

# DRAFT: NOT FOR CIRCULATION

## Chapter 1

### The origins of cyberspace

David J. Pym<sup>1</sup>

#### Abstract

Cyberspace is a romantic term, introduced in the elegant science-fiction writing of William Gibson, but the concepts that make up the environment called cyberspace are the stuff of real science, with origins that can be traced to ancient Greece. Much has been written about the origins of cyberspace, including a comprehensive sourcebook by Hook and Norman. In this chapter, I try to take a rather conceptual view of what constitutes cyberspace, tracing the origins of the ideas from 4<sup>th</sup> Century BCE Greece to the modern Internet-supported interaction space — throughout my discussion, I will seek to elucidate the concept of ‘space’ and how it helps us think about the cyberworld. On the way, I consider the literary origin of the word and the mathematical and logical theory that is required to build models of cyberspace.

#### 1. Defining cyberspace

According to the *Oxford English Dictionary* (OED) [42]:

cyberspace |'saɪbəspeɪs|  
noun [mass noun]  
the notional environment in which communication over computer networks occurs. *I stayed **in cyberspace** for just a few minutes.*

According to *The Oxford Dictionary of Science Fiction* [26]:

the entirety of the data stored in, and the communication that takes place within a computer network, conceived of as having the properties of a physical realm; ... .

My purpose here is to unpack and explain these definitions, which are wholly consistent with each other. In contrast to the approach of Hook and Norman [18], who achieve an enormously impressive coverage of relevant material, my guiding principle is to explore the sense in which cyberspace is a ‘space’, a concept that is well-understood in mathematics and physics, and the understanding of which in those fields is alluded to in the term ‘cyberspace’.

---

<sup>1</sup> David Pym is Professor of Information, Logic, and Security at UCL and is a Turing Fellow at The Alan Turing Institute, London.

# DRAFT: NOT FOR CIRCULATION

The modern world is more-or-less wholly dependent for its operation on networks of communicating computers. These computers come in all shapes and sizes. They may be small devices embedded in everyday objects such as watches, household appliances, or cars, or personal laptops and workstations, or vast datacentres supporting the infrastructure of cloud computing. What is central to their function is communication — that is, the transmission of data — over local networks, wider corporate or government networks, or the Internet itself (<https://en.wikipedia.org/wiki/Internet>).

In more detail, the Internet consists of a global network of networks which connect computers around the world and which use a collection of communications protocols, including TCP/IP (Transmission Control Protocol/Internet Protocol, the basic language that the connected computers use to communicate with one another) and others such as OSPF (Open Shortest Path First), BGP (Border Gateway Protocol), and RIP (Routing Information Protocol) which together describe how *data* should be decomposed into packets (the basic units of data handled by the protocol), addressed, transmitted, routed, and received. Data (which can be treated as a singular, plural, or mass noun) consists of blocks ('bits' and 'bytes') of binary numbers. Data itself has no inherent meaning, but data is used to *represent* features of physical or abstract worlds; that is, *information*.<sup>2,3</sup>

When thinking of the Internet, it is often tempting to conflate it with the World Wide Web (WWW). While this is quite understandable in many ways — and the distinction is often blurred in common discourse — it is a conceptual mistake. The WWW — with its language of Universal Resource Locators, or URLs — organizes data-representations of information in a highly structured way and is just one of many applications that are supported by the Internet.<sup>4</sup>

---

<sup>2</sup> It is not my purpose here to explore the definition of 'information'. The literature on the subject is substantial. For our purposes, the usual understanding of a generally well-educated reader — see, for example, the definition provided in the OED [42] — will suffice. For a philosophically sophisticated discussion that is beyond the scope of this article, see, for example, Floridi's *The Philosophy of Information* [14]. See also Zins' *Conceptual approaches for defining data, information, and knowledge* [41].

<sup>3</sup> TCP/IP: [https://en.wikipedia.org/wiki/Internet\\_protocol\\_suite](https://en.wikipedia.org/wiki/Internet_protocol_suite);  
OSPF: [https://en.wikipedia.org/wiki/Open\\_Shortest\\_Path\\_First](https://en.wikipedia.org/wiki/Open_Shortest_Path_First);  
BGP: [https://en.wikipedia.org/wiki/Border\\_Gateway\\_Protocol](https://en.wikipedia.org/wiki/Border_Gateway_Protocol);  
RIP: [https://en.wikipedia.org/wiki/Routing\\_Information\\_Protocol](https://en.wikipedia.org/wiki/Routing_Information_Protocol).

<sup>4</sup> The distinction between the WWW and the underlying Internet is an essential one, but there would be no WWW in its current form without the underlying communications architecture. Some of the key concepts of the WWW — for example, the concept of *hypertext* and the underlying ideas of *distributed systems* [12,33] — have, as described in Berners-Lee's research proposal [6] and book [7], prior histories in computer science. The linking of hypertext to the Internet through 'http' (HyperText Transfer Protocol) and URLs and HTML (HyperText Markup Language) are Berners-Lee's definitive contributions.

# DRAFT: NOT FOR CIRCULATION

The distinction between data and information — which is directly analogous to, though not identical to, the distinctions between syntax and semantics that are made in logic and linguistics — is very important in our context. While it is data that the Internet processes, and which is used in the WWW to represent information in structured ways, it is information to which human beings relate.

Indeed, the WWW is often described as being an ‘information space’, but what is the concept of a *space* that is being invoked here? In fact, the concept of space is quite delicate. It is in mathematics — in particular, in geometry and topology — that it has been richly developed. The idea starts with the familiar three-dimensional environment in which everyday objects have relative position and direction and its more-or-less intuitive generalization to the four-dimensional environment, often called ‘space-time’ in which such objects also have relative position in time.

Although a formal mathematical definition is not needed for our purposes, the mathematical concept has quite strongly influenced the informal concept that I shall need.

According to the OED [42]:

space |speɪs|  
noun [mass noun]

...

Mathematics: a mathematical concept generally regarded as a set of points having some specified structure. One of the most important examples of the mathematical concept is that of *topological space*.

Topological spaces are a way of describing geometrical properties and spatial relationships that are unaffected by the continuous change of shape or size of figures. An important example of a topological space is *metric space* in which there is assigned a distance — with, essentially, the familiar intuitive meaning of physical distance — between each pair of points.

These ideas of proximity translate not only to the network architecture of the Internet, but also to the data-representation of information in the WWW (though the situation is quite complicated and here I am simplifying matters greatly).

For our present purpose, the concept of space that is useful, and which builds on the mathematical ideas mentioned above, derives from a key concept in computer science, namely *distributed systems* [11].

According to the OED [42]:

distributed system |də'strɪbjʊəd ,sɪstəm|  
noun  
a number of independent computers linked by a network.

Examples of distributed systems include the following:

## DRAFT: NOT FOR CIRCULATION

- The Internet itself — a vast collection of interconnected networks of computers. Individual computers connected to the Internet interact by passing messages, which they do by employing a common means of communication (which will be described later on in this chapter).
- Intranets — localized parts of the Internet that are managed by identified organizations which, typically, enforce local management and security policies to control access and use. Intranets are connected to the wider Internet by special-purpose computers called routers which also employ the common means of communication.
- Cloud computing infrastructure — vast data centres consisting in hundreds of thousands of servers provide storage and computer services for vast quantities of data that is fed from the intranets of large numbers of clients that are logically and physically widely distributed. A key problem here, and in distributed systems more generally, is to maintain the consistency of the different copies while simultaneously maintaining the robustness of the service; see, for example, the CAP (Consistency, Availability, Partition-tolerance) theorem in the theory of distributed systems, also known as Brewer's theorem ([https://en.wikipedia.org/wiki/CAP\\_theorem](https://en.wikipedia.org/wiki/CAP_theorem)), which establishes that only two of the CAP properties can be maintained simultaneously. Handling this problem is important in maintaining the experience of cyberspace.
- Mobile and ubiquitous computing (and the Internet of Things) — laptops, phones, cameras, and wearable devices such as watches and spectacles, as well as cars, domestic appliances, smart meters, and electricity substations; indeed, almost everything on which modern society depends are integrated into distributed systems. They reside on local networks that communicate with other local networks and devices, either directly ('peer-to-peer') or via network-based servers, all employing the common means of communication.
- The global banking system — the intranets belonging to each of the world's banks must not only provide services locally (logically locally, even if not physically locally) to the banks' own customers but must also communicate with one other in order to support the transactions upon which the world's commerce depends.
- Online games with multiple players — each player's local, or home, computer runs a client copy of the game, which communicates with the central game server, which communicates with other players' local games, and coordinates the overall interaction between all of the players.

From a mathematical perspective, the basis of computer science and to which I shall return later, the key concepts of distributed systems are the following [2,5,10,11,12]:

- *Locations*: a collection of linked places, be they physical or virtual, that constitutes the basic architecture of a system. Individual computers, file stores, and so on exist at locations within a distributed system, but locations are also the places within computers at which the CPU (Central Processing Unit) and other components reside;

## DRAFT: NOT FOR CIRCULATION

- *Resources*: the entities that a system uses — consumes, creates, moves — in the course of its operations. Examples of resources include the memory locations at which data is stored, the processor cycles available to perform computations, and human operators required to manage and maintain systems; and
- *Processes*: the collection of activities, which are mostly concurrent, that constitutes a system's operations, and so delivers its services. Examples of services include a bank's customer-facing website, streaming films, and the multitude of system-level services provided by a computer's operating system in order to perform computations, manage the keyboard and screen, manage a computer's memory and storage, send and receive email, etc.

Additionally, a specific system, described using these components, resides within an *environment*, and resource transfers between the system and its environment characterize the service that the system provides.

So, here the relevant 'space' consists in the distribution of resources around the locations of the system and the presence of processes that manipulate those resources. This definition might seem a bit restrictive, but, in the distributed systems metaphor, the presence and activities of a human being (using a computer, posting to Facebook, downloading a file, and so on) simply amounts to the presence of a process.

The term 'cyberspace' is derived from 'cybernetics' and 'space', and the meaning of the term depends essentially, though quite implicitly, on the distributed systems metaphor. The term 'cybernetics' was introduced in the late 1940s by Norbert Wiener [38,39].

According to the OED [42]:

cybernetics |sʌɪbə'netɪks|  
plural noun [treated as sing.]  
the science of communications and automatic control systems in both machines and living things.

ORIGIN 1940s: from Greek *kubernētēs* 'steersman', from *kubernan* 'to steer'.

There are several aspects of this definition that are important for the idea of cyberspace. That it refers to 'automatic control systems' may perhaps seem rather restrictive, but I think that should be seen as a consequence of the perspective of the age of the definition: in the 1940s, although the idea of automated control of machines was well understood, the scope of the information technological revolution that was to come had not been anticipated. It mentions also communications. As I have described, the concept of communication, and a common means of supporting it, is a key aspect of distributed systems.

So, now I have all the components I need to give a definition of *cyberspace*.

# DRAFT: NOT FOR CIRCULATION

According to the OED [42]:

cyberspace |'saɪbəspeɪs|

noun [mass noun]

the notional environment in which communication over computer networks occurs. *I stayed **in cyberspace** for just a few minutes.*

Let's try to understand this rather concise definition using the concepts we've considered so far. First, what are 'computer networks'? The appropriate metaphor here — which I have discussed at some length above — is that of distributed systems. Computers, be they servers, workstations, laptops, phones, controllers embedded in cars, aeroplanes, or refrigerators — or even entire data centres — are resources that reside at locations.

Second, what does 'communication over' mean? Computers residing at locations communicate with other such devices residing at other locations using wired and wireless connections. These connections transfer data between located devices using the TCP/IP protocol.

Last, what is meant by the 'notional environment'? It seems that this is where the presence of human interpreters becomes essential. The distributed systems metaphor completely accounts for the infrastructure and its processing of data, so the 'notional environment' can only be something that is experienced by the users of the infrastructure.

Users provide the interpretations of data and its movement around the infrastructure that constitute the 'information environment'. Now, in principle, every data item is a discrete entity and the collection of all such items in the (albeit vast, massively interconnected) infrastructure of the Internet is finite and so can be counted.

From the perspective of the users, however, things look very different. This is for two reasons. First, end-users (as opposed to users who are systems professionals) primarily perceive information, not data. A picture received on a phone may be a finite collection of pixels, but it represents an image of the physical world of substances and qualities. Second, the exchange of information mediates communication between humans, in a shared social space that is created by the technology and its users, and that communication is almost never wholly captured by the data that is exchanged.

Only with this last component is the definition of cyberspace — and its characterization as a mass noun — really meaningful.

The origin of the word 'cyberspace' does not lie in hard science. Rather, it was coined in science fiction, by William Gibson, first in a short story 'Burning Chrome' [15] in 1982 and reused a little later in his celebrated novel 'Neuromancer' [16] in 1984. Not only does Gibson introduce the term, but he also offers a definition:

# DRAFT: NOT FOR CIRCULATION

‘A consensual hallucination experienced daily by billions of legitimate operators, in every nation, by children being taught mathematical concepts ... . A graphic representation of data abstracted from the banks of every computer in the human system. Unthinkable complexity. Lines of light ranged in the nonspace of the mind, clusters and constellations of data.’

It can be seen that Gibson’s definition, albeit expressed in a novelist’s style, anticipates more-or-less all of our analysis: a highly complex distributed system, the representation of data, and the presence of human minds. Indeed, it seems to capture very directly the experiences of humans who engage in ‘immersive’ or ‘virtual reality’ games with other players who may be physically located in many distributed locations, but who together inhabit a shared environment of data that they collectively, and consistently, interpret as the ‘world’ of their game.

Cyberspace is also an important component of conflict in the modern world [28,34] and, consequently, the world’s military and defence agencies have considered the significance of what is increasingly known as the ‘cyber battlespace’ for their strategies and operations. Indeed, some of them have even attempted to formulate their own definitions; for example, the US Department of Defense in 2008 [30]:

‘a global domain within the information environment consisting of the interdependent network of information technology infrastructures, including the Internet, telecommunications networks, computer systems, and embedded processors and controllers’.

Here it is important to understand that the use the term ‘domain’ refers not to that which is usual in computer science, where it describes a collection of addresses within the Internet, such as everything with a ‘.com’ or ‘.uk’ suffix, but rather it refers to a domain of warfare, the other four being land, sea, air, and space. Again, it is clear that a distinguishing feature of the cyber domain is its combination of the virtual/digital and the physical.

In summary, then, what have I described about cyberspace so far?

- First, that it is a concept that builds on the physical and logical infrastructure provided by the Internet.
- Second, that while the Internet processes data, it is information — that is, interpreted data — that is the medium of cyberspace.
- Third, that the interpretation of data, and the processing of information, is performed by humans, who are themselves essential components of cyberspace. Together, the human participants inhabit the shared *social space* that is an essential component of cyberspace.
- Fourth, that there is an essential interplay between — indeed, a merging of — the physical and the virtual.

# DRAFT: NOT FOR CIRCULATION

So far, I have concentrated on unpacking the concepts that constitute cyberspace, and their associated language. These concepts have, however, a substantial back-story through human history, and it is long and rich.

*Before embarking on the story, I should note that it is not possible in a short chapter such as this one to represent fully and acknowledge all of what is a vast literature. Accordingly, the sources I reference are intended only to be suggestive of the literature, and I apologize unreservedly to anyone who feels unjustly treated. I note also that I am not a professional historian and I make no claim to historical completeness in this article.*

## **2. Cyberspace in the ancient and early modern world: beacons and semaphores**

Travelling between widely separated cities, by walking, riding horses, and sailing in ships takes a long time. Messages sent by these means are therefore slow to arrive. For example, during the negotiation of the Treaty of Westphalia in Münster and Osnabrück in 1648, it took two weeks for a letter to reach Stockholm.<sup>5</sup> Consequently, governments and others throughout history have sought ways of communicating more rapidly.

Perhaps the simplest form of rapid long-distance communication is the *beacon*, a fire lighted on hill to give warning of, say, an approaching enemy. A sequence of beacons on a chain of hills can give rapid warning over long distances — it takes just a few minutes to light a fire, and the signal then travels at the speed of light — but, of course, the language of communication is rather restricted.

In the 4<sup>th</sup> Century BCE, the Greek military strategist Aineias Taktikos described a partial solution, the ‘Greek hydraulic telegraph’ — as explained at, for example, [https://en.wikipedia.org/wiki/Hydraulic\\_telegraph](https://en.wikipedia.org/wiki/Hydraulic_telegraph), where more detail and further

---

<sup>5</sup> The Treaty was a huge diplomatic effort: 176 negotiating teams representing 194 agents; the French brought a 150-strong entourage, including pastry chefs, priest confessors and dancing instructors. The negotiators had constantly to check with their capital: a letter took 2 weeks to reach Stockholm, 4 weeks to reach Madrid. R. Boyes, ‘Treaty that created ‘the soil of despair’’, *The Times*, 24 October 1998.





Figure 1: The Greek hydraulic semaphore  
([https://en.wikipedia.org/wiki/Hydraulic\\_telegraph](https://en.wikipedia.org/wiki/Hydraulic_telegraph))

references can be found.<sup>6</sup> Fire torches were used, by the sender, to initiate and synchronize, with the receiver, a connection between operators at observation points on hills with clear lines of sight between them. Each hill had an identical container, with a valve or spigot at the bottom, filled with water and with a vertical rod floating in the water. The rods were marked with codes at points along their length. The set-up is depicted in Figure 1.

Once the connection was synchronized, each operator would open the valve until the water had emptied to the point marking the required code, at which point the operators would close their valves and simultaneously lower their torches. The length of time the sender's torch remained raised determined a specific, predetermined message. In principle, such a system could be used to send messages in full written language, but, in practice, it would seem likely that the need for efficiency would dictate a small, fixed set of possible messages.

Does the technology of beacons and semaphores, as developed from the Greek world, support something that corresponds the concept of cyberspace? First, there is an underlying physical infrastructure, chains of torch beacons and data-processing water containers, which form a network of communication routes. This network supports logical connections between the individuals wishing to communicate with other individuals at other locations. Second, while the system of beacons and containers transmit data from location to location, it is the humans who send the messages (not the operators of the infrastructure) and

---

<sup>6</sup> 'Telegraph' means 'distance writing' in Greek.

## DRAFT: NOT FOR CIRCULATION

who interpret the data, giving it meaning as information. Last, there is indeed an essential interplay between the physical and virtual. I think I must conclude that 4<sup>th</sup> Century BCE had a form of cyberspace.

A major advance in the development of cyberspace occurred in late 17<sup>th</sup> Century France. 'Le système Chappe', developed by Claude Chappe (see, for example, [https://en.wikipedia.org/wiki/Claude\\_Chappe](https://en.wikipedia.org/wiki/Claude_Chappe) and also Standage's delightful book, 'The Victorian Internet' [35]), was a nationwide semaphore network used for (relatively complex) government and military communications [35]. At its greatest extent, it connected Paris to Amsterdam and Calais to the north, to Mainz, Strasbourg, and Venice to the east, to Marseille, Perpignan, and Bayonne to the south, and to Nantes, Brest, and Cherbourg to the west.



Figure 2: A replica of one of Chappe's semaphore towers in Nalbach, Germany

Chappe's system was a network of towers (see Figure 2) each of which supported two arms that rotated into different positions. The positions included codes for letters and numbers as well as for control signals used to verify the correctness of the reproduction of the messages as they passed from one tower to the next.

# DRAFT: NOT FOR CIRCULATION

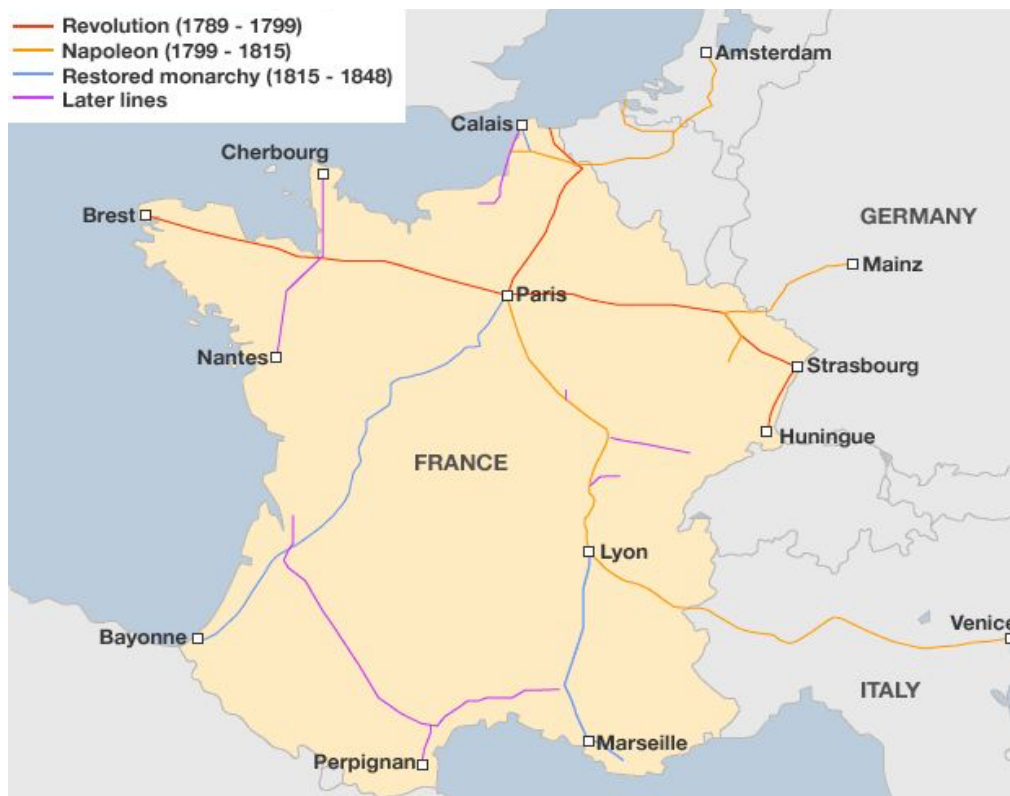


Figure 3: Development of the Chappe Telegraph, 1793-1854 (BBC News, Hugh Schofield, 17 June 2013, <http://www.bbc.co.uk/news/magazine-22909590>)

Somewhat later, in 1838, an English civil engineer called Francis Whishaw also proposed a hydraulic telegraph. Whishaw's hydraulic telegraph [43] was based on the levels of water observed in vessels connected by water-filled pipes: a change in level at one end, representing the sent message, is reflected at the receiving end of the pipe with no perceptible time delay. This proposal represents a possible improvement in both speed and reliability, but essentially the same sense of a cyberspace as the Greek version.

Amusingly, the term semaphore persists in the modern world of information technology in the theory and practice of concurrency, in which two computer programs execute at the same time while attempting to use shared resources. In this context, a semaphore refers to a variable, or other abstract data type, that is used to control access to a resource between concurrently executing processes.

### 3. Some key concepts

Before considering, in the sections that follow, cyberspace in the modern world, it is important to mention three of its other key precursors, both of which have contributed to its technological and social infrastructures.

- The commercially available printing press. Printing presses have existed in China and Korea for around 1800 years, but it is perhaps Gutenberg's

# DRAFT: NOT FOR CIRCULATION

introduction of a commercially available, and well-promoted, service that marks the entry of printing into a 'space' of communication. Copies of single handwritten manuscripts, representing resources of knowledge and typeset using 'standard' characters, could be mass produced and circulated widely. Thus, knowledge could be shared around many distinct locations, commented on, modified, and further shared.

- Postal services. Although commercial printing presses provided a means of mass-producing information resources, so that they might be consumed by many different individuals residing at many different locations, the realization of this sharing requires a process for circulating copies of manuscripts. Postal services provided the first reliable such processes and, in so doing, adumbrate some concepts that are important in modern cyberspace. These include the following (with no particular historical period or timeline implied):
  - *addresses*: the sender of package writes a code on the package that specifies the destination and recipient of the package; the service provider interprets the code in order to execute the process of delivering the package: for this to work, addresses must be written in an agreed, or at least recognizable, format;
  - *routing protocols*: packages might be collected from widely distributed starting points, such as letter boxes, then taken to a local collecting point and combined in to large groups of packages that are moved to a distant collecting point (possibly involving many such steps), from where individual packages are delivered to their final destinations: to make this work, the provider of the postal services must implement processes that collect, sort, and distribute packages; and
  - *supporting infrastructure*: the service provider must provide the equipment (postage stamps, letter boxes, bags, vehicles, buildings, etc.) and personnel to collect, sort, and distribute the packages: the provider may also make use of other services, such as stage coaches, trains, and aeroplanes, and must agree terms of service with them.

A postal service, viewed as service to its users, also has two key features:

- *mass availability*: the postal service is available to all who are able to purchase the tokens, such as stamps, required to access the service;
  - *service guarantees*: state actors, such as monarchs or governments, might provide guarantees, with supporting policies, that packages will be delivered to their intended destinations and recipients and that they will be undamaged in transit.
- Telephone systems. Such systems also provide examples of the importance of the concepts of addresses, routing protocols, and

# DRAFT: NOT FOR CIRCULATION

supporting infrastructure. It also demonstrates mass availability and, at least implicitly, service guarantees.

Agar's *The Government Machine* [1] provides an excellent general contextual discussion for this perspective, exploring the mechanization of government work in the United Kingdom from the nineteenth to the early twenty-first century.

## 4. The beginnings of cyberspace in the modern world: semaphores and telegraphs

A key technology of the Victorian period, the telegraph system developed in the 1830s and 1840s by Samuel Morse, provides another example of the importance of the concepts addresses, routing protocols, and supporting infrastructure. It also demonstrates mass availability and, at least implicitly, service guarantees.

According to the OED [42]:

telegraph, n. — a system of or instrument for sending messages or information to a distant place; v. — to signal (from French *télégraph*)

I have explained that the idea of an optical telegraph dates from the ancient world. Perhaps the most well-known early precursor to the Internet, however, and certainly the one with the strongest resemblance, is the telegraph systems developed in the 19<sup>th</sup> Century. The history of this 'Victorian Internet' has been elegantly and captivatingly described in Standage's book [35], which, as we shall see, helps us to understand the significance of 'The Victorian Internet' for the origins of the concept of cyberspace.

Figure 4 illustrates the major global telegraph connections around the world in 1891. Compare with the modern map of telecommunications cables given in Figure 5.

- First, there is an underlying network infrastructure:
- Second, there are several key resources, placed around the locations of the network infrastructure, upon which the operation of the telegraph depends:
  - the network cables used to connect different points around the world;
  - the electrical devices that generate and receive the electrical signals that are used to encode messages for transmission across the network;
  - the human operators of the devices who translate between natural language and the encoded messages;
  - the paper used to write down messages to be encoded and messages that have been decoded.
- Third, the communication of messages between points on the telegraph network occurs as a collection of concurrent processes that utilize the resources present at locations around the network.

# DRAFT: NOT FOR CIRCULATION



Figure 4: Telegraph Connections (Telegraphen Verbindungen), 1891, [Stielers Hand-Atlas](#), Plate No. 5, Weltkarte in [Mercator projection](#)

Does the telegraph network of the 19<sup>th</sup> Century support something that corresponds to the concept of cyberspace? Let's consider the concepts of cyberspace that I have identified.

- First, there is an underlying network infrastructure:
- Second, there are several key resources, placed around the locations of the network infrastructure, upon which the operation of the telegraph depends:
  - the network of cables used to connect different points around the world;
  - the electrical devices that generate and receive the electrical signals that are used to encode messages for transmission across the network;
  - the human operators of the devices who translate between natural language and the encoded messages;
  - the paper used to write down messages to be encoded and messages that have been decoded.
- Third, the communication of messages between points on the telegraph network, occurs as collection of concurrent processes which utilize the resources present at locations around the network.

The telegraph system thus came very close to delivering a cyberspace. Many of the features that I have identified were present, but one aspect that is missing, at least in a sufficiently explicit form, is that of the shared social space created by the technology and its users. Although messages could be sent and received very efficiently, there is no way to post information that can be read and contributed to by other participants in the space; there is no way to implement something like Facebook using the telegraph system.

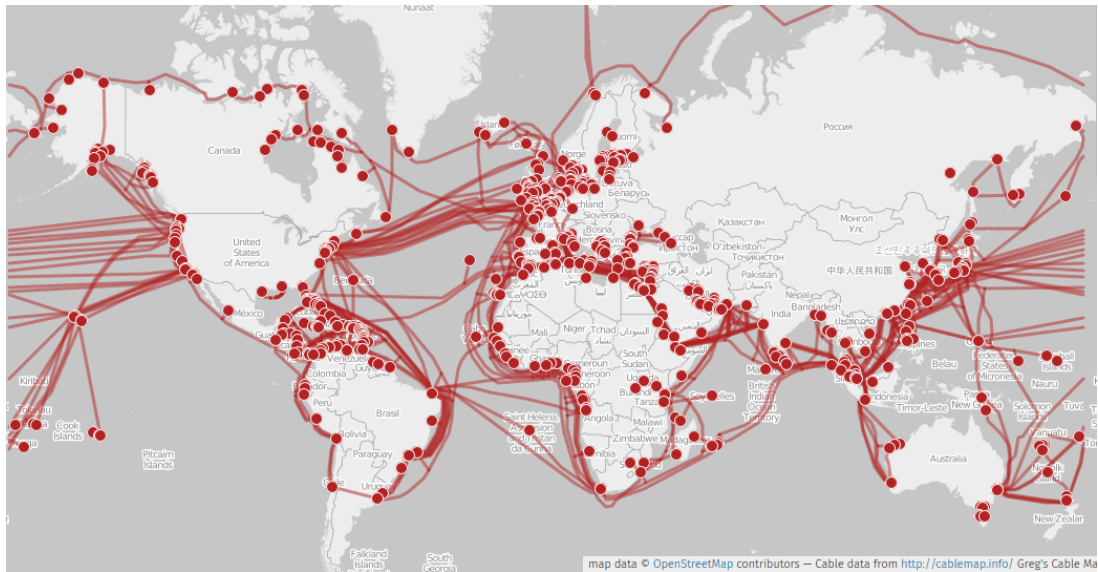


Figure 5: Undersea telecommunications cables, 2015  
(<http://www.submarinecablemap.com>)

## 5. The infrastructure of modern cyberspace

The technology supporting worldwide data communications did not significantly advance from the telegraph (wired and wireless) until the early stages of the development of what would become the Internet. Indeed, as can be seen in Figure 5, the pattern of connectivity even now reflects that of the wired telegraph network (Figure 4).

The ARPANET — which stands for Advanced Research Projects Agency Network (<https://en.wikipedia.org/wiki/ARPANET>), after the United States' Advanced Research Projects Agency (<https://en.wikipedia.org/wiki/DARPA>) that funded its development — was the seed that would eventually grow into the Internet. It was proposed in 1968 and established in 1969, with the first link being between UCLA and Stanford University. Famously, the first message sent between the two sites was 'LO' — the first two characters of 'LOGIN'; the connection failed before the command could be completed.<sup>7</sup>

The ARPANET was an early 'packet-switching' network — in which transmitted data is grouped into blocks, called 'packets', that are of a suitable size (depending on things like the network's 'bandwidth') for transmission across a network — that implemented the TCP/IP protocol, upon which the modern Internet

---

<sup>7</sup> The SAGE (Semi-Automatic Ground Environment) missile defence system, which was developed by MIT's Lincoln Laboratory and which operated in the United States from the late 1950s to the 1980s, is also a seed ([https://en.wikipedia.org/wiki/Semi-Automatic\\_Ground\\_Environment](https://en.wikipedia.org/wiki/Semi-Automatic_Ground_Environment)).

# DRAFT: NOT FOR CIRCULATION

depends. Packet switching ([https://en.wikipedia.org/wiki/Packet\\_switching](https://en.wikipedia.org/wiki/Packet_switching)) stands in contrast to ‘circuit switching’ ([https://en.wikipedia.org/wiki/Circuit\\_switching](https://en.wikipedia.org/wiki/Circuit_switching)), as used in early telephone networks, in which dedicated circuits are established between two points (e.g., two telephone service subscribers) that wish to communicate. Circuit switching, which does not require the overhead of decomposing messages into packets and recomposing after transmission, could be used in the Internet. However, it makes much less efficient use of the available network capacity (or ‘bandwidth’).

Recalling our discussion of semaphore and telegraph systems, it can be seen that although they also required a notion of packet in order to send and receive messages — words are coded as delineated sequences of coded letters, Morse code ([https://en.wikipedia.org/wiki/Morse\\_code](https://en.wikipedia.org/wiki/Morse_code)), they all really worked by establishing circuits between the communicating locations, as with early telephone networks. Packet switching is perhaps the key conceptual advance of the Internet over the telegraph networks.

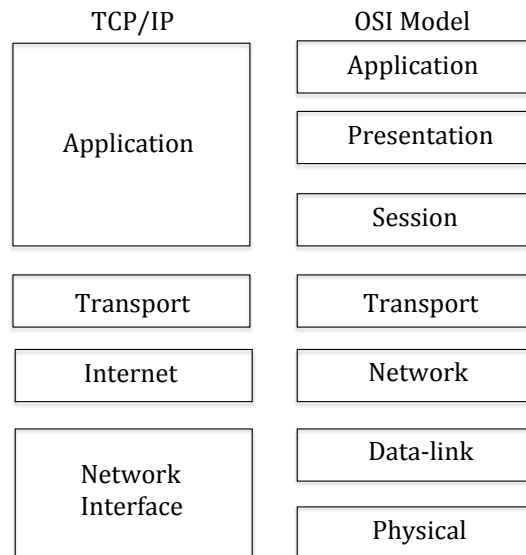


Figure 6: The TCP/IP model and the OSI reference model (many similar diagrams may be found in the literature)

The TCP/IP protocol is one example, a very important example, of a specification of a network communications protocol — recall from our discussion of distributed systems the essential need for ‘common means of communication’ — that is tailored to the underlying physical technology that supports its operation. Such technology is not unique, however, and the Open Systems Interconnection (OSI) model provides a standardized reference — it describes the essential features of the infrastructure of the Internet. Figure 6 illustrates how the components of these protocols are built up, from the physical layer, providing underpinning infrastructure, through logical organizational layers, to the application layer, providing services to users. The TCP/IP model can be seen as



# DRAFT: NOT FOR CIRCULATION

an implementation of the OSI model; for example, the general 'Network' layer described in the OSI model is implemented by the Internet in the TCP/IP model.

Figure 6 depicts how, in a somewhat simplified and quite widely described form, the Internet can be seen as implementing cyberspace. To see this, first recall our summary of cyberspace from the end of Section 1. Note that, for the purposes of this discussion, I am taking cyberspace to be represented by the medium of the WWW and its use by humans.<sup>8</sup>

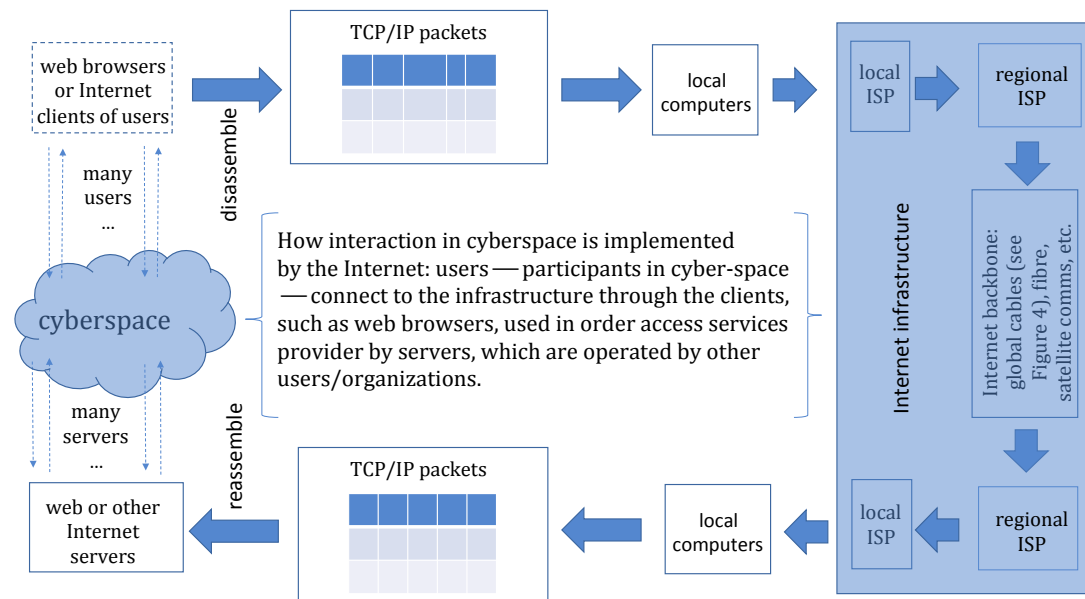


Figure 7: How the Internet implements cyberspace (many similar diagrams may be found in the literature)

Does this implementation of cyberspace deliver what we expect? Recall what we learned about cyberspace in Section 1:

- First, that it is a concept that builds on the physical and logical infrastructure provided by the Internet;
- Second, that while the Internet processes data, it is information — that is, interpreted data — that is the medium of cyberspace;
- Third, that the interpretation of data, and the processing of information, is performed by humans, who are themselves essential components of cyberspace;
- Fourth, that there is an essential interplay between — a merging of — the physical and the virtual.

The first point is clearly supported by this picture: the communication between a web browser and web server, both of which have both physical location and logical location, is implemented by a sequence of flows of data over a physical

<sup>8</sup> For our present purposes, I refrain from considering AI alternatives to humans as interpreters of data to yield information.

## DRAFT: NOT FOR CIRCULATION

network, but that physical network supports interpretations of that data relative to a logical architecture, which organizes the information into useful systems (of knowledge, understanding, and so on). The second and third points reside in the users' interpretations as information of the data that flows between the web browser and the web server. Finally, the fourth point summarizes the overall relationships between the components of the diagram in Figure 7: data and information have both physical and logical location; data is processed at physical locations, around the loop between browser and server, but interpreted at logical locations by the human users — in, we might say, *Bewusstseinslagen*.<sup>9</sup>

Within cyberspace itself, as implemented by the Internet, the structural organization provided by the distributed systems model is not always the most helpful. Rather, it is sometimes more useful to infer information about cyberspace in terms of what the statistical structure of the data and its flows tells us about patterns of use. In this context, topological modelling approaches such as that suggested in [25] may also be helpful.

Examples of this kind of analysis include answering questions about the density of Internet use in different countries around the world (see, for example, Figure 8), which social networking sites are more popular in which countries, from where most phishing attacks originate, and with what levels of intensity, and so on. Figure 8 illustrates Internet users in 2015 as a percentage of a country's population.

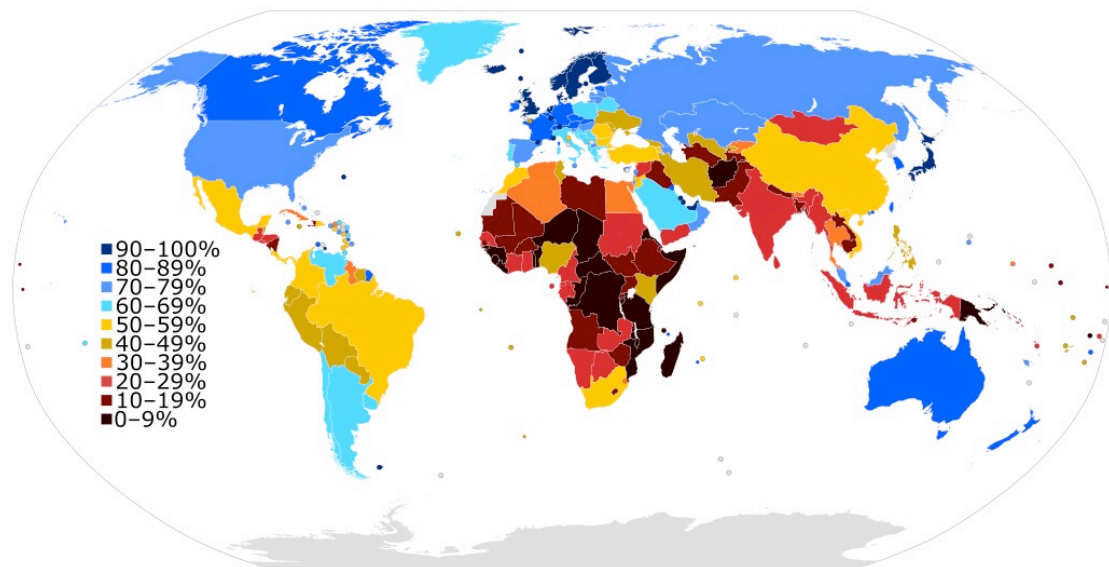


Figure 8: Internet users in 2015 as a percentage of a country's population (source: International Telecommunication Union).

---

<sup>9</sup> *German*: a state of consciousness or a feeling devoid of sensory components (Merriam Webster).

# DRAFT: NOT FOR CIRCULATION

How did a world of communication based on semaphores and Morse code sent over telegraph cables become a world dependent on Internet-supported cyberspace? The answer really is the story of the development of modern computer science (though see [19] for a useful perspective): I cannot hope to do justice to that in this chapter. Rather, I hope to provide a conceptual framework for understanding and reasoning about cyberspace that is applicable to all of these stages in the history of cyberspace.

## 6. Modelling and reasoning about cyberspace

*This section is primarily intended for those readers with a more mathematical, or at least philosophical, background and, in particular, for those with an interest in logic. Nevertheless, I hope that all readers who are willing to encounter a little formalism will be able to appreciate the value of the perspective I describe.*

As I have described, *distributed systems* provide a model of computation in which information-processing devices are located on networks and communicate with one another, and with their environments, and co-ordinate their actions by passing messages between one another and between themselves and their environments. The resulting interaction of these components of systems and their environments delivers the systems' services to their clients.

Mathematically, distributed systems can be described using the following concepts, as described in Section 1:

- *Locations.* Mathematically, locations are described using topological structures that give a useful account of the (physical or virtual/logical) notion of 'place' and 'connections between places'. The leading example is perhaps directed graphs, but mathematically other structures can also be used [11]. The concept of location, and its intended mathematical characterization in this context, provides the topological component that is a core part of the concept of space discussed above.
- *Resources.* Mathematically, resources are modelled by abstract algebraic structures called 'partial monoids'. These gadgets are sets that come with an operation, which has a unit or neutral element, for combining some, but not all, of their elements (used in [11,20]). Perhaps the most important example of such a monoid is given by the set of *natural numbers* (with 0) less than or equal to a specified maximum, *max*. Combination is addition, with unit 0. The combination of two numbers *m* and *n* is defined just in the case that  $m + n$  is less than *max*. Another important example is given by the 'stack' and 'heap' in computer memory (RAM) [20].
- *Processes.* Mathematically, processes are described using structures called 'transition systems' and an important class of examples of transition systems are described by 'process algebras' [2,11,23].

# DRAFT: NOT FOR CIRCULATION

The key idea is that the state of system is described by a triple  $L, R, E$  consisting in the configuration of the system's locations,  $L$ , the distribution of its resources,  $R$ , around its locations, and a description,  $E$ , of the processes that are currently executing. When an action occurs during that execution the resources are manipulated, perhaps being consumed, created, or moved to new locations.

Again, as described in Section 1, systems exist within environments with which they interact (i.e., they are part of an ecosystem). This interaction is typically described in terms of the incidence of events in and out of the model.

- *Environment*. Mathematically, the incidence of actions from the environment upon a model and, conversely, the incidence of actions from a model on the environment can be represented simply using probability distributions. Perhaps the paradigmatic example of this is the arrival of entities, be they people or packets of data, in a queue, where the arrivals at the queue are described using the negative exponential distribution [29], which has just one parameter, the 'arrival rate'.

With this machinery in place, I have, mathematically speaking, all I need to describe the logical and physical architecture of the Internet; that is, the infrastructure of cyberspace:

- *Locations* in the Internet are given by a range of examples that are relevant for my discussion.
  - The physical network graph: Figures 3, 4, and 5 give examples of network graphs; the global scale network connects regional and national networks, which, in turn, connect organizational and domestic networks.
  - The virtual network graph: Organizations may be distributed across the world and yet appear to be a single network location; for example, a multinational corporation may have physical presence in many countries, but its networks may all be part of the same family of IP addresses, so that they appear as part of the same network location even though they are in many different physical locations.<sup>10</sup>
  - The locations of the human users of the Internet, participants in cyberspace, and the devices with which they interact and upon which the services they use depend.
- *Resources* in the Internet are things like computers, some providing computation, some providing network management and some providing storage, peripheral devices such as printers and scanners, security devices

---

<sup>10</sup> An IP address is a numerical label (a sequence of numbers representing a 32-bit or 128-bit number) assigned to each device connected to a network that uses the Internet protocol (IP) for its communications. An organization may, for example, own all of the IP addresses that begin with a given sequence of numbers.

## DRAFT: NOT FOR CIRCULATION

such as IDSs and IPSs,<sup>11</sup> and people, such as programmers, system administrators and end-users (the participants in cyberspace).

- *Processes* in the Internet are the things that happen. For one example, an individual computer's operating system is a program that executes continuously in order to provide all the computer's services to its users; screen, keyboard, network connection, application execution, and so on. For another, the services provided by the network of servers and routers that support the operations of an Internet Service Provider; and, for another, the HR, financial, and other business processes followed by the users of information and management systems.

And, finally:

- *Environment* in this context provides a way for the modeller to focus on a particular part of the Internet, or indeed of cyberspace, while retaining an appropriate representation of the rest of the network on that specific part. While the part of interest is modelled in detail, using the concepts of location, resource, and process as described above, the interaction of that part with the rest of the network is modelled simply in terms of the incidence of events across the boundary of the part modelled in detail.

Although I have described a framework for describing the underlying conceptual and technical infrastructure of cyberspace, I have not yet provided a way to describe cyberspace itself: even if we consider that humans and their interaction with the architecture, and, indeed, other humans, can be described using located resources and processes, we still lack a natural way to talk about the interpretation of data and how humans reason about it. For that, we are going to take our final step, to *logic*.

Logic is the science of reasoning. It is studied within computer science, mathematics, and philosophy. It means the same thing in all of these areas, although they each tend to emphasize different aspects of its study; they interact with one another very fruitfully. In computer science, in particular, it is important to become accustomed to the idea that there is no single, all-encompassing system of logic that is well adapted to all of the different kinds of reasoning that are needed.<sup>12 13</sup> The discussion in [47] of the use of logic as a modelling technology may be useful for some readers.

---

<sup>11</sup> Intrusion Detection Systems and Intrusion Prevention Systems ([https://en.wikipedia.org/wiki/Intrusion\\_detection\\_system](https://en.wikipedia.org/wiki/Intrusion_detection_system)).

<sup>12</sup> That is not to say that many things expressible in one kind of logic cannot be expressed in another; rather, that it may not be convenient to do so.

<sup>13</sup> In addition to the classical propositional and predicate logic that is routinely taught to undergraduates in computer science, mathematics, and philosophy, we can add, among other things, modal, temporal, and epistemic systems, and their higher-order, intuitionistic, and substructural variants. The many volumes of the *Handbook of Philosophical Logic* and the Stanford

# DRAFT: NOT FOR CIRCULATION

Mathematical models of distributed systems of this kind are closely associated with ideas from logic. Logic is the science of reasoning and one of its key ideas is that of *truth*.

In logic, truth is a very precisely defined concept. It relies on a few key ideas: syntax, semantics, and interpretation. Syntactic entities are interpreted as semantic entities, just as data is interpreted as information. Truth is a property of a logical formula (which is a syntactic entity) relative to a *model*. A model is a mathematical structure that describes relationships between semantic entities.

In modal logic [9], which is perhaps the key tool in the logician's kit for reasoning about action, the ideas of *necessity* and *possibility*, so-called 'modalities', can be expressed. The key to understanding these ideas comes from some beautiful work initiated by Saul Kripke [22], the application of which in systems modelling is discussed in, for example, [2,10,11,33].

The main idea is that truth is defined relative to a *world*. The concept of a world is philosophically quite delicate, but for our present purpose we can think of it as a state of knowledge or the state of a system. A collection of such 'possible worlds' that might be taken as a place to give meaning to formal logical expressions can be seen as a space in the sense that we have already discussed. In fact, this kind of semantics can be formulated explicitly in terms of topological spaces, which are perhaps the prototypical mathematical example of the concept of a space [45,46].

What is most important for my story here is that the set of all worlds,  $W$ , must come with a partial ordering<sup>14</sup> on its set of elements, so that we define truth relative to models  $\mathcal{M}$  of the form  $(W, \leq)$ .

Given a world  $w$  in the set  $W$ ,

$$w \vDash_{\mathcal{M}} \phi$$

denotes that the logical formula  $\phi$  is 'true in the state  $w$ '. For example, if  $w$  is a state in which there are precisely three apples and two oranges, then the formula *More (Apples, Oranges)*, which is intended to mean that there are more apples than there are oranges, is true at  $w$ . But, if  $w$  is any state in which has at least as many oranges as apples, then the formula is not true there.

---

Encyclopaedia of Philosophy [36] provide starting points for exploring these topics.

<sup>14</sup> A partially ordered set formalizes the intuitive concept of an ordering of the elements of a set. A 'partial order' on the set relates pairs of elements of the set in such a way that the relationship between the elements of the pair is reflexive (every element is related to itself), anti-symmetric (no two elements precede each other in the ordering), and transitive (if the higher of one pair is below the lower of another pair, then the lower of the former pair is below than the higher of the latter pair). Not all pairs of elements of the set need be related by the order; such pairs are 'incomparable' in such an order.

# DRAFT: NOT FOR CIRCULATION

Suppose now that our states are ordered as follows:  $w \leq v$  just in the case that  $v$  has more apples than  $w$ .

In general,  $w \models_{\mathcal{M}} \Box \phi$  denotes that the formula  $\phi$  is *necessarily* true at the state  $w$  in the model  $\mathcal{M}$ . This is defined as follows:  $w \models_{\mathcal{M}} \Box \phi$  just in the case that, for *every* state  $v$  such that  $w \leq v$ , it is the case that  $w \models_{\mathcal{M}} \phi$ . So, in our little example of apples and oranges, the formula  $\Box \text{More (Apples, Oranges)}$  is true if  $w$  is a state in which there are precisely three apples and two oranges because any state that is beyond  $w$  must have more apples than oranges. Note that we may choose to consider many possible models. Different models will, in general, have different sets of possible worlds.

Similarly,

$$w \models_{\mathcal{M}} \Diamond \phi$$

denotes that the formula  $\phi$  is *possibly* true at the state  $w$  in the model  $\mathcal{M}$ . This is defined as follows:

$$w \models_{\mathcal{M}} \Diamond \phi \text{ just in the case that, for } \textit{some} \text{ state } v \text{ such that } w \leq v, \text{ it is the case that } \\ w \models_{\mathcal{M}} \phi.$$

So, now supposing that states are ordered so that  $w \leq v$  just in the case that either  $v$  has more apples than  $w$  or  $v$  has more oranges than  $w$ , then there is state beyond  $w$  at which

$$\text{More (Oranges, Apples)}$$

is true, so that

$$\Diamond \text{More (Oranges, Apples)}$$

is true at  $w$ .

What has logic got to do with cyberspace? Just about everything, actually, if one believes that logic provides a good, or at least useful, account of human reasoning. Moreover, computers are inherently machines that implement logic. Assuming that at least, then logic provides the essential link between people and the systems that support cyberspace, as described in Figure 9.

# DRAFT: NOT FOR CIRCULATION

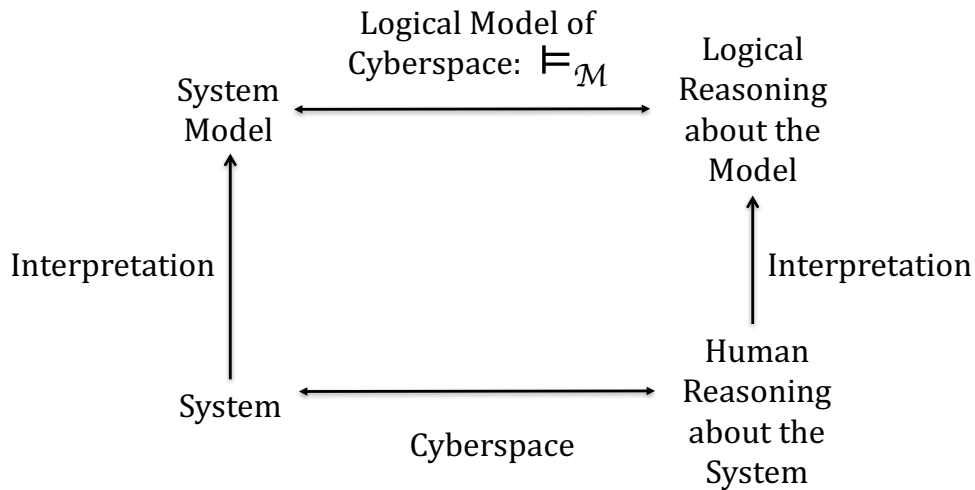


Figure 9: Reasoning about systems and reasoning about system models

The diagram indicates the relationship between reasoning about systems (including about other humans within the system) and logical reasoning about mathematical models of the system. The ideal situation is when this diagram ‘commutes’; that is, for a given system, the logical formalization of human reasoning about its properties corresponds exactly to logical reasoning about a formal model of the system. Such a situation is very rare indeed, and really only works out in the context of very specific reasoning tools (see, for example, discussions in [3,20,10,27,33]).

I have already given an example of a model in the context of distributed systems. It is built out of descriptions of the system’s locations, resources, and processes. Triples of locations, resources, and processes are the states of the model of a system. The ordering of the states is then given by the evolution of the model as *actions* occur [2,11,23].

Actions are the basic building blocks of processes<sup>15</sup>, one of the core components in our model of the infrastructure of cyberspace. When actions occur, the state of a system changes. For one example, a computer program may perform a ‘write’ action to put value, a number representing a *resource*, to a memory, *location*. For another example, a human user of a computer may give the ‘print’ command, so causing data to be copied from the computer’s memory to the printer, followed by the consumption on ink and paper and resources, and the creation of a document.

Thus we have  $L, R, E \leq M, S, F$  just in the case that  $L, R, E$  can evolve to become  $M, S, F$  by some action. The notion of logical truth supported by such a mode, written as  $L, R, E \models_{\mathcal{M}} \phi$ , is read as ‘the process  $E$ , executing with respect to resources  $R$ , at location  $L$ , has property  $\phi$ .’

---

<sup>15</sup> Technical note: for an elegant explanation of the structure of processes, including concurrency, nondeterminism, and recursion, see Robin Milner’s ‘Communication and Concurrency’ [23].



# DRAFT: NOT FOR CIRCULATION

Then we can define versions of the necessity and possibility modalities that are parametrized by actions. The counterpart to the necessity modality is  $[a]\phi$ , which is read as ‘the formula  $\phi$  is necessarily true after the action  $a$ ’. Similarly, the counterpart to the possibility modality is  $\langle a \rangle \phi$ , which is read as ‘the formula  $\phi$  is possibly true after the action  $a$ ’. More formally, we define them as follows:

- $L, R, E \models_{\mathcal{M}} [a]\phi$  holds just in the case that, for *every* evolution of  $L, R, E$  to  $M, S, F$  by the action  $a$ , we have that  $M, S, F \models_{\mathcal{M}} \phi$  holds;
- $L, R, E \models_{\mathcal{M}} \langle a \rangle \phi$  holds just in the case that, for *some* evolution of  $L, R, E$  to  $M, S, F$  by the action  $a$ , we have that  $M, S, F \models_{\mathcal{M}} \phi$  holds.

These definitions explain how logical reasoning about (the data held by) systems interacts with the actions performed by the system; in particular, they begin to explain, in terms of *information* processing, how humans interact with the system and other humans.

In order to see an example of all of this, think about the picture of the implementation of cyberspace given in Figure 6. Suppose a user, sitting at a computer in their home, is using a web browser to access an online store (let’s call it ‘BigRiver’, say). The user is looking at the webpage for OUP’s ‘Handbook of Cybersecurity’ and clicks on the ‘Buy now’ button. If the user’s bank account has sufficient funds and if the BigRiver website has given the correct information about the availability of the book, then, provided everything works as it should, the book will be sent to the user. We can describe this situation logically as follows:

$home\_computer, bank\_account, BigRiver\_website \models_{Cyberspace} \langle buy\_now \rangle book\_sent$

That is, located at their home computer, with the resources available in their bank account, while running the BigRiver website process, the user may click on the ‘Buy now’ link and it is possible, if all goes well, that the book will be sent to their specified address.

Note that this logical assertion describes a state of affairs in the part of the diagram described as ‘Cyberspace’. We could use similar logical assertions to describe the (many) states of affairs that must obtain in other parts of the diagram in order for the assertion about Cyberspace to be realized.

At this point, the reader might be forgiven for thinking that the logical language that explains all this is rather impoverished. I would agree, although I would note that, in fact, the framework I have suggested, as developed, for example, in [2,3,10,11,23], can express a great deal, albeit somewhat tediously. A challenge for the community of logicians, if logic is to demonstrate what I believe is its full potential as a modelling technology, is to develop concise and powerful representations of logical reasoning.

# DRAFT: NOT FOR CIRCULATION

In fact, the world of (modal) logics for reasoning about actions (performed by agents) is much richer than I have so far suggested and includes *epistemic logics*, for reasoning about agents' knowledge of systems, *doxastic logics*, for reasoning about agents' beliefs, and *temporal logics*, which incorporate a representation of time (system events, such as the sending and receiving of messages, occur at relative points in time). All of these systems of logic build on the basic ideas sketched above. *The Stanford Encyclopaedia of Philosophy* (<https://plato.stanford.edu>) provides a great deal of information about these systems of logic, and there is a vast literature in computer science that is concerned with their use in reasoning about systems.

Perhaps the most important of these many possibilities for our story are the epistemic and doxastic logics. Roughly speaking, in these logics modalities are parametrized not by actions but rather by agents, who of course can perform actions. Agents may be humans or system processes, and epistemic and doxastic logics (again, see [36]) allow us to reason about their knowledge and beliefs. Exploring the use of these logics, and their relationships with tools from behavioural economics (e.g., [4]), game theory (e.g., [8]), and psychology (e.g., [21]), to reason about the behaviour of agents in cyberspace would be another chapter in exploring the origins of cyberspace.

Of course, individuals and organizations do not always behave 'logically' when they interact with one another and with the systems that support cyberspace. They behave in ways that others may consider to be irrational. This observation does not really undermine the perspective presented here. It has very little to do with logic in the sense that I have described, which is about the mechanism by which conclusions are drawn from chosen assumptions. Some assumptions may lead to what may be described as 'irrational' behaviour, even though the logical mechanisms may be perfectly sound.

## 7. Summary

I have sought to explain the origins of 'cyberspace' historically, linguistically, and conceptually. I have explained how the idea of cyberspace derives from a complex combination of physical and logical structure, which supports complex interactions between and amongst humans and information-processing machines and I have given a conceptual and mathematical framework for modelling the conceptual and technical infrastructure of cyberspace.

I have also explained how logic can provide tools, based on our approach to modelling the infrastructure of cyberspace, for capturing how humans and other agents reason about cyberspace, and so, to some extent at least, how they experience cyberspace. Experience, of course, involves more than logical reasoning alone. Exploring that dimension would be yet another chapter in exploring the origins of cyberspace.

I have explained how the essential features of cyberspace have been part of the human experience, 'a consensual hallucination experienced daily by billions of

# DRAFT: NOT FOR CIRCULATION

legitimate operators', for a very long time and, as the science fiction writers continue to predict, we can expect that there is much more to come. I have not discussed questions of security in cyberspace; that is the topic of the rest of this handbook. Again, Standage's book [35, Chapter 7, 'Codes, Hackers and Cheats'] provides a delightful starting point.

## Acknowledgements

I am warmly grateful to Tristan Caulfield and Jonathan Spring for their thorough and thoughtful advice on drafts of this article.

## Bibliography

1. J. Agar (2016). *The Government Machine*. MIT Press.
2. G. Anderson and D. Pym (2016). A Calculus and Logic of Bunched Resources and Processes. *Theoretical Computer Science*, 614:63-96.
3. K. Apt, F. de Boer, and E.-R. Olderog (2010). *Verification of Sequential and Concurrent Programs*. Springer.
4. A. Baltag and B. Renne. Dynamic Epistemic Logic. In *Stanford Encyclopedia of Philosophy* (E. Zalta, Principal Editor). <https://plato.stanford.edu/entries/dynamic-epistemic/>. Accessed 3 July 2018.
5. J. Barwise and J. Seligman (1997). *Information Flow: The Logic of Distributed Systems*. Cambridge University Press.
6. T. Berners-Lee (1989, 1990). *Information Management: A Proposal*. CERN, March 1989, May 1990, W3C Archive, <https://www.w3.org/History/1989/proposal.html>.
7. T. Berners-Lee (2000). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Harper Business.
8. K. Binmore (2007). *Playing for Real: A Text on Game Theory*. OUP USA.
9. P. Blackburn, M. de Rijke, and Y. Venema (2001). *Modal Logic*. Cambridge University Press.
10. T. Caulfield and D. Pym (2015). Modelling and Simulating Systems Security Policy. In *Proc. SIMUTools 2015*, ACM Digital Library, doi: [10.4108/eai.24-8-2015.2260765](https://doi.org/10.4108/eai.24-8-2015.2260765).
11. M. Collinson, B. Monahan, and D. Pym (2012). *A Discipline of Mathematical Systems Modelling*. London: College Publications.
12. G. Coulouris, J. Dollimore, T. Kindberg, and G. Blair (2011). *Distributed Systems: Concepts and Design*. Pearson.
13. H. van Ditmarsch, J. Y. Halpern, Wiebe van der Hoek, and Barteld Kooi (Editors) (2015). *Handbook of Epistemic Logic*. London: College Publications.
14. L. Floridi (2011). *The Philosophy of Information*. Oxford University Press.
15. W. Gibson (1982). *Burning Chrome*. New York City: Omni, 4(10), July 1982.
16. W. Gibson (1984). *Neuromancer*. New York City: Ace.
17. K. Hafner and M. Lyon (1998). *Where Wizards Stay Up Late: The Origins of the Internet*. Simon & Schuster.

## DRAFT: NOT FOR CIRCULATION

18. D. Hook and J. Norman (2002). *Origins of Cyberspace*. Novato, California: History of Science.com.
19. W. Isaacson (2014). *The Innovators: How a Group of Inventors, Geniuses, and Geeks Created the Digital Revolution*. London: Simon & Schuster UK.
20. S. Ishtiaq and P. O'Hearn (2001). BI as an assertion language for mutable data structures. *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 14-26. ACM.
21. E. Kirchler and E. Hoelzel (2017). *Economic Psychology: An Introduction*. Cambridge University Press.
22. S. Kripke (1963). Semantical Considerations on Modal Logic. *Acta Philosophica Fennica* 16:83-94.
23. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
24. MIT Artificial Intelligence Laboratory, *The JAIR Information Space*, [MIT Artificial Intelligence Laboratory](http://www.ai.mit.edu/projects/infoarch/jair/jair-space.html), 10 June 1998, <http://www.ai.mit.edu/projects/infoarch/jair/jair-space.html>.
25. K. Ohmori and T.L. Kunii (2007). The Mathematical Structure of Cyberworlds. *Proc. 2007 International Conference on Cyberworlds*. IEEE Computer Society. doi: 10.1109/CW.2007.19
26. J. Prucher (2007). *Brave New Worlds: The Oxford Dictionary of Science Fiction*. Oxford University Press.
27. D. Pym, D., J. Spring, and P. O'Hearn (2018). Why Separation Logic Works. *Philos. Technol.* <https://doi.org/10.1007/s13347-018-0312-8>
28. T. Rid (2013). *Cyber War Will Not Take Place*. London: Hurst and Company.
29. S. Ross (2014). *Introduction to Probability Models*. 11<sup>th</sup> Edition. Academic Press.
30. N. Schachtman. 26 years after Gibson, Pentagon defines 'cyberspace'. *Wired* 5 May 2008, <https://www.wired.com/2008/05/pentagon-define/>. Accessed 29 May 2018.
31. H. Schofield. How Napoleon's semaphore telegraph changed the world. BBC News, 17 June 2013, <http://www.bbc.co.uk/news/magazine-22909590>. Accessed 29 May 2018.
32. <http://www.theshulers.com/whitepapers/internet-whitepaper/>. Accessed 29 May 2018
33. H. Simon. *The Sciences of the Artificial*. MIT Press, 1996.
34. P. Singer and A. Friedman (2014). *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford University Press.
35. T. Standage (1998). *The Victorian Internet*. London: Weidenfeld & Nicolson, 1998.
36. Stanford Encyclopaedia of Philosophy. <https://plato.stanford.edu>. Accessed 29 May 2018.
37. J. Garson. Modal Logic. In *Stanford Encyclopedia of Philosophy* (E. Zalta, Principal Editor). <https://plato.stanford.edu/entries/logic-modal/>. Accessed 3 July 2018.
38. N. Wiener (1948). Hermann & Cie, editors, *Cybernetics; or, Control and communication in the animal and the machine*. Paris: Technology Press.
39. N. Wiener (1950). *Cybernetics and Society: The Human Use of Human Beings*. Houghton Mifflin.
40. N. Wilkinson and M. Klaes (2012). *An Introduction to Behavioural Economics*. Palgrave.

## DRAFT: NOT FOR CIRCULATION

41. C. Zins. Conceptual approaches for defining data, information, and knowledge. *Journal of the American Society for Information Society and Technology* 58(4). <https://doi.org/10.asi.20508>.
42. The Oxford English Dictionary. <https://en.oxforddictionaries.com/definition/cyberspace>. Accessed 29 May 2018.
43. [https://en.wikipedia.org/wiki/Hydraulic\\_telegraph](https://en.wikipedia.org/wiki/Hydraulic_telegraph).
44. *The Mechanics Magazine, Museum, Register, Journal, and Gazette*. October 7<sup>th</sup>, 1837—March 31<sup>st</sup>, 1838.
45. A. Tarski. *Logic, Semantics, Metamathematics: Papers from 1923 to 1938*. Oxford, Clarendon Press (1969). Oxford, Clarendon Press (1969).
46. J. C. C. McKinsey and A. Tarski. The Algebra of Topology, *Annals of Mathematics* 45 (1944), 141–91.
47. D. Pym. Resource semantics: logic as a modelling technology. *ACM SIGLOG News*, April 2019, Vol. 6, No. 2, 5-41.