*, T. Bynum, W. Maner and J. Fodor. New Haven/RCCS.
2. Bynum, Terrell W., Walter Maner, and John L. Fodor, eds. 1992. *Computing and Human Values: Proceedings of the 1991 Conference*. New Haven: Research Center on Computing and Society.
3. Dunlop, Charles and Rob Kling, eds. 1991. *Computerization and Controversy*. Boston: Academic Press.
4. Friedman, Batya and Terry Winograd, eds. 1989. *Computing and Social Responsibility: A Collection of Course Syllabi*. Palo Alto. Computer Professionals for Social Responsibility.
5. Johnson, Deborah G., and John W. Snapper. 1985. *Ethical Issues In the Use of Computers*. Belmont, CA: Wadsworth Publishing Company.
6. Mitcham, Carl, and Philip Siekevitz, eds. *Ethical Issues Associated with Scientific and Technological Research for the Military*. New York: New York Academy of Sciences.
7. Parker, Donn, Susan Swope, and Bruce Baker. 1990. *Ethical Conflicts in Information and Computer Science, Technology and Business*. Wellesley, MA: QED Information Sciences.
8. Winograd, Terry. 1989. Strategic computing research and the universities. In *Directions and Implications of Advanced Computing*, Vol. 1, ed. J. Jacky and D. Schuler, 18-32. Norwood, NJ: Ablex. Reprinted in Dunlop and Kling, 1991, 704-716.

▷ PARAMEDIC ETHICS FOR COMPUTER PROFESSIONALS

W. Robert Collins and Keith W. Miller

Most computer professionals know that difficult ethical issues may arise in their work. We believe that these professionals want to "do the right thing." They accept their responsibilities as moral agents and they recognize that their special technical skills give them power and responsibilities. However, the will to act ethically is not sufficient; computer professionals also need skills to arrive at reasonable, ethical decisions. In this article we suggest a set of guidelines to