Ethics II

Ethics and Computer Science

- Last class, we discussed some historical scenarios
- We'll now consider some issues as they apply to computer science and engineering
- We'll look at project-specific issues and those that demand professional skill
- We'll consider ethical considerations towards both the public at large and towards our employer

Project Goals

- Is it good for the country or world if this system should exist?
- What about likely spinoffs or follow-ons?
- If you weren't paid to do it, would you want to live in the resulting world?

Example: Ballistic Missile Defense

- It seems obvious that it would be nice to be able to shoot down incoming nuclear missiles
- However, one can argue that the existence of such a technology makes a nuclear war *more* likely
- Note: *all* ballistic missile defense systems require a great deal of software

MAD

- For about 70 years, we've avoided nuclear war through "MAD"—Mutually Assured Destruction
- Both the U.S. and the U.S.S.R. had enough capability to absorb a devastating first strike and still destroy the other country
- Anti-ballistic missile (ABM) systems change the equation

MAD vs. ABM

- Any ABM system is imperfect—some percentage of missiles *will* get through
- If a first strike knocks out a lot of one side's missiles, the counterstrike will be smaller
- This in turn means that the ABM system will be more effective; the counterstrike may not destroy enough to deter whomever launched the first strike
- This creates an incentive for a massive surprise attack. . .
- (Of course, U.S. missile subs are largely invulnerable to preemptive strikes, which complicates matters even more.)

A Multilateral World

- MAD was a strategy for a two-party world: the U.S. (plus NATO, though that's a complication I won't go into) versus the U.S.S.R. (China had little or no ICBM capability.)
- Many more potentially hostile powers have nuclear bombs and missiles now
- Is a limited ABM system—one too small to destabilize the balance with Russia—now more rational?

Issues

- Given the strategic balance, is it good or bad for the world for an ABM system to exist?
- (This issue is one reason, among several, why the U.S. and the U.S.S.R. signed a treaty in 1972 drastically limiting ABM systems.)
- What about a limited ABM system, aimed at smaller states (e.g., DPRK), accidental launches, or the *Dr. Strangelove* scenario?
- Is defense more moral than offense, and especially more moral than MAD?
- Who gets to decide? What is the ethical obligation for employees, including computer scientists?

This is Hard!

- Answering these questions requires knowledge of game theory, psychology, diplomacy, and perhaps intelligence data
- Are programmers qualified to decide?
- (We'll get to the software issues in a few minutes)
- But—if offered a job working on it, you have to decide if it's right for you. What are your criteria?

Example: Deep Packet Inspection

- Some ISPs are deploying “Deep Packet Inspection” technology.
- Some countries are deploying “Deep Packet Inspection” technology.
- Is this good or bad?
- (What is Deep Packet Inspection?)

Deep Packet Inspection

- Ordinary firewalls work on packet headers: IP addresses and port numbers.
- (A *port number* is more or less an identifier for a specific service on a computer. The web is on port 80, mail is received on port 25, etc.)
- *Deep Packet Inspection* (DPI) is technology that permits examination of the *payload* of packets: what the actual message is.
- A DPI-based firewall could perhaps block web traffic that appeared to contain forbidden content
- Is there a privacy issue?
- Does this violate consumer expectations for Internet service?
- Should the free market settle this? (Is there an effective market for broadband consumer Internet?)

The Network Stack

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Link
1	Physical

DPI works on *layer 7* data.

DPI and Countries

- “The Iranian regime has developed. . . one of the world’s most sophisticated mechanisms for controlling and censoring the Internet. . . ” (all quotes from WSJ, 6/22/09)
- “China’s vaunted ‘Great Firewall’ . . . is believed also to involve deep packet inspection.”
- “Britain has a list of blocked sites, and the German government is considering similar measures. In the U.S., the National Security Agency has such capability”
- “The Australian government is experimenting with Web-site filtering to protect its youth from online pornography”
- “Internet censoring in Iran was developed with the initial justification of blocking online pornography”

Rationale

- “Mr. Roome of Nokia Siemens Networks said the company ‘does have a choice about whether to do business in any country. We believe providing people, wherever they are, with the ability to communicate is preferable to leaving them without the choice to be heard.’”
- “Nokia Siemens Networks provided equipment to Iran last year under the internationally recognized concept of ‘lawful intercept’”
- “Content inspection and filtering technology are already common among corporations, schools and other institutions, as part of efforts to block spam and viruses, as well as to ensure that employees and students comply with computer-use guidelines. Families use filtering on their home computers to protect their children from undesirable sites, such as pornography and gambling.”

Is Working on DPI Ethical?

- It has many very legitimate uses
- It's also a technology that can be and has been misused
- Who is responsible for making the ethical decision? Programmers? Corporate executives who sell the product? Users of it?
- What is the right answer?

Stuxnet

- The world's first (detected) cyberweapon, aimed at the Iranian uranium enrichment centrifuge plant in Natanz
- Allegedly created by the US and Israel
- Could infect many computers, but with very high probability would only cause damage to the centrifuge plant
- Was developing Stuxnet—or working on it—ethical?

Offensive Use of Cyberspace

Cyberexploitation Hacking into computers to spy on foreign companies and/or governments

“Preparing the Battlefield” Hack in, and plant back doors and other tools in case the computer or device may be needed later

Cyberattack Use cyberattacks in addition to or in place of “kinetic” weapons

Stuxnet: Pros and Cons

Iran can't be trusted with the Bomb	The US and Israel already have the Bomb
Stuxnet was better than an airstrike, which certainly would have killed people	Stuxnet, when reverse-engineered, taught others how to create cyber-weapons
International law on the use of force is well-understood	A precedent has been set: cyber-weapons are a legitimate tool. Are they "weapons", within the meaning of international law?

Is Stuxnet Ethical?

- There's been very little public discussion about targeting and use philosophy (which is quite *unlike* nuclear weapons)
- Stuxnet exploited stolen digital certificates—and certificates are the root of trust and security on the Internet
- "Sadly, the scientists are not pulling back the reins"...I don't think I ever saw anyone question what was being done." (Anonymous former government worker, quoted in Zetter's *Countdown to Zero Day*.)
- But—it did delay Iran's nuclear weapons program, and it didn't kill anyone

ABM Systems Redux: Can We Build One?

- An ABM system requires a massive amount of software
- You can never really test the system, since its behavior will depend on the precise timing of the precise inputs—radar signals, number of incoming missiles, enemy decoys, how many computing nodes have already been knocked out or are misbehaving because of radiation, etc.
- The 1970s and 1980s ABM systems required nuclear-armed missiles —all controlled by this large, complex, untestable software system. . .
- (Stuxnet was much simpler—and in fact the code was frequently updated during the attack.)
- What is the ethical response?

Professional Ethics

- Different professions have specific ethical principles
- Example (AMA): “An individual’s opinion on capital punishment is the personal moral decision of the individual. A physician, as a member of a profession dedicated to preserving life when there is hope of doing so, should not be a participant in a legally authorized execution.”
- Example (ABA): “As a representative of clients, a lawyer performs various functions. . . As advocate, a lawyer zealously asserts the client’s position under the rules of the adversary system. . . A lawyer’s representation of a client, including representation by appointment, does not constitute an endorsement of the client’s political, economic, social or moral views or activities.”
- The computing profession has several codes of ethics, too

The APA Code of Ethics

- Generally applicable principles, e.g., “Psychologists seek to promote accuracy, honesty and truthfulness in the science, teaching and practice of psychology. In these activities psychologists do not steal, cheat or engage in fraud, subterfuge or intentional misrepresentation of fact.”
- The patient is paramount: “Psychologists take reasonable steps to avoid harming their clients/patients, students, supervisees, research participants, organizational clients and others with whom they work, and to minimize harm where it is foreseeable and unavoidable.”
- Computer-related: “Psychologists who offer services, products, or information via electronic transmission inform clients/patients of the risks to privacy and limits of confidentiality.”

An Obligation for Confidentiality?

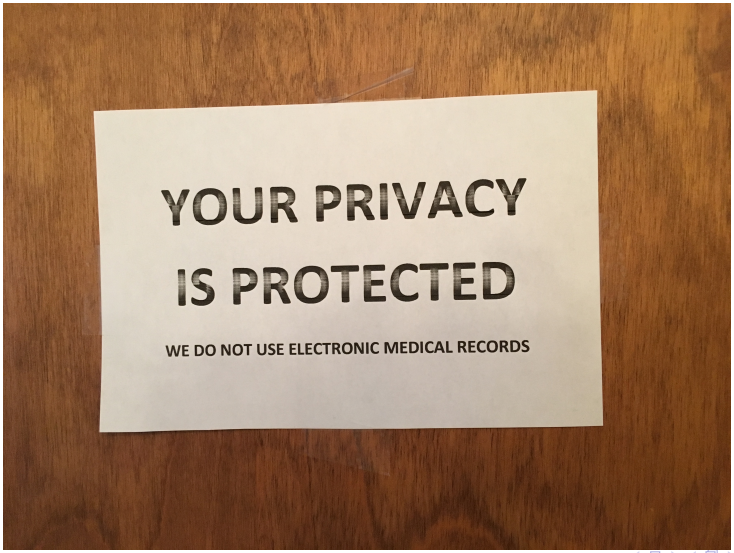

The ACM Code of Ethics

- (The Association for Computing Machinery is the oldest professional organization in the field, founded in 1947.)
- “A computing professional should. . . Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing.”
- “Avoid harm. . .
“‘Harm’ means negative consequences. . . “Well-intended actions. . . may lead to harm.”
- “Respect privacy
“Only the minimum amount of personal information necessary should be collected in a system.”
- “Give comprehensive and thorough evaluations of computer systems and their impacts, including analysis of possible risks.”

ACM Software Engineering Code of Ethics

- “Moderate the interests of the software engineer, the employer, the client and the users with the public good.”
- “Approve software only if they have a well-founded belief that it is safe, meets specifications, passes appropriate tests, and does not diminish quality of life, diminish privacy or harm the environment. The ultimate effect of the work should be to the public good.”
- “Disclose to appropriate persons or authorities any actual or potential danger to the user, the public, or the environment, that they reasonably believe to be associated with software or related documents.”
- “Identify, document, collect evidence and report to the client or the employer promptly if, in their opinion, a project is likely to fail, to prove too expensive, to violate intellectual property law, or otherwise to be problematic.”

What Do These Excerpts Say?

- First: you have a responsibility to society
- Second: that you must conduct your professional life in accordance with this principle
- Third: report some issues that appear to violate these codes
- In particular, you have to *honestly* assess a system design, especially if there are risks to others

Special Roles Have Special Concerns

- (SAGE/Usenix/LOPSA Code of Ethics) “I will access private information on computer systems only when it is necessary in the course of my technical duties. I will maintain and protect the confidentiality of any information to which I may have access, regardless of the method by which I came into knowledge of it.”
- On most computers, a system administrator can override any protection mechanisms.
- Often, it is *necessary* to do so to keep things running smoothly

Computer Security Work

- My specialty is computer security. To do that, I often have to find security holes
- Me: “I have the best job in the word—I get to think evil thoughts and feel virtuous about it.”
- Is what I do ethical? In other words, am I really virtuous?

Digging Down

- What if the hole is in someone else's system?
- What if the only way to test it is to exploit it? Example: `http://www.example.com?acctnum=1234567`
- Andrew "Weev" Auernheimer did more or less that to an AT&T site—and was sentenced to 41 months in jail for hacking
- (Conviction overturned on appeal on technical grounds. Despite the fact that Auernheimer writes for a neo-Nazi publication, he was represented pro bono on appeal by a Jewish lawyer—because that's consistent with the legal code of ethics.)

Should Security Holes be Reported or Saved for Use?

- (Most system penetrations are due to buggy code)
- Many governments look for security holes in commercial software
- Many governments buy them, too, on the open market
- Is it better—more ethical—to report the holes or to leave them in place?
- ☞ One helps the defense; the other helps the offense.

Do They Help?

- Many systems are never updated (and some aren't updatable)
- ☞ If a hole is reported and patched, the bad guys will reverse-engineer the patch to discover the hole, and use that information to attack unpatched systems
- On the other hand, just because you don't disclose the hole you've found doesn't mean it will remain unknown
- ☞ The vendor may find and patch it, rendering your attacks useless
- ☞ Or your enemies may find it, too, and use it to attack your systems

What Should Security Researchers Do With Holes?

- Publicly disclose all details?
- Notify the vendor?
- Sell them to the security company?
- Sell them to the (legal) hacking company?
- Sell them to a government? Should they distinguish among the different governments?

Public Disclosure

- If the bug is publicly know, people can take precautions
- But—if it's known, bad guys can exploit it
- Should it be disclosed?
- Should it be disclosed some time after notification of the vendor?
- Note that disclosure or the threat of disclosure often speeds up fixes

Bug Bounties

- Many companies reward researchers for (confidentially) disclosing security holes in their products
- Is this ethical?
- Or is it blackmail?
- (It's generally agreed that bug bounties are reasonably effective)

Uber and the Bug Bounty

- Uber paid someone $100,000 for finding keys (on Github!) that allowed access to databases with information on 57 million customer and driver accounts
- Uber did not report the data breach, though that is arguably required by the laws of most states
- The CSO and the responsible lawyer were fired for not handling the case properly—even though the previous Uber CEO, Travis Kalanick, had signed off on the deal
- Of course, Kalanick himself was forced out for doing ethically questionable things
- ☞ Partly for that reason, many people assumed that Uber paid off the “hacker” to cover up the incident
- What was the proper course of action here?

What About Research Funding?

- Do funding sources drive research?
- A major responsibility of a professor at a research university is to bring in grant money
- Does it matter if my money comes from the National Science Foundation instead of DARPA (Defense Advanced Research Projects Agency)? What about DARPA instead of IARPA (Intelligence Advanced Research Projects Agency)?
- What if the project has ethically good goals? (Note: I am *not* stating or assuming that DoD, the NSA, etc., are evil.)
- Does it matter if your goals match theirs?

Research Directions

- Should you pick research projects because they're socially desirable (according to whatever your metrics are)?
- (Rogaway specifically suggests this for cryptographers.)
- Should you avoid projects that are easily diverted to bad ends (again, for whatever value of "bad" you hold)?
- What is your responsibility if someone—a bad guy, a government spy agency, a terrorist—uses your ideas for their goals?
- (How about encryption?)

On the Other Hand. . .

- Are academics (and by extension, scientists and engineers) qualified to make moral judgments?
- Are they better off doing what they’re expert at?
- For academics: should they simply pursue knowledge without trying to change the world?

Who Were They?

- Allan J. McDonald?
- Roger Boisjoly?
- Bob Ebeling?

“My God, Thiokol, When Do You Want Me to Launch, Next April?

- McDonald: Director of the Space Shuttle Solid Rocket Motor Project for Thiokol; opposed launch
- Boisjoly: warned of the O-ring problem six months earlier; opposed launch
- Ebeling: sounded the initial alarm
- Thiokol and NASA management wouldn’t listen to them, and went ahead

The explosion of the Space Shuttle Challenger.

Normalization of Risk

- The Challenger disaster and the subsequent investigation showed that when there were warning signs but nothing went wrong, NASA came to believe that nothing would go wrong the next time
- That was fixed—for a while.
- And then the same phenomenon recurred

The Space Shuttle Columbia

ELLIOT: After the Challenger accident, did NASA indeed correct this cultural problem?
Mr. HARRIS: Well, they did for a while because the next 87 flights actually were all successful flights. But it wasn't entirely a cultural change, it was also partly luck because actually during that time as well, other problems, not o-rings, but other problems like o-rings emerged. The classic, of course, was foam falling of the external fuel tank of the space shuttles during life off. That was not supposed to happened. It was outside what was allowable but people said, well, it keeps happening and hitting the shuttle and nothing bad happens except little digs that we can replace later during repair so let's not worry about it. Well, as we all know, three years ago, the space shuttle Columbia had one of these pieces of foam hit the leading edge of one of its wings and that led to the second space shuttle disaster. The same cultural problem within NASA.
(https://www.npr.org/templates/story/story.php?storyId=5176563)

Back to Software

- A project will fail because it is too complex, or with too little time or money
- A system design is likely to prove unreliable

☞ This can be very hard to prove, especially because it's often a subjective judgment based on experience

- A system poses privacy risks

☞ In some countries, that can be a matter of law, too

- What about normalization of risk?

What to Do?

- Go to your management?
- Go public? What about confidentiality agreements?
- Inform legal authorities?
- Resign?

To What Extent are These Principles Honored?

- We've all seen and used awful computer systems
- What is the proper balance between cost and {function, reliability, security, privacy, etc.}?
- Who draws that line?
- What is the responsibility of the individual?

Questions?

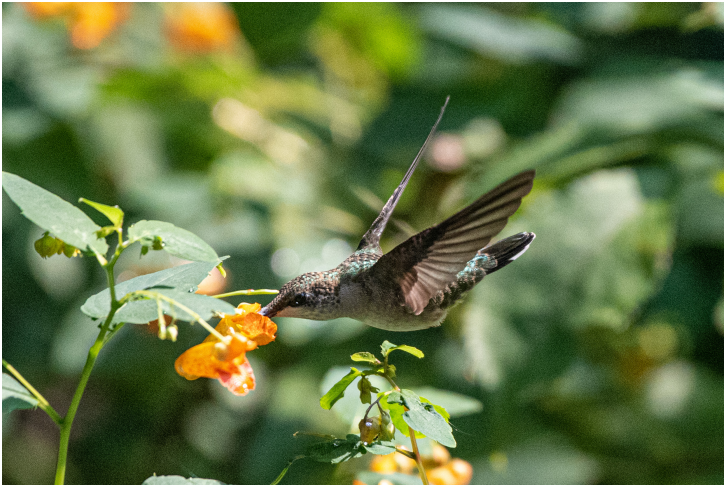

(Ruby-throated hummingbird, Central Park, September 14, 2021)