November 2000

Dear Fellow Members:

Our small Maine village, thanks to a grand 19th-century linguistic gesture, has an "opera house." The well-bricked building on the main street, however, has nothing to do with the music of Verdi and Puccini. Instead, the 500-seat auditorium of the Camden Opera House was built for a gaggle of itinerant jugglers, ragtimers, mesmerists, and other vaudevillians whose circuit through rural Maine brought levity and astonishment to gray winter weekends.

But in the first half of the 20th century a new technology, riding in on a silver screen, trampled these old halls into disuse. Why go see a nameless brood of touring comics when, for the same price, you could see Kay Tucker and other world-class stars of the cinema? As the century progressed, these old halls sank into disrepair.

But then something happened. In communities across New England, local citizens rose up, raised funds, and restored the plaster florets and wooden moldings to their gilt and primary-color splendor. Maybe it's an antiquarian impulse to preserve our history. Maybe it's a Yankee thrift that throws nothing away. But maybe it's a sense that, as the silicon-chip technologies invade our lives, we're losing our sense of place. How better to reassert that feeling of community, that civic camaraderie, than to reconstitute the public spaces of our lives?

So it seems ironically fitting that, last month, the Camden Opera House hosted the Camden Technology Conference–an astonishing aggregation of techies, entrepreneurs, venture capitalists, media commentators, and newly minted corporate executives of the digerati generation. Fitting, too, was the village-style topic we were debating: What does it mean to be human in a digital age?

The conference, the fourth in as many years, was convened by such luminaries as John Sculley (former CEO of Pepsi-Cola Company and of Apple Computer) and Robert Metcalfe (inventor of Ethernet and founder of 3Com Corporation). It drew together such different folks as Ira Glasser of the American Civil Liberties Union, Li Lu of Himalaya Capital Partners (better known as one of the principal student leaders of the Tiananmen Square uprisings in 1989), and Bill Joy, cofounder of Sun Microsystems. And it took steady aim at questions of isolation and connectedness, privacy and community, tribalism and globalism, work and play, creativity, spirituality, and ethics.

Ethics, in fact, ran like an obbligato throughout the sessions. If there was a subtext to the three-day meeting, it came from Joy's widely read article in last April's *Wired* magazine titled "Why the Future Doesn't Need Us." In it, he argued that self-replicating, fully conscious robots may one day become so sophisticated as to render humans obsolete. Opening the Camden conference, cybercommentator John Perry Barlow asked how such a robotic consciousness "could have a conscience"–unless, he noted, we first had one ourselves to give it. Ira Glasser, recalling the core convictions that have always defined the human community, called for "a revival of values." Mary Furlong, who started SeniorNet and heads ThirdAge Media, insisted that "spirituality really matters" and that, in the face of digital experiences, "people want a sense of soul." That

point was picked up by Leonard Sweet, founder of SpiritVenture Ministries and professor of evangelism at Drew University, who noted that since "humans have been hardwired by their creator to be creative," we have a "moral responsibility to evolve."

But the audience hung most attentively on Joy's words. Focusing on the intersection of genetic technology, robotics, and nanotechnology, he talked about soaring possibilities and sobering dangers. He predicted that advances in hardware and software would combine to improve computing speed by a factor of one trillion in the next few decades–a multiplier roughly equivalent to the difference in explosive power between a matchstick and a hydrogen bomb. As speed increases, he noted, costs will come down and computing power will be so democratized that everyone will have it. Among the results: a "post-scarcity society" that will "create sufficient material wealth to end material poverty in this century."

The downside? Once computer productivity makes human labor obsolete, and work is no longer needed to produce wealth, where will people find a "life worth living" that provides them with dignity and value? More important, how will we ensure that the new, democratized tools for creativity are used to develop rather than to destroy–through, perhaps, a new plague or "designed disease"? And what if a successor species of humanoid robots threatens to displace us?

All good questions. But behind them lay Joy's central and Faustian conundrum: Should we continue to pursue knowledge for its own sake? Or are there things we simply should not seek to know?

One might think, with all the potential wealth at stake in the Silicon revolution, that his answer would be, "Yes, but proceed with caution." Instead, this audience heard him saying something else.

"Maybe the things that are too dangerous to do we should not do," he said–noting, by way of example, that "as scientists, we cannot be complicit" in the development of technologies that lead to genocide. He acknowledged that among the core values of our age is the conviction that "knowledge is good" and "access to knowledge is good." Yet he insisted that "even if the search for truth is our highest value, there must be some exceptions." Or, put another way, "It's not always desirable to know everything."

That, of course, is an ethical position, a choice of one moral argument over another. And that's the point to which he had come by the end of his talk. Needed, he told his audience, is "ethics education in our scientific organizations" to help "reduce the risk" of the dangers that technological innovation is creating.

And what might the digital age look like if "knowledge for knowledge's sake" were *not* always allowed to prevail? What if science really *could be* conducted in an atmosphere of ethics? What if the development of speed, productivity, and wealth were principally valued for the good they could do the entire human community? And what if the goal behind every new line of technological inquiry were first to do no harm? That, in fact, was a point Joy had made in his

Wired magazine article. "Scientists and engineers," he wrote, should "adopt a strong code of ethical conduct, resembling the Hippocratic oath."

Beyond the prow of his intriguing suggestion lies an ocean of uncharted implications. In my comments to this same audience, I tried to sketch out some of them. Here they are, very much in draft form and inviting comment.

Such a code, to begin with, should recognize that while most technology really is morally neutral, those who create it, use it, and benefit or suffer from it have specific, direct responsibilities to the moral community. In other words, the code must address people more than technology. To do so in practical, effective, ways, the code must:

1. Grow up from within the private-sector technological community. If the code comes only from academics, it risks drowning in arcane disputes and the footnote mentality. If it comes only from governments, it risks relying only on the blunt instruments of law and compliance. But ethics is not about footnotes or compliance. It's about moral values. The code needs a values-based language of public discourse so robust that it can galvanize thought into action. Yet it must be so nuanced that it can handle the ambiguities of breakthrough discoveries happening at breakneck speed. Above all, like every code, it needs broad support from those who sign it.

2. Grow up right beside the technologies it addresses. Admirable though Bill Joy's case is, it's something of an afterthought. History will observe that only in April 2000 did the Silicon community publicly focus on the ethics at the intersection of a trio of technologies that were already decades old. Yet history will also note that when James D. Watson established the Human Genome Project at the National Institutes of Health in 1990, he insisted that some funding (currently 5 percent) be set aside for the so-called "ELSI" questions–the "ethical, legal, and social implications" of this research. Spending to date on the ELSI of the genome: $58.3 million, not after the fact but during the work. Fortunately, that's a pittance in today's venture-capitalized world. Doing good ethics, after all, is a bargain compared to doing good technology. Conclusion: If we want to do it, we can afford it.

3. Create ethical frameworks specific to cyberinnovation. Such frameworks can provide guidance on a range of right-versus-wrong issues of the Hippocratic "do no harm" variety. More important, they can help address more subtle right-versus-right issues. Example: Imagine that your bank, without your knowledge, crunches data from thousands of computerized records of your checks and credit-card payments. It produces a highly accurate profile of your attitudes and habits. Then, looking over the insurance you hold, it finds you're exposed in some areas and overpaying in others. So it offers you better, cheaper coverage. Will you be grateful for the savings–or angered at the invasion of your privacy? Both sides are highly charged with emotion–and both can easily be justified as "the right thing to do." How is a bank–or any other computer user–to think this through?

4. Promote awareness about digital ethics. In a word, the code must help educate, both within and beyond the cybercommunity. Here the motto for the techies is chillingly explicit: Do

it, or have it done to you. Why? Because the stakes are enormous. As John Perry Barlow told the audience in Camden, we may be witnessing a "fundamental metamorphic shift" more important than anything humans have seen since they captured fire. That can be scary. And fear, in a democracy, can translate into public mandates that could freeze up innovation. So predictable is this response, in fact, that you can almost express it as a formula:
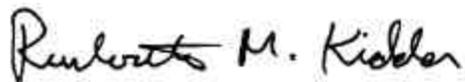
$$t - e \longrightarrow (po)^{x}$$

Or, *technology* minus *ethics* yields *exponential increases* in *public outrage*. Doubt it? Look at the future of genetically modified foods, now the subject of intense consumer outrage–and restrictive legislation–all over Europe.

What's needed? An ethics initiative that addresses the "ELSI" questions *before* they get embedded in the technologies. It needs to be practical and powerful, visible and voluble. It doesn't need to be expensive. But it needs to happen *now.* It can benefit from governmental persuasion, academic research, and nonprofit vision. In the end, however, it needs to be of, by, and for the innovators. Above all, it needs strong leadership from the Siliconese themselves.

Why? Because it's a lot easier to bang the dents out of a can from inside than from outside. Fail to address ethics by values-based initiatives from within, and it will be imposed by penalty-based laws from without. Self-regulation is ethics. Imposed regulation is law. And if you think such laws would be too hard to impose–that "they wouldn't dare"–then imagine how a few small changes in tax structure could effectively prevent entrepreneurs from cashing in on their inventions. How long would it take before the impulse to innovate were entirely sapped?

Admittedly, such laws would be harsh. Yet if Joy worries about a kind of robotic genocide, I worry about a failure to grasp the importance of ethics. So should the digerati. If the people gathered in Camden last month want to avoid harsh rules from a terrified world, the formula is straightforward: "Techie, regulate thyself."

Sincerely,

Rushworth M. Kidder