Solutions to Homework 2

Please note that some of the following solutions contain more information than you needed to provide to receive full credit. Others only provide examples, rather than an exhaustive list, of correct answers.

1. The number of points you get for this answer will depend on how well you support your opinion on whether or not the first-sale rule is appropriate. Here are two sample answers.

The first-sale rule is not appropriate in the digital realm:

Copyright law gives the owner of copyright the exclusive right "to distribute copies or phonorecords of the copyrighted work to the public by sale or other transfer of ownership, or by rental, lease, or lending." The reason that the first-sale rule is an effective mechanism for limiting that right in the analog world of books, CDs, DVDs, *etc.*, is that the content of such analog works is bound to and delivered in a physical object. It is this physical object that the owner of the copy is allowed to transfer under the first-sale rule, and once he has transferred this copy he no longer possesses it. Furthermore, the first-sale rule does *not* give the owner of a particular copy the right to create more copies; the only way for new copies to get into circulation in the first place is for the copyright owner to create them and place them into circulation, and this process has non-zero unit cost in the world of physical objects. The facts that each copy is possessed by exactly one person at a time and that unit production costs are non-zero are crucial to the logic of the first-sale rule. It is because the content of an analog work is bound to a physical object that both of these facts hold. This logic breaks down in the digital realm, where "content is liberated from medium."

The first-sale rule is appropriate in the digital realm:

The standard argument against the applicability of the first-sale rule to digital works is that the owner of a digital work can *both* transfer it to someone else and keep it and continue using it. That is because it is in the nature of digital works that the owner of a copy can duplicate that copy, perfectly and at zero cost. However, technical-protection technology may someday advance to the point at which that is not the case. Microsoft and other major players in the computer industry are developing "trusted platforms" that will supposedly be able to enforce security policies of various kinds, including copyright law. If these platforms work as advertised and are widely adopted, then an application such as an e-book reader could disallow use of a work on an owner's computer while that work is on loan (or after it has been given) to someone else.

2. Multi-channel retailing. (See lecture notes from January 28, 2003.)

3. Technically, Kazaa differs from Napster in that all of its major functions are done in a *distributed* fashion. Content that's available to be shared is stored on many "supernodes" all over the world, and any Kazaa peer can choose to become a supernode. A requester can send its request to what it thinks is a nearby supernode that can satisfy the request, but, if the supernode cannot satisfy the request, it can pass it along to another peer that it knows about. Because content that's in high demand will be listed by many supernodes, and requesters can contact these supernodes directly, there is no straightforward technical way to shut the service down. You could shut down one supernode, but the material listed on that node would probably be listed on other supernodes in places all over the world. New supernodes are created all the time, and shutting them all down will be impractical. In Napster, the database of online users that matched requesters of files with owners of those files was a centralized service. So it presented a single

point of vulnerability to technical (and legal) attacks. Kazaa presents no such single point of vulnerability.

Legally, there are two main reasons that it is difficult to shut down Kazaa. First, both the software vendor and many of the supernode owners are outside the U.S. Thus, it's not clear that U.S. copyright law (or the copyright law of any single country, for that matter) applies. Various powerful countries can form treaties in order to harmonize their copyright laws, but Kazaa supernodes can easily be established in countries without strong copyright industries that are unlikely to be parties to such treaties. The second reason that it's hard for law enforcement to shut down Kazaa is that millions of supernodes would have to be prosecuted. This is an inherently more difficult prosecution than the one against a centralized Napster service.

4. Here is a possible realization of this scheme for web-based broadcast of encrypted pages:

i. A customer U and a content server use a standard security protocol, *e.g.*, SSL, to create a "session key" K_U and transfer payment from U to the server.
ii. The server sends $k' = E(k, K_U)$ to U.
iii. U's browser computes $k = D(k', K_U)$, downloads the encrypted broadcast content, decrypts it using k, and displays it.

Technical obstacles to the effectiveness of this scheme include:

a. While the decrypted content is in use on U's machine, why can't U print, save, e-mail it to a friend, or otherwise "redirect" it?
b. At some point, the "master" decryption key k is going to be in use on U's machine. If a hacker breaks into U's machine while k is in use, he can steal k.
c. A poorly designed or implemented player application may neglect to delete decrypted content or decryption keys when it shuts down, particularly when it shuts down under an error condition or other unusual circumstances. So the windows of opportunity to misappropriate content or keys may be wider than they seem in (a) or (b) above.
d. The browser and player application that are needed for this distribution scheme might interact in insecure ways with other software running on the system. For example, the system may do "checkpointing," *e.g.*, periodically creating "snapshots" of all files currently in use so that, if a crash occurs and a time-consuming process is interrupted, it can be restarted at the last checkpoint rather than at the beginning. (The auto-save feature of a text editor does a simple form of checkpointing.) Obviously, checkpointing could lead to extra copies of plaintext or decryption keys being created and stored. Similarly, interaction with a file back-up system could lead to security problems.

Here is a possible realization of the overall scheme for cable-TV broadcast.

i. Subscribers $U_1, U_2, \ldots, U_N$ are given decryption keys $K_{U1}, K_{U2}, \ldots, K_{UN}$ when they sign up for service, and these keys are stored in their set-top boxes when the boxes are installed.
ii. When a program encrypted with key k is broadcast, the stream $k_1 = E(k, K_{U1})$, $k_2 = E(k, K_{U2})$, …, $k_N = E(k, K_{UN})$ is also broadcast. (Notice that the amount of per-subscriber information that must be created and broadcast is small; k_i is just a single number and hence much smaller than the encrypted program.)
iii. Subscriber U_i's set-top box decrypts the i^{th} element of the stream to recover $k = D(k_i, K_{Ui})$, uses k to decrypt the program, and displays the program.

One serious technical objection to this scheme is

e. The subscriber keys K_{Ui} are very long-lived. If a box is successfully copied, the copy can be used to decrypt *all* future programs until new keys (which might entail new boxes) are assigned.

(a)-(e) above are just representative answers, and you will get 5 points for any similar correct observation.

5. A) Layering and TCP/IP. HTTP is an application-layer protocol, and it is therefore not responsible for making the connection between client and server. It assumes that the layers below it work and fulfill their guarantees; in particular, it is built on top of the transport-layer protocol TCP which is responsible for a "reliable" connection between two computers over the Internet.

B) Camp mentions several security goals, and for each there are numerous examples of how basic HTTP fails with respect to that goal and numerous solutions. Listed below are some examples.

Confidentiality: HTTP does not guarantee secrecy because HTTP requests and responses are sent in plain text over the Internet, and anyone "eavesdropping" on the conversation in some way (*e.g.*, packet sniffing) can obtain the information exchanged. Using some form of encryption is a solution, because encryption technology attempts to prevent anyone but the sender and intended recipient from making sense of the information.

Availability and Scalability: Because HTTP runs on TCP/IP, it is possible that a reliable connection cannot be formed, that the network somehow fails or becomes too congested, or that the web server becomes subject to a malicious attack, *e.g.*, a TCP denial-of-service attack described in the textbook. Ways of solving this problem described in the text include load balancing (distributing material across several web servers, eliminating a single point of failure) and providing an offline way of checking information and conducting business (through a phone service or storefront).

Authentication: Standard HTTP does not validate identities. Although the client and server have purported IP addresses, not even that information is necessarily true; thus, it's hard to identify with whom one is communicating. Identity is important when providing access to private information; it may be nice to let people view bank account information, but there must be a way to check who is requesting it. The use of authentication schemes like passwords and digital signatures can solve this problem.

Integrity: By means of chance or malice, packets can become lost or corrupted in transit. Standard HTTP does not provide an error-detection mechanism (although some amount of error-detection is provided by TCP). Hash functions, encryption, and digital signatures (as discussed in class) can be used to detect when content has been changed.

Nonrepudiation: There is not enough information carried in a standard web-page exchange to prove that an action has or has not occurred, especially given the possible security problems discussed above. Digital signatures can be used to assure provenance and prevent repudiation. If a public key can be reliably associated with a unique signer, he cannot later claim that he did not sign a document using the corresponding private key.

6. There are many possible answers to this question, and there are many issues that you could have addressed in your answer. Some of these are discussed below. You will receive credit based on how well you justify your arguments.

The possibility contemplated in this question is that both creators of information (including data subjects and other parties to transactions) and collectors of information (including businesses as well as governments and law-enforcement agencies) could be given legal rights and responsibilities with respect to that information, potentially including, but not limited to, ownership rights. Thus, information in transaction records and corporate or government

databases could be viewed as analogous to copyright material in some respects and therefore potentially amenable to technical, legal, and business practices that work in the distribution and use of copyright material. This is a very controversial proposition, and there are numerous arguments both for and against it.

A) Technical-Protection Services

One argument that TPSs can be used to safeguard privacy and civil liberties if the TIA program goes forward is as follows: The fact that information about commercial transactions and other online activities of individuals is collected by the government does not necessarily imply that all of the information need be stored unencrypted and accessible to all government agencies all of the time for all purposes. Just as encryption, access-control systems, rights-management languages, and other computer-security technology is used by copyright owners to control the use of digital works while those works are resident on consumers' machines, these technologies could potentially be used on behalf of individuals to control the use of their personal information while that information is resident on government machines. For example, each person's travel records (including dates and times at which he entered and left the country, as well as where he was going or coming from) could be stored in encrypted form and only decrypted and made available to appropriate government officials if specific travel patterns, origins, or destinations known to be associated with national-security threats were detected. Most citizens' personal information would therefore remain private most of the time.

The standard argument against a TPS approach to protecting privacy in the presence of a government effort like TIA is that no one has yet been able to make a technical-protection system work on the kind of massive scale that would be necessary. Note that DRM systems for mass-market entertainment content have so far been notoriously easy to circumvent. For TPSs to work as privacy safeguards in TIA, someone would have to be responsible for key management, audit trails, and other forms of technical quality control. The idea of having a federal agency take that responsibility is problematic, because it is the federal government's potential abuse of power that the TPS is supposed to be preventing. In any case, there is currently no organization, public or private, that is known to be able to build and operate secure information systems on the necessary scale.

B) Intellectual-Property Law

It has often been suggested that individuals should be given property rights in the transaction records and other information about them that is created by their activity, particularly when that information has commercial value. This proposal is often put forth as a way of achieving a better balance of power between individuals (currently perceived to be powerless) and the organizations that collect information about them. If information about customers were the property of those customers, and a merchant wanted to sell it to a business partner, then he would actually have to convince his customers that his use of their information would provide benefit to them, *e.g.*, through enhanced service or better prices, and he would have to pay them to use it.

If one accepts the idea that personal information can be a form of intellectual property, then the owners of that property have rights, and the federal government, in particular, cannot simply take that property and do whatever it wants with it. Even if technical protection were infeasible and arbitrary amounts of personal information wound up in TIA databases, an individual could assert that his property rights had been violated if his information were used without his consent and without compensation. Just as warrants for searches, authorizations for wiretaps, and "eminent-domain takings" currently provide legal means for governments (at the local, state, and federal levels) to gain access to personal property, similar rules could be developed for use of personal

information in a TIA database without the owner's consent in circumstances in which national-security concerns dictated that it was appropriate.

The main argument against the intellectual-property rights approach to managing personal information is that there really are some fundamental differences between "information" and "intellectual property" as the latter is defined in U.S. law. Copyrights, patents, and trademarks have evolved to govern creative works and inventions. They explicitly do not govern *facts* such as who took a certain flight on a certain date. Furthermore, intellectual-property law was created in order to promote *dissemination* of creative works (and, more generally, "progress in science and the useful arts"), while the type of information-ownership rights being suggested here seem primarily aimed at keeping information private and *not* letting it be disseminated. If privacy of personal information is to be a U.S. citizen's "right," then that should that right not be inalienable in the same way that first-amendment rights are? It is in the nature of intellectual-property rights that they *are* alienable, and, in particular, creators of works sell their rights to employers and distributors all the time. Owners in weaker market positions sell their rights for less than owners in stronger positions. Do we want privacy rights to be subject to the same market forces? Finally, who exactly would own information that is created by many parties? For example, who would own an electronic transaction record that could not come into existence without the efforts of a merchant, a customer, a credit-card issuer, an ISP, and an application-software vendor?

C) Business Models

In **Database Nation**, Garfinkle offers the following bird's eye view of the past 38 years of data-privacy policy in the U.S.

> "Back in 1965, the United States government stood at a computational crossroads. On the table was a proposal to create a massive government database. But when details of the project reached the public, the project was terminated. Instead, the U.S. Congress held hearings on the threat of computers to privacy, a U.S. government commission formulated the idea of data protection, and a (relatively) small part of the U.S. government's executive branch was given the mission to enforce a new set of laws.
>
> "We blew it. A national database could have headed off the excesses of the credit reporting industry. If the system has allowed strong user controls, or had avenues for redress, it further could have prevented the sea of errors that exist in the plethora of private databanks today. Moreover, with a public system, uses of the data for purposes other than those originally intended would have been debated in public, rather than proposed and approved behind closed doors."

The proponents of TIA could use Garfinkle's text as a starting point for a marketing campaign. The gist of what he's saying amounts to "massive amounts of information are being collected anyway, probably primarily by companies but perhaps by governments as well. As a society, we should admit that it's happening, make it official, and regulate it. Checks and balances on database usage should be part of the TIA-authorization legislation, and new uses that are proposed for the database should be openly debated."

From a business perspective, one potential use is the streamlining of various major financial and legal processes that individuals engage in. If there were an official dossier on every adult in the U.S. that contained all important facts about his financial and legal status, then it should be trivial for all qualified people to get mortgages, business loans, visas for travel to exotic places, permission to adopt children, licenses to run small businesses (such as restaurants and bars)

where required, *etc.* Now, these procedures are often extremely time-consuming, costly, and seemingly arbitrary. There are no nationwide standards, numerous state and local governments (as well as banks and other large businesses) impose complex rules, and people often wind up hiring lawyers primarily to avoid having to sort through all of the rules.

Americans might be more comfortable with the prospect of a national database of personal information if it were accompanied by a federal mandate that anyone whose database record indicated that he was in sufficiently good financial and legal standing automatically qualified for certain mortgages, business loans, *etc.* There is currently a great deal of complacency about the TIA proposal; many people assume (perhaps incorrectly, given the error rates of some of the relevant technologies) that, because they are not terrorists, no computer program could ever find evidence that they are. Still, many people are wary of centralized government databases. Can they be swayed by the proposition that a *good* central-database record could be a universally acceptable, useful credential?

There are at least two compelling arguments against the proposition that such a database record would be either useful or universally acceptable. One is the fact mentioned in part (A) above that no one has yet been able to build and deploy a secure information system at this scale. How could the government mandate the use of this information by other organizations if it cannot credibly claim that the information is accurate and secure?

The second argument against such a proposition is that it is politically infeasible. Federalism is a key feature of U.S. political culture, and local and state governments will not relinquish their rights to determine what constitutes adequate documentation of financial or legal status. Similarly, large businesses (including banks) will not relinquish their rights to determine their own business practices, and professionals such as accountants and lawyers, who currently profit by helping people negotiate complex and seemingly arbitrary rules, will not regard “streamlining” as a good thing.