that it can take its bit packets from many different physical substrates, and deliver those packets for use by many different higher-level services.

The Reliability of the Internet

The Internet is remarkably reliable. There are no "single points of failure." If a cable breaks or a computer catches on fire, the protocols automatically reroute the packets around the inoperative links. So when Hurricane Katrina submerged New Orleans in 2005, Internet routers had packets bypass the city. Of course, no messages destined for New Orleans itself could be delivered there.

In spite of the redundancy of interconnections, if enough links are broken, parts of the Internet may become inaccessible to other parts. On December 26, 2006, the Henchung earthquake severed several major communication cables that ran across the floor of the South China Sea. The Asian financial markets were severely affected for a few days, as traffic into and out of Taiwan, China, and Hong Kong was cut off or severely reduced. There were reports that the volume of spam reaching the U.S. also dropped for a few days, until the cables were repaired!

Although the Internet *core* is reliable, the computers on the edge typically have only a single connection to the core, creating single points of failure. For example, you will lose your home Internet service if your phone company provides the service and a passing truck pulls down the wire connecting your house to the telephone pole. Some big companies connect their internal network to the Internet through two different service providers—a costly form of redundancy, but a wise investment if the business could not survive a service disruption.

The Internet Spirit

The extraordinary growth of the Internet, and its passage from a military and academic technology to a massive replacement for both paper mail and telephones, has inspired reverence for some of its fundamental design virtues. Internet principles have gained status as important truths about communication, free expression, and all manner of engineering design.

The Hourglass

The standard electric outlet is a universal interface between power plants and electric appliances. There is no need for people to know whether their power is coming from a waterfall, a solar cell, or a nuclear plant, if all they want to do is to plug in their appliances and run their household. And the same electric outlet can be used for toasters, radios, and vacuum cleaners. Moreover, it will instantly become usable for the next great appliance that gets invented, as long as that device comes with a standard household electric plug. The electric company doesn't even care if you are using its electricity to do bad things, as long as you pay its bills.

The outlet design is at the neck of a conceptual hourglass through which electricity flows, connecting multiple possible power sources on one side of the neck to multiple possible electricity-using devices on the other. New inventions need only accommodate what the neck expects—power plants need to supply 115V AC current to the outlet, and new appliances need plugs so they can use the current coming from the outlet. Imagine how inefficient it would be if your house had to be rewired in order to accommodate new appliances, or if different kinds of power plants required different household wiring. Anyone who has tried to transport an electric appliance between the U.S. and the U.K. knows that electric appliances are less universal than Internet packets.

The Internet architecture is also conceptually organized like an hourglass (see Figure A.3), with the ubiquitous Internet Protocol at the neck, defining the form of the bit packets carried through the network. A variety of higher-level protocols use bit packets to achieve different purposes. In the words of the report that proposed the hourglass metaphor, "the minimal required elements [IP] appear at the narrowest point, and an ever-increasing set of choices fills the wider top and bottom, underscoring how little the Internet itself demands of its service providers and users."

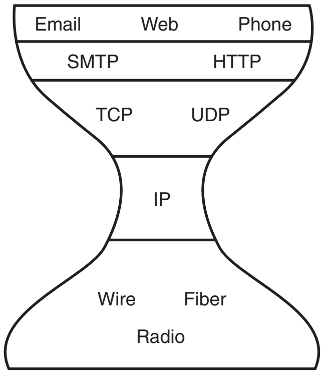

FIGURE A.3 The Internet protocol hourglass (simplified). Each protocol interfaces only to those in the layers immediately above and below it, and all data is turned into IP bit packets in order to pass from an application to one of the physical media that make up the network.

For example, TCP guarantees reliable though possibly delayed message delivery, and UDP provides timely but unreliable message delivery. All the higher-level protocols rely on IP to deliver packets. Once the packets get into the neck of the hourglass, they are handled identically, regardless of the higher-level protocol that produced them. TCP and UDP are in turn utilized by even higher-level protocols, such as HTTP ("HyperText Transport Protocol"), which is used for sending and receiving web pages, and SMTP ("Simple Mail Transport Protocol"), which is used for sending email. Application software, such as web browsers, email clients, and VoIP software, sit at a yet higher level, utilizing the protocols at the layer below and unconcerned with how those protocols do their job.

Below the IP layer are various physical protocol layers. Because IP is a universal protocol at the neck, applications (above the neck) can accommodate various possible physical implementations (below the neck). For example, when the first wireless IP devices became available, long after the general structure of the Internet hourglass was firmly in place, nothing above the neck had to change. Email, which had previously been delivered over copper wires and glass fibers, was immediately delivered over radio waves such as those sent and received by the newly developed household wireless routers.

Governments, media firms, and communication companies sometimes wish that IP worked differently, so they could more easily filter out certain kinds of content and give others priority service. But the universality of IP, and the many unexpected uses to which it has given birth, argue against such proposals to re-engineer the Internet. As information technology consultant Scott Bradner wrote, "We have the Internet that we have today because the Internet of yesterday did not focus on the today of yesterday. Instead, Internet technology developers and ISPs focused on flexibility, thus enabling whatever future was coming."

Indeed, the entire social structure in which Internet protocols evolved prevented special interests from gaining too much power or building their pet features into the Internet infrastructure. Protocols were adopted by a working group called the Internet

THE FUTURE OF THE INTERNET—AND HOW TO STOP IT

This excellent book by Jonathan Zittrain (Yale University Press and Penguin UK, 2008) sees the vulnerabilities of the Internet—rapidly spreading viruses, and crippling attacks on major servers—as consequences of its essential openness, its capacity to support new inventions—what Zittrain calls its "generativity." The book reflects on whether society will be driven to use a network of less-flexible "appliances" in the future to avoid the downsides of the Internet's wonderfully creative malleability.

Engineering Task Force (IETF), which made its decisions by rough consensus, not by voting. The members met face to face and hummed to signify their approval, so the aggregate sense of the group would be public and individual opinions could be reasonably private—but no change, enhancement, or feature could be adopted by a narrow majority.

The larger lesson is the importance of minimalist, well-selected, open standards in the design of any system that is to be widely disseminated and is to stimulate creativity and unforeseen uses. Standards, although they are merely conventions, give rise to vast innovation, if they are well chosen, spare, and widely adopted.

Layers, Not Silos

Internet functionality could, in theory, have been provided in many other ways. Suppose, for example, that a company had set out just to deliver electronic mail to homes and offices. It could have brought in special wiring, both economical and perfect for the data rates needed to deliver email. It could have engineered special switches, perfect for routing email. And it could have built the ideal email software, optimized to work perfectly with the special switches and wires.

Another group might have set out to deliver movies. Movies require higher data rates, which might better be served by the use of different, specialized switches. An entirely separate network might have been developed for that. Another group might have conceived something like the Web, and have tried to convince ordinary people to install yet a third set of cables in their homes.

The magic of the hourglass structure is not just the flexibility provided by the neck of the bottle. It's the logical isolation of the upper layers from the lower. Inventive people working in the upper layers can rely on the guarantees provided by the clever people working at the lower layers, without knowing much about *how* those lower layers work. Instead of multiple, parallel vertical structures—self-contained silos—the right way to engineer information is in layers.

And yet we live in an information economy still trapped, legally and politically, in historical silos. There are special rules for telephones, cable services, and radio. The medium determines the rules. Look at the names of the main divisions of the Federal Communications Commission: Wireless, wireline, and so on. Yet the technologies have converged. Telephone calls go over the Internet, with all its variety of physical infrastructure. The bits that make up telephone calls are no different from the bits that make up movies.

Laws and regulations should respect layers, not the increasingly meaningless silos—a principle at the heart of the argument about broadcast regulation presented in Chapter 8.

End to End

"End to End," in the Internet, means that the switches making up the core of the network should be dumb—optimized to carry out their single limited function of passing packets. Any functionality requiring more "thinking" than that should be the responsibility of the more powerful computers at the edge of the network. For example, Internet protocols could have been designed so that routers would try much harder to ensure that packets do not get dropped on any link. There could have been special codes for packets that got special, high-priority handling, like "Priority Mail" in the U.S. Postal Service. There could have been special codes for encrypting and decrypting packets at certain stages to provide secrecy, say when packets crossed national borders. There are a lot of things that routers might have done. But it was better, from an engineering standpoint, to have the core of the network do the minimum that would enable those more complex functions to be carried out at the edge. One main reason is that this makes it more likely that new applications can be added without having to change the core—any operations that are application-specific will be handled at the edges. This approach has been staggeringly successful, as illustrated by today's amazing array of Internet applications that the original network designers never anticipated.

Stupid Networks

Another way to understand the Internet's end-to-end philosophy is to realize that if the computers are powerful at the edge of the network, the network itself can be "stupid," just delivering packets where the packets themselves say they want to go. Contrast this with the old telephone network, in which the devices at the edge of the network were stupid telephones, so to provide good service, the switching equipment in the telephone office had to be intelligent, routing telephone signals to where the network said they should go.

Separate Content and Carrier

The closest thing to the Internet that existed in the nineteenth century was the telegraph. It was an important technology for only a few decades. It put

The Victorian Internet

That is the title of an excellent short book by Tom Standage (Berkley Books, 1999), making the argument that many of the social consequences of the Internet were seen during the growth of the telegraph. The content-carrier conflict is only one. On a less-serious level, the author notes that the telegraph, like the Internet, was used for playing games at a distance almost from the day it came into being.

the Pony Express out of business, and was all but put out of business itself by the telephone. And it didn't get off to a fast start; at first, a service to deliver messages quickly didn't seem all that valuable.

One of the first big users of the telegraph was the Associated Press—one of the original "wire services." News is, of course, more valuable if it arrives quickly, so the telegraph was a valuable tool for the AP. Recognizing that, the AP realized that its competitive position, relative to other press services, would be enhanced to the extent it could keep the telegraph to itself. So it signed an exclusive contract with Western Union, the telegraph monopoly. The contract gave the AP favorable pricing on the use of the wires. Other press services were priced out of the use of the "carrier." And as a result, the AP got a lock on news distribution so strong that it threatened the functioning of the American democracy. It passed the news about politicians it liked and omitted mention of those it did not. Freedom of the press existed in theory, but not in practice, because the content industry controlled the carrier.

More on Information Freedom

The SaveTheInternet.com Coalition is a pluralistic group dedicated to net neutrality and Internet freedom more generally. Its member organizations run the gamut from the Gun Owners of America, to MoveOn.org, to the Christian Coalition, to the Feminist Majority. Its web site includes a blog and a great many links. The blog of law professor Susan Crawford, `scrawford.net/blog`, comments on many aspects of digital information freedom, and also has a long list of links to other blogs.

Today's version of this morality play is the debate over "net neutrality." Providers of Internet backbone services would benefit from providing different pricing and different service guarantees to preferred customers. After all, they might argue, even the Postal Service recognizes the advantages of providing better service to customers who are willing to pay more. But what if a movie studio buys an ISP, and then gets creative with its pricing and service structure? You might discover that your movie downloads are far cheaper to watch, or arrive at your home looking and sounding much better, if they happen to be the product of the parent content company.

Or what if a service provider decides it just doesn't like a particular customer, as Verizon decided about Naral? Or what if an ISP finds that its customer is taking advantage of its service deal in ways that the provider did not anticipate? Are there any protections for the customer?

In the Internet world, consider the clever but deceptive scheme implemented by Comcast in 2007. This ISP promised customers unlimited bandwidth, but then altered the packets it was handling to slow down certain data transmissions. It peeked at the packets and altered those that had been generated by certain higher-level protocols commonly (but not exclusively) used for downloading and uploading movies. The end-user computer receiving these altered packets did not realize they had been altered in transit, and obeyed the instruction they contained, inserted in transit by Comcast, to restart the transmission from scratch. The result was to make certain data services run very slowly, without informing the customers. In a net neutrality world, this could not happen; Comcast would be a packet delivery service, and not entitled to choose which packets it would deliver promptly or to alter the packets while handing them on.

In early 2008, AT&T announced that it was considering a more direct violation of net neutrality: examining packets flowing through its networks to filter out illegal movie and music downloads. It was as though the electric utility announced it might cut off the power to your DVD player if it sensed that you were playing a bootleg movie. A content provider suggested that AT&T intended to make its content business more profitable by using its carrier service to enforce copyright restrictions. In other words, the idea was perhaps that people would be more likely to buy movies from AT&T the content company if AT&T the carrier refused to deliver illegally obtained movies. Of course, any technology designed to detect bits illegally flowing into private residences could be adapted, by either governments or the carriers, for many other purposes. Once the carriers inspect the bits you are receiving into your home, these private businesses could use that power in other ways: to conduct surveillance, enforce laws, and impose their morality on their customers. Just imagine Federal Express opening your mail in transit and deciding for itself which letters and parcels you should receive!

Clean Interfaces

The electric plug is the interface between an electric device and the power grid. Such standardized interfaces promote invention and efficiency. In the