

50
digital ideas
you really need to know



Tom Chatfield

50 digital ideas

you really need to know



Tom Chatfield

Quercus

New York • London

© 2011 by Tom Chatfield

All rights reserved. No part of this book may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without permission in writing from the publisher, except by reviewers, who may quote brief passages in a review. Scanning, uploading, and electronic distribution of this book or the facilitation of the same without the permission of the publisher is prohibited.

Please purchase only authorized electronic editions, and do not participate in or encourage electronic piracy of copyrighted materials. Your support of the author's rights is appreciated.

Any member of educational institutions wishing to photocopy part or all of the work for classroom use or anthology should send inquiries to Permissions c/o Quercus Publishing Inc., 31 West 57th Street, 6th Floor, New York, NY 10019, or to permissions@quercus.com.

ISBN 978-1-62365-181-7

Distributed in the United States and Canada by Random House Publisher Services
c/o Random House, 1745 Broadway
New York, NY 10019

www.quercus.com

Contents

Introduction

FUNDAMENTALS

- 01 The internet
- 02 The world wide web
- 03 Internet service provision
- 04 Email
- 05 Personal computing
- 06 Servers
- 07 Browsers
- 08 Markup languages
- 09 Search

TOWARD THE DIGITAL PRESENT

- 10 Web 2.0
- 11 Netiquette
- 12 Blogging
- 13 Aggregation
- 14 Chatting
- 15 File sharing
- 16 Streaming media
- 17 Rich internet applications
- 18 Wireless
- 19 Smartphones

THE DARK SIDE

- 20 Malware
- 21 Spam
- 22 Privacy
- 23 The deep web
- 24 Hacking
- 25 Cyberwar

PLEASURE AND LEISURE

- 26 Social networking
- 27 Games consoles
- 28 Mashups
- 29 Culture jamming

BUSINESS AND GOVERNMENT

- 30 E-commerce
- 31 Online advertising
- 32 Analytics
- 33 Optical character recognition

- 34 Machine translation**
 - 35 Location-based services**
 - 36 Virtual goods**
 - 37 eGovernment**
-

TRANSFORMATIVE TRENDS

- 38 Crowdsourcing**
- 39 Free software movement**
- 40 Digital distribution**
- 41 Cloud computing**
- 42 Going viral**

SHAPING A DIGITAL FUTURE

- 43 Virtual worlds**
- 44 Avatars**
- 45 Net neutrality**
- 46 Semantic web**
- 47 Augmented reality**
- 48 Convergence**
- 49 The internet of things**
- 50 Distraction**

Glossary

Introduction

What do we mean by “digital”? In one sense, we simply mean something whose ultimate existence is a string of ones and zeroes. A digital object, unlike the analogue stuff of the world or even of early computing, is a binary pattern.

Out of this simple fact have come some of the most transforming cultural forces of the second half of the 20th century and, now, the first half of the 21st. For perhaps the most remarkable property of digital objects—whether their ones and zeroes encode words, music, images, applications, web browsers, or databases containing much of the store of human knowledge—is their compatibility with each other.

For the first time in human history, it is possible almost endlessly to copy and distribute words, sounds, images, ideas; and it is possible to access, adapt and create all of these on the same devices.

The history of digital ideas goes back to well before the 20th century. It is in part the millennia-old story of mathematics, as well as the story of centuries of ingenious mechanical design preceding the last century’s developments in the fields of electronics and computing.

I have chosen here, however, to focus on the more recent past—and on those aspects of the digital present that seem most powerfully to be shaping our future. As a result, this is overwhelmingly a book written about and around the internet.

As I write these words, in 2011, around two billion people—almost one third of humanity and one half of the planet’s adult population—have some form of access to the internet. Thanks to the spread of mobile internet access, this figure will continue to soar over the coming decade, as will the use of those online services steadily reshaping much of what it means to be a member of a modern society.

The news is not all good: growth, amelioration and freedom are not to be taken for granted, and indeed are not enjoyed by many of the world’s digital citizens. The internet is as powerful a force in the hands of many of the world’s repressive and censorious regimes as it is in the hands of those using it to liberate, educate, connect and delight.

Similarly, it is a breeding ground for the best and the worst of humanity alike: the scammers and the selfless; the entertainers and the mockers; the entrepreneurs and the predators. But this only makes understanding its history, structures, potentials and possible futures all the more vital.

01 The internet

The internet is best thought of not so much as a technology as an infrastructure: a mass of interlinked hardware and software—from deep sea cables and telephone wires to desktop computers and mobile phones—connecting an increasingly large proportion of the world’s computing devices. Many services operate via the internet—including, perhaps most famously, the world wide web—but the internet itself predates these. It is the vast, physical network within which much of modern digital culture exists.

The history of the internet dates back to the Cold War and the communications networks developed by America in the wake of Russia’s launch of the satellite Sputnik—the first man-made object ever to orbit the earth—in 1957. Shocked by its rival’s achievement, America poured resources into developing new communications technologies. In particular, the aim was to develop communications networks able to function even if a disaster destroyed large parts of the network itself.

This early research culminated in a 1968 report entitled “Resource Sharing Computer Networks,” which laid the foundations for the first computer network based on a system known as “packet switching”—where all data transmitted between machines is broken down into small blocks or packets. The first computer network built using this technology was called the ARPANET (Advanced Research Projects Agency Network), and began operations in 1969, initially connecting four sites at the University of California in Los Angeles, the Stanford Research Institute, the University of California in Santa Barbara and the University of Utah.

ARPANET grew rapidly. By 1970 it had reached the east coast of America. The first ever piece of electronic mail was sent across it in 1971. By September 1973, 40 machines across America were connected to the network, and the first method for transferring computer files between them—known simply as “File Transfer Protocol” or FTP—had been implemented. In December 1974, the word “internet” itself was used for the first time in a paper by Vinton Cerf, Yogen Dalal and Carl Sunshine as an abbreviation for the term “internetworking.” It referred to what was becoming a revolutionary idea: a global communications meta-network, itself formed by combining multiple networks of machines that used the same protocol for sharing packets of information. This “network of networks” idea is perhaps the core idea behind the modern internet.

“The internet is becoming the town square for the global village of tomorrow.”

Bill Gates

Protocols Cerf, Dalal and Sunshine’s 1974 paper formulated two ideas that would be central to the future of the internet: “Internet Protocol” (IP) and “Transmission Control Protocol” (TCP). Between them, these protocols explained the precise way in which data should be broken down into packets and sent between computers. Any computer that used these methods—usually abbreviated to TCP/IP—should, in theory, be able to communicate with any other computer also using them. Internet protocols defined the route that the data should travel along to connect computers at particular locations—or “IP addresses” as they would come to be known—while transmission control protocol ensured that the packets of data were sent in a reliable, ordered manner.

IPv6

Since 1981, every computing device attached to the internet has had a unique numerical address allocation to it via the fourth version of the Internet Protocol system on which the net is based: IPv4 for short. Today, however, the phenomenal growth in internet-based resources means that IPv4 is running out of numbers. IPv4 is able to support “only” 4,294,967,296 unique internet addresses—the maximum size of a 32-digit binary number. The newest, sixth version of Internet Protocol is based on 128-digit binary numbers, allowing more than a billion times as many addresses again as version four. Switching the internet over to a new protocol is a huge challenge, especially for older physical hardware, and it is being performed and tested in several stages. But it is an increasingly urgent task, with “address exhaustion” for IPv4 dated for some regions to mid-2011.

Over the next decade, much work was done ensuring that it was possible for as many different types of computer as possible to connect successfully to each other using TCP/IP. This culminated in 1983 with the switch of all the computers on the ARPANET system to TCP/IP rather than the older packet-switching system they had previously used. In 1985, America’s National Science Foundation commissioned its own network of computers, designed to run in American universities using TCP/IP, and in 1988 this network was opened up, allowing other emerging computer networks to connect to it. The TCP/IP protocol made it easy for different kinds of machine and network to link together, and the end of the 1980s saw the launch of the first commercial Internet Service Providers (or ISPs), offering network access to companies and individuals.

“The Internet is the first thing that humanity has built that humanity doesn’t understand, the largest experiment in anarchy that we have ever had.”

Eric Schmidt

Growth By the early 1990s, large parts of the world—led by universities and research institutions—had independently begun to use computer networks based on TCP/IP protocols, something that made extremely easy for these networks to connect to each other, share files and data, and send electronic mail. It was only, however, with the invention of the “world wide web” in 1989 by Tim Berners-Lee that the tools it was possible to use via the internet began to become genuinely accessible to ordinary computer users as well as academics and researchers. Through the 1990s, the number of people connected by the internet more than doubled each year on average, a rate that only slightly lessened through the next decade. It is estimated that, at the end of 2010, approximately two billion people, almost one third of humanity, were connected.

Net governance

Set up in America in 1998, the Internet Corporation for Assigned Names and Numbers (ICANN) is the single most important organization for regulating the future of the internet. Its primary responsibility is managing Internet Protocol address spaces: that is, ensuring the allocation and management of a unique address for every site and service using the internet. It also allocates blocks of addresses and what are known as “top-level domains” to different parts of the world: for example, allocating websites that end in “uk” to British-based websites, in “au” to Australian

websites, in “com” to commercial operations, and so on. ICANN is a charitable organization and must balance the needs of different parts of the globe while attempting to preserve a coherent, cohesive internet structure: a considerable challenge due to the net’s increasingly international nature.

The world wide web is explained in more detail in the next chapter. With its arrival, and the rapid spread of websites and browser technology, the internet began to play the role that most people would recognize today: a truly global network, connecting individuals and nations through an increasingly vast proliferation of computing devices. But its future cannot be taken for granted. The internet relies on continued cooperation between nations and individuals; and also on the continued construction and upgrading of hardware—that is, cables and computer servers—capable of handling the increasingly vast amounts of information being exchanged online, in the form of not only words, but videos, images, sounds and complex applications.

the condensed idea

the internet connects everything

timeline

- 1969** US Department of Defense creates ARPANET
- 1974** Word “internet” is used for the first time
- 1982** ARPANET switches to Transmission Control Protocol (TCP) and Internet Protocol (IP)
- 1985** US National Science Foundation commissions its network
- 1990** First dial-up commercial internet access
- 1991** World wide web is invented

02 The world wide web

Although many people treat them as synonymous, the world wide web is not the same thing as the internet. The web is just one of a number of services that use the internet. Others include everything from file sharing and online games to video chat or email. The web is, however, perhaps the single most important service within modern digital culture, for the world wide web is founded on the principle that anyone with an internet connection should be able to move freely between websites and create their own if they wish. In this sense, it's as much a set of principles as a technology.

The idea for the web was born in 1989, in a research paper by the British engineer and computer scientist Tim Berners-Lee. In it, Berners-Lee outlined his concept for a “universal linked information system” that would “allow a place to be found for any information or reference which one felt was important.” This system would operate via the established structure of the internet—but the key words in Berners-Lee’s scheme were “universal” and “linked.” Anyone had to be able to create sources of information on his system, and these sources had to be created in such a way that it was possible for everyone to find them, use them and move seamlessly between them.

Tim Berners Lee: a digital life

Sir Tim Berners-Lee—knighted in 2004 for his achievements—is revered today as “the father of the web.” Perhaps still more remarkable than his ingenuity, however, is his guardianship of the principles of universality of freedom embodied in his creation. Born in London in June 1955, he studied physics at Oxford before working in telecoms, software engineering and technical design, becoming a fellow at CERN in Switzerland in 1984, where in 1989 he first conceived of the web. Once his brainchild had launched, rather than profit from restricting its use, in 1994 Berners-Lee founded the World Wide Web Consortium: an organization devoted to ensuring the web’s continued openness and the maintenance of its universal standards. He remains one of the world’s leading advocates of the open, transparent use of data, and an advocate of maintaining the web’s evolution through the 21st century as a still more potent force for linking and creating knowledge.

By December 1990, Berners-Lee had—with the help of his Belgian colleague, Robert Cailliau— assembled all the components necessary for a fully functional realization of his proposal at their place of work, the physics research institute CERN in Geneva, Switzerland. Three essential components were involved: the world’s first “page” of digital information (so that there was something to look at); the first “browsing” program, which allowed people to view the information on this page from their own computer terminals; and the first web server, the computer on which the web page was “hosted.” The host computer would function a little like a digital noticeboard: the page of information was posted on it and then anyone with a browser program could simply connect to this host computer and view the page on the noticeboard. And as many people as had browser programs could look at the same time: the page itself simply remained on the host, displayed to all.

Today, even conservative estimates put the number of web pages in existence at over a trillion, a large proportion of which it is now almost impossible to locate—despite the immense power of modern search engines—yet billions of sites are in use and accessible to anyone with a computer and a web browser. It’s an astonishing testament to the power of Berners-Lee’s original vision and to the determined efforts of the company he set up in 1994, the World Wide Web Consortium, in order to maintain common and open standards for the whole of the web.

The first website

The old joke about who the owner of the first telephone was going to talk to doesn’t quite apply to the web, but it comes close. Fittingly enough, then, the first ever website—created in December 1990 by Tim Berners-Lee—was simply a few pages of linked text describing what the “world wide web” was all about. And it had what was also the world’s first web address: <http://nxoc01.cern.ch/hypertext/www/theproject.html>. The world wide web, it explained, is something that “merges the techniques of information retrieval and hypertext to make an easy but powerful global information system.” And you can still view a version of the original site online today via the World Wide Web Consortium’s current site, www.w3.org.

Hypertext HTML or, to give it its full name, HyperText Markup Language, is the set of rules that underpins every page on the world wide web. As the name suggests, it’s like a “hyper” or enhanced version of ordinary writing. This is why every web address today begins with the letters http: they stand for HyperText Transfer Protocol, and refer to the system that allows information in the form of hypertext to be exchanged between computers. The words printed on the pages of a book are text in the ordinary sense: mere words on paper. They become a hypertext when this ordinary text is “marked up” by the simple addition of a small number of “tags” embedded in the electronic version of the text. Each tag contains a special, additional piece of information about a piece of text, telling any web browsing program how it should be presented, what other pages elsewhere on the web it should connect to and so on.

In the first version of HTML, there were just 20 basic ways in which information on a page could be tagged. Perhaps the most crucial idea of all was that each page was to be given its own unique “base address,” and that it would then be possible to specify the addresses of other pages in a form “relative to the address of the current document.” What this meant in practice was that a simple tag could be used to connect any one page to any other, by telling a web browser to create a link between the unique address of one page and another page. Most people using the web today will never see a single line of actual HTML: but in its latest version it remains the underpinning of every website that’s out there.

Web browsers Browsers are the dedicated programs through which users access the web. Today, a host of different browsers are available, including such famous brands as Internet Explorer, Firefox, Chrome and Safari. Part of the genius of the web is that any browser should be able to grant any user access to pretty much every website there is.

Browsers allow users to navigate their way between different websites and web pages, either by following links between them or by entering a specific address. The very first browser was initially

called WorldWideWeb—and later renamed Nexus—and allowed users to do little more than view different web pages and move between them. Over the years, a host of more sophisticated functions have been built into browsers, allowing for more and more complicated effects to be achieved through websites: from the use of complicated style sheets to, today, streaming sounds and video within web pages and even running complicated interactive applications such as games entirely within the windows of a web browser.

“Anyone who has lost track of time when using a computer knows the propensity to dream, the urge to make dreams come true and the tendency to miss lunch.”

Tim Berners-Lee

Web servers Web pages can be viewed by anyone with an internet connection and a web browser, but the code that they are made up of needs to be actually stored on a computer somewhere for it to work. The computers on which websites are run are known as web servers and are responsible for delivering the content of a website over the internet to anyone who wishes to access it through a browser.

What this means is that a web server literally “serves” a copy of a web page to every user whose browser is set to that page’s unique address. If two people do this at the same time, each is simply given a copy of the page to view in their browser; and the same is true for a hundred, a thousand or even millions of people. Provided, of course, the server is powerful enough to keep up with the demand. The more complicated a page or the more people who want to view it, the more powerful the server hosting it needs to be. In the case of the world’s most visited websites, such as Google’s homepage, the servers required to grant everyone access to the site and process the queries they type in run to thousands upon thousands of machines, serving hundreds of millions of user requests.

the condensed idea

the web is open and universal

timeline

- 1989** Tim Berners-Lee first proposes the web
- 1990** First working version of the web
- 1991** Web becomes publicly available via the internet
- 1994** World Wide Web Consortium is founded

03 Internet service provision

Internet Service Providers, or ISPs, are the backbone of digital culture providing individuals or companies with internet access in return for a fee. In the early days of the internet, before the world wide web, it was ISPs who brought to the general public a technology that was still new and daunting to access. Even after the far more accessible world wide web took off, ISPs remained a dominant force, controlling the email and online experiences of millions of internet users. Today, their role is less visible, but they remain an integral part of much of the world's access to online services—and powerful players in any debate over the future direction of the net.

The first Internet Service Providers began to appear in the late 1980s but they did not offer access to what we think of today as the “internet” or even its full early incarnation. The direct ancestors of the modern internet—NSFNET and ARPANET—were at this time reserved for use by higher educational institutions and defense contractors. And so US-based companies such as UUNET (founded 1987) and Netcom (founded 1988) began to offer paying customers access to systems based on the more informal UUCP, or Unix-to-Unix Copy, network: a system that allowed networked machines to call each other individually and exchange information, such as news and mail.

The world wide web did not even exist at this point, meaning that paying for internet access principally allowed users to send and receive emails and to take part in discussion forums such as Usenet, an influential early community of digital words and ideas. Until this point, digital culture outside of major institutions had relied on individuals with computers accessing so-called Bulletin Board Systems (BBSs) by dialing them directly with modems over ordinary phone lines. BBSs continued to grow their user numbers until the mid-1990s, when a combination of widening internet access and the growth of the world wide web finally overtook them.

Offering commercial access to the UUNET in the late 1980s proved a useful source of revenue for expanding network systems and capacities, but much debate continued to surround the commercial opening up of the NSFNET and ARPANET systems. Finally, in 1992, Congress voted to allow commercial operations to begin on the NSFNET. This signaled the beginning of commercial access to the “true” infant internet, which until then had been largely a non-profit scholarly tool.

How fast?

One of the most confusing and potentially controversial areas in internet access is speed: how fast will the connection provided by an ISP download and upload data; and how well will it cope with the increase in online traffic during hours of peak demand? In Britain, for example, ADSL broadband—provided over the phone network—has a theoretical maximum download rate of 8 Mbps. But this is rarely achieved, with rates ranging from 4 to as low as 1 Mbps more common. Some countries, such as South Korea, have invested heavily in dedicated cables providing a far faster national broadband network than others. ISPs there are able to offer an average of around 12 Mbps, compared to Japan's 8 Mbps and Britain's average of just under 4 Mbps. Speeds are

expected to increase dramatically over the coming decade but so are the amounts of data transferred by the average internet user, thanks to the increasing quantity and quality of video streamed online. To meet demand and maintain speeds, many ISPs believe that substantial government investment across the world is required in parallel to the commercial sector's efforts to keep up.

First full access From 1992, ISPs were able to sell “dial-up” access to the full internet as it then existed. Dial-up access is so called because—much as with Bulletin Board Systems—it involves a computer connecting to the internet via an ordinary phone line, attached at one end to the computer's modem, and at the other to a receiving modem attached to the computer belonging to the ISP.

The first company to offer a full dial-up internet service was The World, which did so from August 1992, having offered dial-up access to the UUNET since 1989. Other companies followed in quick succession, and with the launch in 1993 of the first graphical browser for the world wide web, global online growth began to take off.

In 1996, internet service provision saw the emergence of a new term, “broadband,” which appeared prominently in an advertising campaign for American ISP MediaOne. The term has no precise technical meaning, but was used to show the fast internet access rates MediaOne was able to offer via its cable modem service: modems that operated not through ordinary telephone wires, but along the cable television infrastructure, allowing far faster access speeds. Cable modems had been in limited use since the early 1990s, but it was not until the second half of the decade that the notion of broadband access began to take off and constitute a major market sector for ISPs.

America Online

America Online or AOL was one of the most impressive success stories of the early days of the commercial internet. From the late 1980s, AOL offered one of the first and most appealing of the early “walled garden” internet services. For an hourly charge—which became a monthly fee in 1996—users gained access to chatrooms, games, email on the company's own network, an instant messenger service and other proprietary content. The service was extremely user friendly and well marketed, and gained a huge number of users through the 1990s, reaching around 35 million subscribers by 2002. But the model of safe, controlled internet service provision fell apart in the 2000s in the face of the ongoing expansion of the web and the proliferation of free online services. By early 2011, AOL had only around four million subscribers remaining—and was busily reinventing itself as a content-based business, with purchases including the world's most-read blog, the Huffington Post, for over \$300 million.

ADSL arrives Toward the end of the 1990s, a second broadband technology began to become an affordable option for domestic users: an Asymmetric Digital Subscriber Line or ADSL. The technology had existed since the end of the 1980s but had been extremely expensive as it involved using a complex form of digital signal processing to send information along an ordinary telephone wire at a far higher rate than was possible via an ordinary dial-up connection.

All of this has meant a gradual transition in pricing models for internet access from paying by the minute for dial-up to paying a flat monthly fee for broadband access that was always on. ISPs—themselves pay for the internet access they pass on to their customers—usually from a larger ISP, with access to larger sectors of the internet. Ultimately, this money goes to funding the structure of the internet itself, spread between the numerous private companies that maintain the high-capacity lines sometimes referred to as the “backbones” of the internet, as well as the regional networks, consortia and connections that form the vast modern spread of the internet.

“Access to the internet itself has become a human right.”

Michael L. Best

Today, both more varied and more powerful ways of connecting to the internet are becoming widespread: from high-speed cables to more advanced forms of ADSL-style encoding, as well as high-speed wireless access. For ISPs, tasked with maintaining reliable access to the internet for a rapidly growing number of users who are also each using increasingly data-intensive services, the future holds a number of challenges to do with maintaining both profits and service quality in the face of this growth. New business models supported by advertising are being attempted by some, while others are attempting to make money by charging to prioritize certain kinds of traffic—a key issue in the “net neutrality” debate.

the condensed idea

Those controlling internet access wield great power

timeline

- 1987** Limited public paid access to internet
- 1992** First full paid-for net access
- 1996** Term “broadband” is coined
- 2002** AOL subscribers peak

04 Email

Electronic mail—the sending of text-based messages between computers—is one of the first and most fundamental of all digital ideas. In fact, it predates the existence of both the internet and the world wide web, dating back to almost the earliest occasions on which one computer was linked to another. For almost as soon as the idea of connecting computer terminals was possible, people set out to do something that has for millennia driven civilization forward: communicate with each other.

In the early days of computing, machines were huge and extremely expensive, and it was usual for one “mainframe” to have a number of terminals linked to it through which many people could use the single, central computer. From 1961, an early mainframe computer known as an IBM 7090 was in use at Massachusetts Institute of Technology (MIT). This computer had what was then a radical operating system, which allowed several users to log into the computer from different terminals and transfer files they were using onto the computer’s central disk.

It gradually became clear that this ability to transfer files onto a central disk meant that this operating system could in effect be used like a mail box. You could, for example, write a message in a file and then upload it to the computer’s central disk with a filename like “FOR TOM.” This was like posting a letter addressed to a particular person. I would then log onto a different terminal, search the computer’s central disk, see the “FOR TOM” file and thanks to its name would open it to read the message that had been left in it for me.

By 1965, this system was being sufficiently widely used that a specific “mail” command was created on the operating system. This effectively automated the established mailing process, meaning that any user could now send a message to any other specific user, and the file containing this message would be placed in a specific “mailbox” location on the main disk. All you had to know was the unique numbers identifying a particular user and you could send a message to them. This would be indicated the next time they logged on to a terminal by the message “you have mailbox.”

Introducing @ The next significant developments in email came with the creation of the Advanced Research Projects Agency Network (ARPANET) in 1969 at MIT—the precursor to the internet. The rapid expansion of this network hugely increased the use of electronic mail and began to generate further innovations in the mailing process. In 1971 the now-standard “@” sign—usually referred to as the “at” sign—was used for the first time to denote the particular location that someone sending or receiving mail was “at,” an innovation that formed part of a new email system implemented by programmer Ray Tomlinson. Tomlinson’s system was the first to allow messages to be sent between different host computers, rather than simply between different users on the same mainframe computer. In the modern sense, it marked the first true email system.

@ around the world

Until its first use in emails in 1971, the @ symbol was an obscure accounting symbol used to indicate pricing levels in accounting. Since then, it has become one of the world’s most widely

used symbols and has gathered a bewildering and colorful variety of different descriptions in different languages. While in English it is simply called the “at sign,” others are more poetic. In Italian, it is *chiocciola*, “the snail,” thanks to its shape, while the Finnish language thinks it looks more like a curled-up cat (*miukumauku*). Russian leans toward a dog (*sobaka*) and the Chinese sometimes call it *xiao laoshu* or “little mouse.” But perhaps most colorful of all is the German interpretation: *Klammeraffe* or “spider-monkey.”

The next year, 1972, one of the founding fathers of the internet—Larry Roberts, who had led the development of the ARPANET, designed the first fully operational email management program, able automatically to read, respond to, file and manage emails. Numerous other email management programs quickly followed but every one remains essentially the same in the core functions it performs.

Email apparatus

A host of options are now considered essential for any email program. Top of many lists is the ability to send and receive files—documents, photographs, databases and so on—as “attachments” to an email. Address Book functions for managing contact details are also almost universal, as is the ability to send a “carbon copy” (cc) or “blind carbon copy” (bcc) of a message to other recipients, and to forward (fw) messages on to others. Many people also use email signatures, automatically putting their business or personal details at the bottom of each email. Newer email systems like Google’s Gmail have also introduced the idea of “threaded” email exchanges, grouping together emails sent backward and forward with particular people for ease of reading. And tagging and filing systems for keeping track of past emails are increasingly important given the many thousands of messages an average user can expect to build up over their lifetime—not to mention filtering and priority systems for ensuring that important messages get read, and trivial ones automatically set aside for casual browsing.

Email software When an email is sent, it arrives not at an individual’s computer but at the online mail server providing your particular email service. If you’re using a piece of software that can process emails—such as Microsoft’s Outlook—this software will log into the mail server via the internet to look for new messages, downloading a copy of any new messages to your computer.

Many people now use “webmail” rather than a software “client” on their computers, meaning that they access their mail server through a web browser. Webmail programs were first demonstrated in 1995, and offer the convenience of being able to send, receive and read mail wherever there is access to the internet and a web browser. Popular webmail providers today include Google’s Gmail, Microsoft’s Hotmail and Yahoo!’s Mail.

“Email never seemed big at the beginning because there weren’t many computers. It was only as big as the network.”

Ray Tomlinson, Inventor of email

For both webmail and mail checked using a software client, there are two dominant modern protocols.

for getting messages off a mail server: Post Office Protocol (POP) and Internet Message Access Protocol (IMAP). POP is a simpler system, treating the server much like a post office: it tends to connect to a server, check for and download new messages, delete old messages, then disconnect. IMAP offers a more sophisticated process, connecting to the server for longer and allowing multiple pieces of client software on different computers to connect with the central server mailbox and synchronize the status of messages between them. In effect, it allows the remote management of the email mailbox on a server, rather than simply the download and sending of messages.

Overkill Given the sheer volume of email now sent and received around the world—even excluding the vast quantities of spam messages demanding ever more sophisticated filtering techniques on the part of email service providers—inbox management is a vital skill in many modern businesses. Popular techniques for managing emails range from “inbox zero” (no messages allowed to linger in the inbox) to “batching” (setting aside time for dealing with large amounts of email in one intense session). The world’s most-used communication method is at once simple and an ever-shifting art.

the condensed idea

Sending words digitally is its own revolution

timeline

- 1965** “Mail” command is first introduced
- 1971** @ sign is first used
- 1972** First email software
- 1995** First webmail

05 Personal computing

The earliest computers were anything but personal: vast, expensive machines with the reach only of elite academic institutions, major corporations or governments. The gradual democratization of computing power has been one of the most fundamental of all digital transitions, and something that's summarized in the very idea of "personal computing"—a device in every home. It has also created a culture based on a historically unprecedented intimacy between people and machines.

The word "computer" itself began as a term for a person who performed calculations. It was first used to refer to a calculating device in the 19th century, and to an electronic calculating device in 1946. These early machines used vacuum tubes to perform calculations, a bulky and energy-intensive system that filled entire rooms. The first fully programmable electronic computing machine, named Colossus, was built in Britain in 1943 and was in operation in 1944 decoding German messages during the Second World War.

From 1955, electronic transistors began to replace valve systems in computers, leading to the development of the first mainframe computers; in the 1970s, integrated circuits and then microchips replaced these transistors, and for the first time it became possible to think about computing systems whose size, costs, simplicity and energy requirements put them in the range of domestic consumers.

"Ready to slash through long routines and come up with answers in milliseconds. The new Hewlett-Packard 9100A personal computer."

1968 advert in *Science* magazine

The microprocessor revolution The phrase "personal computer" itself was featured in several advertisements in the 1960s, but the first machine truly to deserve the title was designed in 1970 by the Kenbak Corporation. Called the Kenbak-1, it cost \$750 in 1971—the equivalent of around six times that amount today. Less than 50 machines were built.

The Kenbak-1 did not have a microprocessor—just a combination of transistor-based circuits. In 1971, however, the American company Intel produced what's often thought of as the first commercially viable microprocessor, the Intel 8080. Around 500 times faster than the Kenbak-1's integrated circuits, this 8-bit chip was the basis of a 1975 computer, the Altair 8800—manufactured by the young US company Micro Instrumentation and Telemetry Systems (MITS)—that would help transform the public perception of digital technology for good.

"If the automobile had followed the same development cycle as the computer, a Rolls-Royce would today cost \$100, get a million miles per gallon, and explode once a year."

Robert X. Cringely

Priced at \$439 in kit form—or \$621 assembled—over a thousand Altair 8800s were ordered within a month of launch, and thousands more over the course of the year. A culture of modification and accessory-building rapidly sprang up around them. But it was on the software side that the Altair 8800 had perhaps the greatest effect, thanks to the release for it of a programming language called Altair

BASIC—the very first product of a new company then called Micro-Soft which two friends, Paul Allen and Bill Gates, founded in response to the computer’s release.

The Altair 8800 had helped to prove something that surprised almost everyone at the time: computers were not simply of interest to academics and experts; they had the capacity to ignite passions across the human spectrum. Building on this success, 1977 saw three other companies release some of the very first mass-produced machines that would go on to sell millions of units and help to establish digital thinking as part and parcel of daily life: the Apple II, Commodore PET 2001 and Tandy TRS-80.

How fast?

The world’s first personal computer, the Kenbak-1, was capable of performing fewer than 1,000 operations per second on its release in 1971. Four years later, an Altair 8800 could perform 500,000. But this was barely the beginning. By 1982, Intel’s new 80286 chip could operate at 2,660,000 operations per second, while the computers it was installed in had over 100 times more memory for running programs in than the Altair. By 1990, personal computer chips were operating at over 40 million operations per second; by 2000, this had risen to over 3.5 billion; and by 2010 to over 140 billion. And these chips were installed in home computers costing less than 10 percent of the price of a 1980s machine, with around 10,000 times the memory capacity. This continuing story of getting more and more for less is one of the most important ongoing drivers of change—and waste—in digital culture.

Toward the present By the mid-1980s, the global personal computer market was worth billions of dollars, boosted by steadily improving hardware and software, and the gradual move toward Graphic User Interfaces (GUIs), such as Mac OS (released 1984) and Microsoft’s Windows (released 1985). These were operating systems that allowed people to interact with their computers via colored, interactive graphics rather than having to type in commands on a blank screen.

IBM released its first personal computer, the IBM PC, in 1981, and its success profoundly affected the personal computer market by moving it rapidly toward a common standard. Previously, software and hardware had been particular to different models and manufacturers, creating a fragmented market, but the huge success of IBM’s machine led to a rash of “IBM-compatible PCs” being released, and eventually to the acronym “PC” itself coming to mean an IBM-compatible machine rather than any personal computer. The 1980s also saw the emergence of a new kind of portable computer, first marketed in 1981 but with little commercial impact until the late 1980s: the laptop.

Apple continued to operate its own software and hardware lines, but IBM-compatible machines dominated. They gradually gained a host of accessories that made domestic computing ever more powerful and universal: dedicated and increasingly powerful graphics and sound cards, CD-ROM drives from the early 1990s, larger and higher performance monitors and, in due course, internet connections.

Touch-screen tablets

The newest addition to the personal computing family is the “tablet”—a slim, portable device that’s primarily controlled by a touch screen rather than a keyboard. Microsoft first demonstrated a tablet PC in 2001; but it was the release of Apple’s iPad in 2010 that fundamentally shifted the idea of what a personal mobile computer could mean. Running a streamlined mobile operating system derived from that used on Apple’s 2007-released iPhone, the iPad suggested a new kind of ease and intimacy in the use of mobile computing devices, and their emergence as platforms not simply for work and functionality, as laptops have traditionally tended to be used, but as fully-fledged media centers, for browsing, reading, watching, listening and playing. Computing is continuing to get considerably more personal.

By the year 2000, more than 100 million personal computers were being purchased a year; and by 2002 more than half a billion were in use worldwide, rising to a billion by 2008. Today, however, personal computers’ dominance of the digital marketplace is shrinking, as they become simply one increasingly commoditized device among many—no longer many people’s sole window into a new technological world.

the condensed idea

Computers are for everyone

timeline

- 1970** First personal computer
- 1974** First commercial microchip
- 1981** First IBM PC
- 1985** Microsoft Windows is launched
- 1991** First Apple PowerBook laptop
- 2010** Apple launches the iPad

Almost every action taken online involves at least one “server” computer. In the simplest sense, a server is simply a system on a network that can provide services to other systems on that network. Today, thanks to the client/server model that underpins the entire structure of the internet, tens of millions of servers exist around the world—clustered, in the case of the biggest web services, within data centers holding thousands of units. This physical infrastructure consumes immense amounts of resources and gives the virtual world a substantial (and potentially vulnerable) real-world footprint.

In a typical client/server relationship online, the client is an individual computer or program such as a web browser, while the server, also sometimes known as a “host,” is the source of the information or service that the client wishes to use. All internet resources run according to the client/server model, meaning that there are a host of different server types: email servers that people connect to in order to use mail, web servers that operate websites, games servers hosting online video games, and so on. Servers also perform many of the higher level functions that keep the basic infrastructure of the internet itself operating.

One such high-level function is performed by servers that deal with the names of different websites, ensuring that the unique numbered addresses of every site on the web are successfully mapped into the addresses that people are accustomed to reading in their web browsers: www.bbc.co.uk and so on. This mapping function is performed across the web by Domain Name System (DNS) servers, which ensure that internet users can type in a web address into a browser and be successfully directed toward the right site.

Serving by proxy

Using a proxy server introduces an additional link into the client/server chain: a “proxy” machine sits between a user’s computer and the service they are using. This has a range of uses, the most common of which is speeding up internet access: the proxy server caches information that is being used frequently by the client (that is, it keeps copies of it) meaning that not everything has to be sent from the actual server. Proxy servers are widely used by Internet Service Providers and data hubs to speed up traffic. But they can also be used to circumvent filtering or censorship. It may be impossible to access a particular website or online resource because its server details are blocked on a particular network, but using a proxy server as a go-between circumvents this. Proxy servers can also allow people to browse the internet anonymously by concealing their unique IP address, effectively acting as a mask between an internet user and anyone wanting to observe and trace their actions online. Chains of proxy servers can make the original source of web traffic almost impossible to trace.

Among DNS servers, “root name servers”—those that operate the most fundamental part of the hierarchy of the domain name system, its “root” zone—are especially important. These are, in effect,

the nodes that ensure the functioning of the internet. Just 13 different root name servers exist around the world, named from A to M in the alphabet—although there are in fact multiple copies of each for security and stability reasons, with a total of around 250 root name server machines operating worldwide.

“If you hack the Vatican server, have you tampered in God’s domain?”

Aaron Allston

Server chains Simply accessing a page on the world wide web can involve numerous servers, as requests are routed from a web browser on your computer through your domestic router or internet connection, then through different levels of domain name servers implementing the protocols governing web addresses to the server or servers on which the website itself is hosted.

If the website or service you’re using is not on the same local network as your computer, your local DNS server will need to get its information from another DNS server—and will have to contact the root name server to find out which primary and secondary name servers have the relevant information about the web address you entered. Only once the actual numerical address of the website you’re trying to visit has been relayed back to your computer can your web browser then contact the site directly and begin to display it on your screen.

Building a server Servers are physically specialized computers that function very differently to office or domestic machines. Their defining characteristic is that they should almost never need to be turned off: a server needs to be able to respond to requests at all times, and is expected to be operational well over 99% of the time (in the case of many businesses, this expectation can rise up to 99.995%). Reliability is thus a key concern, together with security of power supply and the ability to maintain a safe operating temperature. This often means installing backup and failsafe systems.

Who has the most?

The amount of computing power needed to operate the world’s largest websites and internet companies is staggering. Hard numbers are difficult to come by, and are constantly changing, but in 2010 by some estimates Google was thought to operate more than a million servers in order to process the billions of requests that its websites receive. This might on its own represent around two percent of all the web servers in the world, concentrated into purpose-built data centers at key global locations and consuming millions of dollars worth of power. By comparison, the world’s largest computer chip manufacturer, Intel, had around 100,000 dedicated servers in 2010; and Facebook may also be approaching that number thanks to the enormous speed of its growth. Meanwhile, online retailer Amazon—with tens of thousands of servers of capacity, installed in part to ensure its site never crashes even during peak times of sales fever—is leading the world in renting out its server power to paying customers (see [Cloud computing](#) where this is addressed in more detail).

These factors coupled with the need of most major digital services for a large number of servers mean that it makes sense to operate them in large groups, sometimes referred to as “server farms” or “data centers.” These can contain hundreds or even thousands of machines, physically arranged in racks: high metal cabinets that can traditionally store up to 42 servers stacked up vertically. These servers

are machines stripped down to only basic components; a still more pared-down type of server, known as a “blade” due to its extremely thin profile, has taken this still further, with up to 128 blade machines theoretically fitting into each server rack.

“Today in our data centers we have literally hundreds of thousands of servers. In the future, we’ll have many millions of those servers.”

Bill Gates

Server farms consume large amounts of electricity—more in eventual running costs than the purchase cost of the computers themselves—and generate large amounts of heat, meaning that keeping them cool and allowing air flow around the racks is crucial. This has traditionally meant expenditure on air conditioning, fans and raised construction to ensure air flow, but, with the potential energy savings running to millions of dollars, large companies are increasingly experimenting with basing their servers in locations with cold climates, using advanced ecologically friendly air circulation techniques, and even sea rather than land-based centers.

The energy consumption of servers is an increasingly significant environmental issue. Around two percent of America’s total electricity consumption is already spent on information technology, and with the steady global growth of the internet and online companies, the internet itself is becoming one of the world’s fastest growing sources of emissions, thanks in large part to the massive energy consumption of the servers required to keep it running.

the condensed idea

Servers are the engines of the digital world

timeline

- 1970s** First racked servers begin to be used
- 1990** World’s first web server
- 2001** First blade server shipped

- [click Pastry: A Master Class for Everyone, in 150 Photos and 50 Recipes](#)
- [Five Legs \(A-List\) here](#)
- [click Collected Works, Volume 31: Marx 1861-63](#)
- [The Key to the Indian \(The Indian in the Cupboard, Book 5\) pdf, azw \(kindle\), epub, doc, mobi](#)
- [Analog Science Fiction and Fact \(May 2014\) book](#)
- [click CSS Essentials pdf, azw \(kindle\), epub](#)

- <http://studystrategically.com/freebooks/Pastry--A-Master-Class-for-Everyone--in-150-Photos-and-50-Recipes.pdf>
- <http://omarnajmi.com/library/Designing-Information--Human-Factors-and-Common-Sense-in-Information-Design.pdf>
- <http://yachtwebsitedemo.com/books/Transcendentalism-Overtuned--From-Absolute-Power-of-Consciousness-Until-the-Forces-of-Cosmic-Architectonics--Ana>
- <http://interactmg.com/ebooks/The-Key-to-the-Indian--The-Indian-in-the-Cupboard--Book-5-.pdf>
- <http://fitnessfatale.com/freebooks/The-Last-Rendezvous.pdf>
- <http://diy-chirol.com/lib/Unhinged--The-Trouble-with-Psychiatry---A-Doctor-s-Revelations-about-a-Profession-in-Crisis.pdf>