1

A brief history of the Internet

The Internet is the result of the bold effort of a group of people in the 1960s, who foresaw the great potential of a computer-based communication system to share scientific and research information. While in the early times it was not a user-friendly environment and was only used by a restricted community of computer experts and scientists, nowadays the Internet connects more than one hundred million hosts[1] and keeps growing at a pace unknown in any other communication media. From this perspective, the Internet can be considered as one of the most representative accomplishments of sustained investment in research at both the basic and applied science levels.

The success of the Internet is due to its world-wide broadcasting capability that allows the interaction between individuals without regard for geographic location and distance. The information exchanged between computers is divided into data packets and sent to special devices, called *routers*, that transfer the packets across the Internet's different networks. Of course a router is not linked to every other router. It just decides on the direction the data packets take. In order to work reliably on a global scale, such a network of networks must be very slightly affected by local or even extensive failures in the network's nodes. This means that if a site is not working properly or it is too slow, data packets can be rerouted, on the spot, somewhere else. Surprisingly, such a network communication system is realized by a complex interplay of communication protocols, hardware infrastructures, and connectivity architecture that is the outcome of an evolution lacking any central authority. Things have always been changing on the Internet, sometimes gradually and sometimes very rapidly, but always evolving without a precise general design. The Internet is in this sense a major example of a *self-organizing system*, combining human needs and technological capabilities in a cooperative way.

[1] A host is defined as a computer that allows individual users to communicate with other computers through the Internet.

In this chapter we want to provide a brief history of the Internet's development and growth, introducing at the same time the basic terminology and concepts used to describe this system. The main scope of the following pages is to allow the reader to understand the complex nature of the evolution of the Internet, and to highlight its most important technological and organizational ingredients. From this perspective, and for the sake of simplicity, the present chapter is necessarily a sketchy presentation of the Internet's history. The interested reader can dig deeper into this subject in the works of Gillies and Cailliau (2000), Abbate (2000), the Hobbes' Internet Timeline (2000)[2], and the online history by Leiner *et al.* (2000)[3].

1.1 The early times

1.1.1 Distributed communication networks

The history of the Internet began in the early 1960s, at the height of the cold war, when ARPA,[4] an agency created in 1958 by the US Department of Defense to sponsor research projects related to military problems, started to fund programs at universities and corporations concerning the creation of a computer network to access and share data and programs among computers located in different places. This was an appealing project, specially from the perspective of providing secure control over information in the event of large-scale international conflicts.[5]

It is important to recall that at that time computer communications were only point to point, with each network link depending upon the link before it. In such a structure, any part of the system could easily be isolated by knocking out just one of the links. A different and obvious topology for a computer network in the 1960's was a highly centralized one, in which all computers were connected to a central unit handling all data exchanges, the so-called "star-shape" topology. Those centralized systems, however, are even more vulnerable to attacks and failures, since knocking out the central node is enough to disconnect the whole network.

It is possible to devise several other regular topologies, such as those represented in Figure 1.1, though their respective vulnerability to external attacks or internal failures was uknown. With this prospect at hand, Paul Baran at RAND corporation was given a US Air Force grant to investigate how the US military could protect its communication systems from serious damage. In his conclusions, Baran outlined the principle of "redundancy of connectivity" and explored various models for designing communication systems and evaluating their vulnerability. In the ensuing

[2] http://www.zakon.org/robert/internet/timeline/
[3] http://www.isoc.org/internet/history/brief.shtml.
[4] Advanced Research Project Agency. ARPA was renamed DARPA, the Defense Advanced Research Project Agency, in 1972.
[5] Despite the fact that the later studies on survivability and robustness considered explicitly the nuclear threat (Baran, 1964), the initial ARPA was not officially related to building a network resistant to nuclear war.

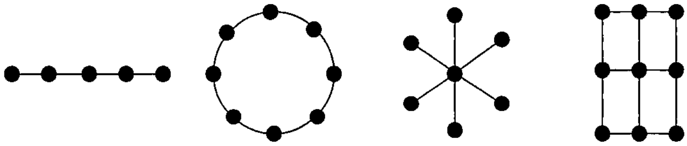

Fig. 1.1 Examples of regular topologies: from left to right, linear, ring, star, and mesh.

report, Baran (1964) proposed a distributed communication system, in which there would be no obvious central unit, every node having the same routing capabilities. With such a configuration, all points surviving a large-scale attack would be able to maintain the contact with the surviving part of the network.

The best design for such a decentralized network is obviously a highly interconnected, distributed network in which each node is connected to all the others. This is a fully connected network with the highest degree of redundancy. The first plan was indeed to connect the mainframe computers at each site directly to all others, but it was soon realized that such a level of redundancy was far too complicated and expensive to be handled. A different approach was thus pursued, designing a *distributed network* with sufficient redundancy, but far from a fully connected topology.

1.1.2 Packet switching technology

The technology capable to handle communications in a distributed network was provided by the work on data packet switches initiated by several groups in the mid-1960s.[6] The group led by Frank Heart, at the Bolt, Beranek, and Newman (BBN) corporation, was awarded, by ARPA in 1968, a contract to develop the Interface Message Processors (IMPs), small machines – the ancestors of modern routers – which were designed to be a part of each mainframe dedicated to form the subnetwork between computers. The IMPs used a technology called *packet-switching*, which parcels data in small chunks called packets, labeled with the destination address. Packets can be sent in any order through any path leading to the same destination and reassembled on arrival. The advantage of a packet-switching system is evident: in a centralized network, such as a star-shaped system, all the information is channeled to the central unit, to be processed and redistributed. In a distributed network with packet-switching technology, however, each node has the authority to originate, pass, and receive messages. In particular, if a node is not working, or working too slowly, a packet can be rerouted through some other nodes. This dynamical rerouting implies that each packet finds its own way through

[6] The first report on packet switching theory, dated in 1961, was written by Leonard Kleinrock while a graduate student at MIT.

the network and allows all nodes to be equivalent, i.e. *peers*. With every node having the same routing capabilities, the network can fully exploit its connectivity redundancy: Only a failure affecting nearly all computers can disable communications over the whole network.

1.1.3 The ARPANET

In 1966 the MIT researcher Lawrence G. Roberts started for ARPA the design of ARPANET, a network initially intended to wire up four major mainframes at universities in Southwestern US: the University of California at Los Angeles (UCLA), the University of California at Santa Barbara, the Stanford Research Institute (SRI), and the University of Utah. In 1968 the ARPANET specifications were laid down, based on the IMP technology, and the first host-to-host message was sent from UCLA to SRI in October 1969. The four planned mainframes were connected by the end of 1969, forming ARPANET, the precursor of the present Internet.

During the 1970–71 period, ARPANET grew to 23 nodes, while work proceeded quickly on designing a functionally host-to-host protocol called Network Control Protocol (NCP), which became fully implemented in 1971–72. Once a reliable working protocol was established, researchers started to develop applications, and in 1972 Ray Tomlison at BBN wrote the basic e-mail message software, which with time has become probably the most widely used application on the network. The same year the Internetworking Working Group (INWG) became the first standard-setting entity to govern the growing network. In 1973 the first international nodes were set up in England and Norway. In 1977 the ARPANET encompassed 107 nodes. Its growth was slow but steady, the scientific community recognizing that the new communication network was going to become something wider than they had ever imagined. This fact called for designs and technologies intended for a larger network and the Internet took off.

1.2 The rapid growth

1.2.1 More networks

In a short time, the ARPANET example was absorbed by other communities, such as the US Department of Energy, which established the MFENET for its researchers in magnetic fusion energy, and the HEPNET for high energy physicists. NASA space physicists followed with SPAN, enlarging the number of purpose-built networks for academic and research communities. Private companies soon also made their first move into the electronic world when BBN opened in 1975 Telenet, a commercial network based on the ARPANET model.

At the same time, other networking technologies were being developed by the computer science community. One such alternative was the CSNET (Computer Science Research Network), providing networks for computer science departments. CSNET benefited specially from the dissemination of the Unix operative system and its built-in Unix User Control Protocol (UUCP). In the early 1980s other important networks based on this technology, such as USENET (Unix User Network)[7] and BITNET (Because It's Time Network), sprang up. The reason for the surge in UUCP-based networks was that UUCP, modems, and the existing telephone lines were a ready-to-use technology of data transport. Also, computer facilities and institutions which were not part of the ARPANET were increasingly aware of the benefits of belonging to a large linked computer system. UUCP networks were, however, rather different than the ARPANET, providing essentially store-and-forward facilities and e-mail, in order to create discussion groups and provide two-way and one-to-many communication. In other words, the user's computer connects to another machine which has filed users' postings in separate topical discussion groups (newsgroups). The user may issue a command requesting the full text of a particular posting, post a follow up or even start a new newsgroup. These e-mail discussion lists constituted another major element in the building of the Internet community, heavily contributing to the international growth of the internetworking principles.

1.2.2 The TCP/IP development

The launch and growth of several different networks was a further stimulus for the development of the key technical idea underlying the modern Internet, namely that of *open architecture* networking. In an open architecture network, the individual networks may be designed in accordance with specific requirements that can be freely selected by each administration entity. Each network, however, communicates with the other networks through a set of protocols that are the same regardless of which network the user or service operates in. In other words, each network stands on its own and no internal changes are required to connect it to the Internet. In this respect, the NCP communication protocol used in ARPANET was not a viable solution, since it did not deal with end-to-end host errors and destinations different than the IMPs on the ARPANET. The first effort to develop an open-architecture network was led by Robert E. Kahn at BBN and Vinton G. Cerf at Stanford. Their research yielded as a final result the TCP/IP protocol. The TCP,

[7] USENET was developed by T. Truscott and J. Ellis while graduate students at Duke University and North Carolina University, respectively. At first it was just a graduate students activity but eventually grew until accommodating a link to join the ARPANET mailing list. Doubtless, USENET has been one of the main factors driving the physical and social self-organized nature of the Internet.

or *Transmission Control Protocol*, converts messages into a bunch of packets at the source and reassembles them back into messages at the destination. The IP, or *Internet Protocol*, however, handles the addressing and routing of single packets across nodes and different networks, providing a unique address space for the Internet. The TCP/IP was adopted as a standard of the US Department of Defense and was shared by other agencies, and in 1983 ARPANET experienced a complete transition from NCP to TCP/IP. The new protocol allowed the original ARPANET to split into two different networks, its military part, MILNET, and an ARPANET supporting only research needs. At the same time the TCP/IP, which was public-domain software, was adopted by other networks, such as the CSNET, to route information exchange with the ARPANET. As the use of the TCP/IP became more and more common, other entire networks began to interconnect and the network of networks was finally born.

1.2.3 The NSF acceleration

Another turning point in the history of the Internet was the National Science Foundation (NSF) program to establish a new transcontinental network and five super-computing centers. The program was driven by the idea that networking and computer resources were indispensable tools for the research community, and the NSF philosophy was to serve the entire higher education community, regardless of discipline. In order to allow all institutions to link up to the network, NSF agreed to pay for the establishment of a connection to its high-speed network (the *backbone*) only if universities provided connections to smaller educational and research institutions. A critical decision of the NSFNET (NSF Network) program was the mandatory use of the TCP/IP protocol, triggering the marginalization of other wide-area network protocols. Another NSF decision, which contributed to the shape of today's Internet, was to encourage commercial network traffic at the local/regional level along with the ban of the commercial use of the NSFNET backbone. This stimulated the establishment of long-haul private networks (national-scale providers) offering alternative backbones for the commercial traffic.

The NSFNET had an enormous impact on the Internet's evolution. The massive funding and the establishment of policy guidelines triggered the transition from a small network limited to the research environment to a full-scale network of networks with solids links also in the commercial community. The hardware technology received an impressive thrust as well, and the TCP/IP was put on its way to becoming the standard of the present. NSFNET co-opted the ARPANET in 1989 and finally, in 1995, reverted back to a pure research network, leaving the national-scale connectivity to private backbone providers.

Cambridge University Press
978-0-521-82698-3 - Evolution and Structure of the Internet: A Statistical Physics Approach
Romualdo Pastor-Satorras and Alessandro Vespignani
Excerpt
More information

1.3 The network of networks: a growing self-organized system

1.3.1 A large-scale infrastructure

From 1990 the Internet has experienced explosive growth. The number of hosts nowadays are counted in tens of millions and new networks and providers are connecting to the Internet on a daily basis (see Figure 1.2). The reasons for this phenomenon are partly due to the fact that personal computers have become an household item, and partly to the advent of the World Wide Web (see Chapter 7), which allows easy access to the huge amount of information stored in the Internet. This increase in scale of the Internet necessarily introduced several new concepts and changes in the underlying technologies. The shift from a few networks with a modest number of hosts to a large number of networks with a wide range of connected hosts has led to the definition of three possible classes of networks. Class A represents large national-scale networks (there are a few of these networks and they have a large number of hosts). Class B networks are regional-scale networks. Class C comprises small local area networks with a limited number of hosts.[8] Since hosts have a unique numeric address, a related name is assigned to each of them, to make it easier for people handling the addresses.

However, with millions of hosts it is impossible to have a single table of all the hosts and their associated names and addresses. This has led to the introduction of the Domain Name System (DNS) that allows the resolution of host names into

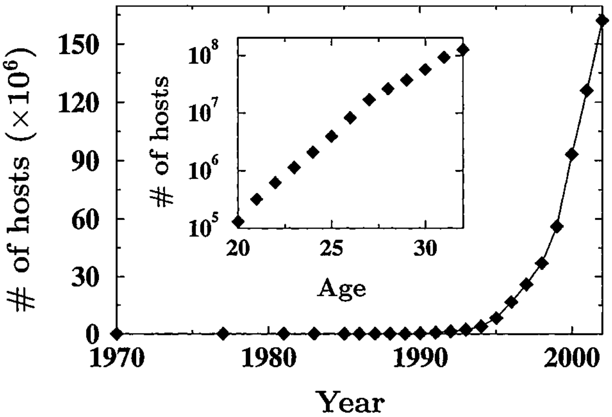

Fig. 1.2 Number of hosts in the Internet starting from 1970. The inset reports the number of hosts as a function of the Internet's age in a logarithmic-linear scale. The linear behavior shows that in the last ten years of life the Internet has grown exponentially, doubling annually. Data from the the Hobbes' Internet Timeline http://www.zakon.org/robert/internet/timeline/.

[8] The class subdivision results in IP addresses allowing different numbers of networks and related hosts, as explained in Chapter 2.

Internet addresses. The same sort of problem has occurred also for routers. Originally, all routers implemented a single distributed algorithm for the routing of data packets (a long list of paths to all addresses). As the number of networks grew, a hierarchical model of routing split the protocol into an Interior Gateway Protocol (IGP), used inside Internet regions, and an External Gateway Protocol (EGP) used to link different regions. With time, new problems related to Internet scalability are faced. New ideas continue to pop up as well, and new solutions to be envisioned. The Internet behaves in this sense as an evolving system, whose appearance is changing over the various phases of its life.

1.3.2 *Self-organization and cooperation*

During its growth the Internet has changed continuously, surviving dramatic changes in technology and evolving to accommodate the exponential increase in users. A key element for these "adaptive" features of the Internet is the dynamic exchange of ideas and the open access to technical standards. A clear example of this attitude is the Request For Comments (RFC) series of notes.[9] Originally these memos were intended to be a fast way to share and distribute ideas, with online files accessible via the File Transfer Protocol (FTP). The very early RFCs presented ideas developed by groups of researchers to the rest of the community, or information on protocols and engineering issues. Nowadays, RFCs are more focused on Internet protocols and are viewed as the "documents of record" for Internet standards. RFCs have a centralized administration for the required protocol number assignment. Standards are set by the Internet Engineering Task Force (IETF), that has working groups on different aspects of Internet engineering.[10] IETF is not, however, a formal group of people but remains open to anyone interested in participating. Indeed, each working group has a mailing list where draft documents are discussed, and only when a large consensus is reached are the documents distributed as RFCs. Once a standard is set, it becomes a sort of commandment on the Internet, since this is the only way to ensure that people using different hardware and software can communicate. This illustrates fairly well the level of self-organization present in the Internet community.

Another amazing example of cooperation in the Internet community is how networks are physically connected together. No governing office determines how routers must be connected. Each network organization, from small campuses to large national-scale providers, decides autonomously its connections and makes arrangements to pass along other's network traffic. From this perspective, it is natural

[9] http://www.rfc-editor.org/.
[10] The IETF has grown in time and resulted in further substructure, such as the Internet Steering Group (IESG) formed by working group directors.

that non-functional network connections (excessive loads, spam e-mail, denied services) are in a short time cut-off by their peers. In other words, a working Internet requires network managers to cooperate, and, despite its seemingly chaotic development, the Internet acts as an efficient communication medium.

The functioning of the Internet also requires large-scale control of many complicated organizational issues and many important entities set the proper guidelines and policies. Examples can be found in the registration authorities that regulate the assignment of domain names and IP addresses,[11] or the handling of security-related problems. As we have shown, however, not one of the organizational or technical groups has ever plotted a global project of the Internet. The Internet is not driven by any supervising agent or authority, nor follows the blueprint of a pre-established architecture. It grows and develops because of cooperation and self-organization, to conform to technical standards and associative needs. Indeed, if we look at the Internet on a coarse grained scale, we see a spontaneously growing system, whose large-scale dynamics and structure are a cooperative effect due to many interacting units aimed at optimizing local communication efficiency.

[11] There are currently four Regional Internet Registries: the Réseaux IP Européens Network Coordination Centre (RIPE), the American Registry for Internet Numbers (ARIN), the Asia Pacific Network Information Centre (APNIC), and the Latin American and Caribbean IP address Regional Registry (LACNIC).

Cambridge University Press
978-0-521-82698-3 - Evolution and Structure of the Internet: A Statistical Physics Approach
Romualdo Pastor-Satorras and Alessandro Vespignani
Excerpt
More information

2

How the Internet works

While the scope of this book is to look at the Internet as a self-organizing complex system and to study its large-scale properties by using a statistical approach, a general knowledge of how the Internet works is necessary to identify the main elements forming the Internet, as well as their respective interactions. To give a physical analogy, if we want to understand the properties of a certain material we first have to know the atomic elements it is composed of, and how they interact. Similarly, it is impossible to approach the Internet without first having some hint of what a router is and how it communicates with its peers.

In this chapter we provide a brief description of the different elements, both hardware and software, at the basis of the Internet's functioning, to allow the non-expert readers to familiarize themselves with the general mechanisms that make it possible to transfer data from host to host. These mechanisms will turn out to be very relevant for understanding problems related with measurement infrastructures and other communication processes taking place in the Internet. Our exploration of the Internet's workings will be necessarily non-technical and, needless to say, the expert reader can freely skip this chapter and use it later on for convenient reference.

2.1 Physical description

In providing a general picture of the Internet, the starting point is the concept of a *computer network.* A network is any set of computers – usually referred to as *hosts* – connected in such a way that each one of them can inter-operate with all the others. The connection among hosts is made possible by two major components: hardware and software. The hardware refers to the physical components of the networks, such as computers and communication lines, ranging from the local telephone lines to fiber optic cables, and even satellite connections, that transfer