



## Mapping technological trajectories as patent citation networks. An application to data communication standards

Roberto Fontana , Alessandro Nuvolari & Bart Verspagen

To cite this article: Roberto Fontana , Alessandro Nuvolari & Bart Verspagen (2009) Mapping technological trajectories as patent citation networks. An application to data communication standards, *Economics of Innovation and New Technology*, 18:4, 311-336, DOI: [10.1080/10438590801969073](https://doi.org/10.1080/10438590801969073)

To link to this article: <https://doi.org/10.1080/10438590801969073>



Published online: 06 May 2009.



Submit your article to this journal [↗](#)



Article views: 556



Citing articles: 61 [View citing articles ↗](#)

## Mapping technological trajectories as patent citation networks. An application to data communication standards

Roberto Fontana<sup>a,b\*</sup>, Alessandro Nuvolari<sup>c</sup> and Bart Verspagen<sup>d</sup>

<sup>a</sup>Department of Economics, University of Pavia, Via San Felice 5, 27100, Pavia, Italy; <sup>b</sup>KITeS, Bocconi University, Via Sarfatti 25, 20139, Milan, Italy; <sup>c</sup>Sant' Anna School of Advanced Studies, Piazza Martiri della Libertà 33, 56127 Pisa, Italy; <sup>d</sup>Department of Economics & UNU-MERIT, Keizer Karelplein 19, Maastricht University, 6211 TC Maastricht, The Netherlands

(Received 22 January 2007; final version received 31 January 2008)

We use patent citation networks to study the dynamics of technical change in Ethernet, a standard for data communication. We first consider the evolution of the broad technological system of local area networks (LANs) and we identify the most significant inventions within the specific subfield of Ethernet. Then we analyse the structure of connectivity of patent citation networks to reconstruct the main technological trajectories in this subfield. Our results suggest that these technological trajectories are characterized by the presence of a number of interconnected technological environments which cluster following an engineering logic. These clusters include the most significant patents related to 'milestones inventions' in the evolution of the LAN technology. Thus, our analysis of the structure of connectivity in patent citation networks seems appropriate to capture both the cumulateness associated to the development of a specific technology and the significant discontinuities that punctuate the trajectory, as well as the technological interrelatedness characterising large technical systems.

**Keywords:** technical systems; technological trajectories; patents; network analysis; data communications

*JEL classification codes:* O30; O33; L96

### 1. Introduction

Extensive literature exists on innovation in 'large technological systems' (Bunn and David 1988). This literature outlines how the dynamics of technical change is affected by the role played by network effects (Katz and Shapiro 1994), technological interrelatedness (Antonelli 1993), and the need for compatibility standards (David and Greenstein 1990). In particular, it has been pointed out that a proper understanding of the evolution of these systems requires an appreciation of the technical bottlenecks arising from the need to integrate different interacting components. In large technical systems, innovation and change at the component level create imbalances between the different parts of the system. These

---

\*Corresponding author. Email: roberto.fontana@unibocconi.it

imbalances act as ‘focusing devices’ (Rosenberg 1976) or ‘heuristics’ (Dosi 1982) for inventing activities.

This paper examines the patterns of technical change in large technical systems as captured by patent citation networks. One of the most common approaches for evaluating the importance of (patented) innovations is to ‘weigh’ them by using patent citations. Indeed, several studies have generally found that patent citations provide a reasonable ‘proxy’ for their technological significance, as they generally appear to be highly correlated with other measures of the value of innovations such as the assessments by technology experts (Jaffe and Trajtenberg 2002). In this paper we argue that, in the specific case of technical systems, this approach may have drawbacks. First, innovation in technical systems is ‘*distributed*’ in the sense that it takes place at the level of individual components across the system. However, single components need to ‘work together’. Thus changes in one specific component trigger changes in other parts of the system. Second, due to these systemic features, innovations tend to be *incremental* in nature and generally occur around well-established technical designs in order to ensure that adequate levels of compatibility are maintained (Antonelli 1992, 12). In this context, an analysis based simply on citation counts may underestimate the significance of specific innovations. Third, innovations tend to occur *continuously*. Both the incremental nature of innovations and their systemic nature contribute to generate a rapid pace of technical change (Antonelli 1993). Again, an analysis based on citation counts may fail to identify those concepts and principles that could act as ‘focusing devices’ for a sequence of inventive activities.

In this paper we consider the case of Ethernet, a standard for data communication. A Data communication equipment comprises all the capital goods enabling the transmission of data among a defined set of computers (such as for example the local area network of a university campus). From a technological point of view, local area networks (LANs) present several points of interest. First, LANs represent the archetypical case of a large technical system in which components must conform to specific communication standards. Accordingly, the history of this technology, whose beginnings can be traced back to the early 1970s, has been punctuated by the emergence and demise of successive communication standards (Ethernet, Token Ring, Fast Ethernet). Second, the technology is characterized by a relatively fast rate of innovation and by strong uncertainty related to the availability of several technological options.

Our analysis is based on the study of citation links among patents in a data communication equipment granted by the US patent office. Our basic assumption is that patent citations represent an indicator of the prior knowledge underlying a specific inventive step. Thus, by examining the structure of the patent citations networks in a technological field, one should be able to identify the main trajectories that have characterized its evolution. Our methodology for the analysis of the patent citation networks follows Hummon and Doreian (1989), who have suggested a number of criteria for the identification of the main paths of ‘connectivity’ in a network which can be understood as the fundamental technological trajectories (Dosi 1982) in a given field (see Verspagen 2007 and Mina et al. 2007 for two recent contributions in this vein).<sup>1</sup>

The analysis of the connectivity of patent citation networks has important implications for assessing the ‘relevance’ of innovations in large technical systems. By extracting the most significant sequences of patents, the structure of connectivity approach allows the identification of the key technological bottlenecks and focusing devices that have played a major role in the evolution of the system. Thus, technologically significant patents should belong to the ‘main paths’ of the citation networks and/or locate at particularly critical ‘junctions’ within those paths. A second merit of this approach is that, by reconstructing

technological trajectories as sequences of patents, it opens the opportunity for a fruitful reconciliation between quantitative and qualitative insights in the study of large technical systems. Indeed, the approach allows to pin down a restricted number of patents, whose content can then be examined in detail. By reading the content of these patents, it is possible to reconstruct the heuristics governing inventive activities and go beyond the assessment of innovation patterns based only on patent counts.

The structure of the paper is as follows. Section 2 illustrates the use of indicators of connectivity structure for the analysis of patent citation networks. In Section 3 we provide a short history and an overview of the evolution of the Ethernet standard based on a detailed inquiry of the engineering and industry literature. Section 4 describes the construction of our patent data-set and gives a preliminary descriptive analysis of its main properties. Section 5 presents the analysis of the patent citation networks using connectivity structure indicators. Section 6 concludes the study.

## 2. Mapping technological trajectories using patent citation networks

Patent citations have been frequently used to measure the ‘importance’ of a specific (patented) innovation. The idea behind this approach is very simple. If a patent is cited very frequently, it means that the patent contains a ‘piece of knowledge’ that forms the basis for several subsequent inventions. Hence, the patent in question ought to be regarded as ‘technologically’ important. Empirical studies have generally confirmed the existence of a positive and significant relationship between the number of citations received by a specific patent and other indicators of technological and economic importance (Jaffe and Trajtenberg 2002).

By looking comprehensively at the evolution of patent citations in a specific technological field, one is typically confronted with a network of patent citations. An intuitive interpretation of these networks is that they can be understood as representing the relationships between the pieces of knowledge contained in the individual patents. Following this reasoning, it should be possible to trace ‘technological trajectories’ in the sense of Dosi (1982) in the evolution of patent citation networks (see Verspagen 2007 for a more extensive treatment).

Additionally, when we take this broader view and consider the network of patent citations in a certain technological field as providing a possible representation of the underlying dynamics of knowledge, it also appears that the traditional approach of measuring the technological importance of patents by looking at the number of their *direct* citations may be fraught with difficulties. As shown by von Wartburg, Teichert, and Rost (2005), assessing the importance of patents, using the number of *direct* citations received, amounts to use ‘centrality’ measures for assessing the importance of nodes in a network. This approach may be clearly justified both on intuitive and on empirical grounds. However, one may also suggest that this exercise ought to be integrated with a study of the whole ‘connectivity structure’ of the network in question. In other words, a proper assessment of the technological significance of a specific patent ought to be based *both* on direct citations *and* on a general characterization of the position of the patent in the overall structure of the network of patent citations.

Several approaches based on network theory may be used for this task. In this paper, we will follow the approach proposed by Hummon and Doreian (1989) which has been already adopted for the analysis of citation networks in different technological fields (Mina et al. 2007; Verspagen 2007). Hummon and Doreian (1989) are interested in identifying what may be considered the main paths of a network which represent the relationship between

pieces of knowledge. These main paths should be viewed as the ‘main flows of ideas’ characterizing the structure of the network in question.

Hummon and Doreian assign a weight to each citation link on the basis of its position in the overall structure of the network. The method is based on the examination of the different ‘search paths’ existing in the network. Search paths are sequences of links that connect the vertices of the network. The method is illustrated in Figure 1A. A more extensive and formal treatment is contained in Verspagen (2007). In the figure the vertices represent the patents, and the edges represent the citations. The edges are oriented from the cited towards the citing patent, so that the arrow represents the direction of the knowledge from the cited to the citing patent. For example, in the figure, patent A is cited by patent C. Patent C in turn is cited by patent D.

A search path is represented by the sequence A–C–D–F–H–J, which would indicate a flow of knowledge from patent A to patent J through several intermediate patents. The simplest indicator proposed by Hummon and Doreian is the so-called *search path link count* (SPLC) measure. The SPLC is calculated by considering all the possible search paths in the

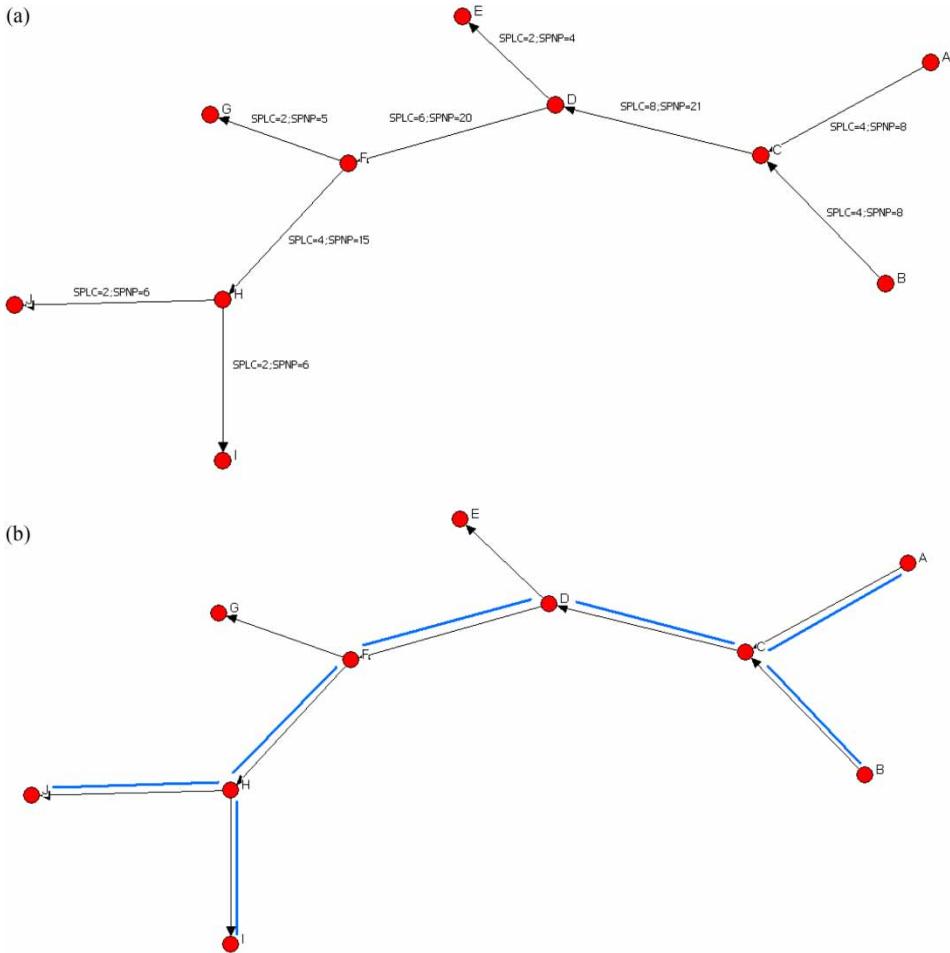


Figure 1. (A) Calculation of SPLC and SPNP. (B) The network of main paths.

network, and counting how often an edge lies on such a search path. Although Hummon and Doreain suggest to include even those search paths involving intermediate vertices as the origin or final destination, one may also use only the search paths from start points to endpoints. In this case, one obtains a value of SPLC of 4 for the edge  $A - C$  (paths  $A \rightarrow E$ ,  $A \rightarrow G$ ,  $A \rightarrow J$  and  $A \rightarrow I$ ), of 8 for the edge  $C - D$  (paths  $A \rightarrow E$ ,  $A \rightarrow G$ ,  $A \rightarrow J$ ,  $A \rightarrow I$ ,  $B \rightarrow E$ ,  $B \rightarrow G$ ,  $B \rightarrow J$  and  $B \rightarrow I$ ) and so on and so forth.

An alternative measure proposed by Hummon and Doreian is the *search path node pair* (SPNP) indicator. Consider the edge  $C - D$ . This edge connects three vertices ( $A$ ,  $B$  and  $C$ ) to its destination (i.e. vertex  $D$ ). At the same time, it is responsible for connecting its origin (vertex  $C$ ) to other seven vertices ( $D$ ,  $E$ ,  $F$ ,  $G$ ,  $H$ ,  $I$ , and  $J$ ). We define the SPNP value for this edge as the product of these values ( $3 \times 7 = 21$ ) because in total 21 distinct pairs of vertices are connected by it. As a result of this multiplication, the SPNP measure 'weights' edges in the middle of a path relatively more than those at the start or at the end. The logic underlying both measures is that the citation links that are responsible for connecting higher number of patents contain the most significant knowledge flows of the citation network.

Once the weights to the links of the network have been assigned using the SPLC or the SPNP procedures (in most of the cases the methodologies lead to similar results), Hummon and Doreian suggest a procedure for constructing the network of 'main paths'. This is represented by the network constructed by moving from each 'start point' towards an 'endpoint' of the network. In doing this, at each node, one must follow the link with the highest SPLC or SPNP value. In the case of a tie, both links are considered as part of the network of main paths. Figure 1B represents the edges of network of main paths for the original network of Figure 1A, as marked by thicker lines. In this case, the construction of the network of main paths does not change if we use as weights of the edges the values of SPLC or SPNP and it involves the deletion of the edges  $D - E$  and  $F - G$ . In other words, this heuristic procedure implements the idea of identifying the most important knowledge flows in the citation network by deleting marginal or redundant edges.<sup>2</sup>

### 3. Data communications technology: a short history and overview

In this paper we apply the patent citation network approach to study the evolution of data communications technology. Given the broad scope of this technological field, we focus on a specific subfield related to technologies and equipment for LANs. Within this subfield we consider in detail the patents related to the Ethernet communication standard. In this section we provide the relevant background information to understand the development of Ethernet and to put it within the more general context of the evolution of technological systems such as LANs.

The simplest way to define a LAN is as a set of communication conduits connecting computers or other devices capable of sending and receiving digital information. These conduits may consist either of ordinary telephone wires or of more sophisticated cabling systems. They may connect devices in different patterns which are referred to as topologies. LANs started to be developed in the early 1970s although experiments with computer networking had originated in the late 1950s and early 1960s.<sup>3</sup> These early LANs were highly heterogeneous and 'proprietary' systems, developed mainly for intra-company use. They provided stable and simple connections between a mainframe (IBM type) or minicomputers (DEC type) and 'dumb' terminals. The major growth in the LAN market was a direct consequence of the advent of microcomputers in the mid-1970s and PCs in the 1980s. While the advent of the microcomputer entailed the possibility of economically providing each user with an end station, the extension of peripherals services (mainly printing) to

single users proved to be very expensive. Tying users together was perceived as a way to reduce the costs while extending the access to an increasing amount of users. This event spurred the diffusion of LANs.

From a technological viewpoint we can identify four phases in the evolution of LAN technology (Christensen et al. 1995). Each phase is associated to the introduction of a particular type of component as well as to the dominance of a specific communication standard. The first phase, from mid-1970s to the mid-1980s, is characterized by the absence of an official communication standard and the presence of many proprietary solutions marketed by different manufacturers. The second phase, from the mid-1980s until the end of the decade, saw the introduction of the equipment called ‘hubs’ and the official ratification of two communications standards (Ethernet and Token Ring). Hubs were introduced mainly to offer a unique point of concentration within offices in order to rationalize the cabling system. The third phase, from the end of the 1980s until mid-1990s, witnessed the growth of the hub market and the establishment of Ethernet as the dominant standard for LANs. As more and more nodes were added to the LAN, congestion increased revealing the shortcomings of hubs in solving new bottlenecks. The introduction of another type of equipment called ‘switch’ and the emergence of fast Ethernet as the high-speed upgrade of Ethernet characterized the fourth phase (Fontana and Nesta 2006).

In the remaining parts of this section we discuss in detail the Ethernet standard. Our aim is to identify the main innovations that have marked its evolution. This type of analysis will provide a roadmap to guide the interpretation of the results of our empirical analysis.

### 3.1. *The evolution of Ethernet*

Together with the diffusion of microcomputers, the invention and commercialization of Ethernet represented a crucial event for the diffusion of LANs. Ethernet was invented in 1973 by Robert Metcalfe at the Xerox PARC research centre in Palo Alto as a part of a project undertaken by Xerox aimed at creating an ‘architecture of information’ in which personal computers (instead of minicomputers), file servers, and laser printers could all be linked by a network. This project was part of a comprehensive strategy by Xerox aimed at extending its core business and embracing the computer technology. The strategy eventually culminated in the development of the Alto PC, later commercialized as Star workstation, in 1973. The establishment of the PARC laboratories, whose activity resulted in the invention of many revolutionary technologies such as laser printers and file servers, was essential for achieving the aim (von Burg 2001, 67–9).

The Xerox network responded to specific technical and economic requirements. First, it was able to transfer data at higher rates than could be achieved by phone wires over WANs because high speed was required to maximize the performance of laser printers. For this reason, Metcalfe decided to employ *coaxial cables* as transmission media, as they could achieve a transmission rate of up to a thousand kilobits per second. Second, it was able to span hundreds of metres to connect all computers deployed throughout a large company campus. To achieve this goal, a *bus topology*, potentially able to branch out in every direction at each point to accommodate additional computers, was chosen. Third, it was *cheap*. Therefore, Metcalfe devised ways of limiting the cost of connection to 5% of the cost of a computer (Metcalfe 1994).

Moreover, the network invented by Metcalfe benefited from three basic developments in data communications which had been achieved in the past.<sup>4</sup> The first development was *packet switching*. Packet switching was a communication method which, by breaking a message into small units (*frames or packets*), allowed the message to be sent in

a discontinuous stream which did not require establishing and maintaining a direct and dedicated link between two nodes for the duration of the transmission as required by the *circuit switching* method applied in telephone networks. The second development was *layering*. Layering consisted of breaking up the communication process into several steps, each performing a specific function in order to reduce the complexity of the whole transmission and enabling manufacturers to focus on the development of protocols within a single layer at a time instead of tackling the whole process at the level of the overall network.<sup>5</sup> The third development was the use of a *network interface* ('transceiver') to connect the computer to the medium of choice and accomplish the tasks of both taking the message from computers and formatting it in the appropriate manner in order to be transported over the LAN to other stations and receiving the messages from the network and reformatting them in a form which was understandable by the local computer.

One of the major problems of designing a network in a non-mainframe environment was to devise a way of allocating access to the channel across all the nodes. To regulate the sequence of accesses and control the duration of each transmission over the network, Metcalfe's designed an 'access method' based on an algorithm which later became known as the carrier sense multiple access with collision detection (CSMA/CD) method.<sup>6</sup>

Metcalfe's ideas led to the implementation of an Ethernet-based LAN in the PARC network in 1974 when, together with David Boggs, he developed and subsequently patented the first transceiver. At first, the outcome of the collaboration between Metcalfe and Boggs remained an isolated attempt.<sup>7</sup> The implementation of the Ethernet on a large scale became possible only with the important developments which took place in the 1980s. The first development was the increase in the speed of Ethernet from 2.94 Mbps to 10 Mbps which occurred after the formation of the DIX Alliance, involving DEC, Intel and Xerox. This alliance was set up in 1979 with the explicit goal of promoting Ethernet as an industry wide *de facto* standard, making the technology available to anybody (von Burg 2001, 102–107). The choice of joining efforts to promote Ethernet as an industry-wide open standard, responded to specific needs and strategies of the three main proponents. These strategies are typical of competition in technical systems characterized by the presence of 'network effects' where the main goal is to reduce costs by widening the installed base of the system components. First, neither Xerox nor DEC saw the LAN technology as their core business. On the contrary, they thought that a wide acceptance of Ethernet would help them promote the diffusion of Ethernet 'complements' such as their computer hardware, the Alto-based office system and the VAX system, respectively. Second, both companies realized that the diffusion of Ethernet would ultimately depend on the availability of other Ethernet components such as transceivers. Promoting Ethernet as an industry wide 'open' standard could lead to the creation of independent suppliers of Ethernet components. Third, among the required components, chipsets were the most important. Intel thought that a strategy of open licensing could help achieve the volumes necessary to spread the high costs of manufacturing and designing them (von Burg 2001, 102–7).

The second development was the starting of an official standardization process within the Institute of Electrical and Electronic Engineers (IEEE hereafter) eventually leading to the official (de-jure) standardization of Ethernet in 1985. Ethernet started to be considered for standard definition by the IEEE Project 802 group in February 1980. This effort overlapped with the attempts of the DIX Alliance to achieve de facto market dominance.<sup>8</sup>

In 1981, 3Com, the Xerox spin-off funded by Metcalfe, shipped the early transceiver and controllers compatible with DEC minicomputers. In the same year, DEC in collaboration with AMD and Mostek started developing early integrated circuits for Ethernet transceivers. It was quickly followed by other manufacturers such as 3Com and Ungermann Bass which

struck a collaboration agreement with Seeq Corp. and Fujitsu respectively. Further milestones were linked to variants to the original design. In August 1981, 3Com redesigned the transceiver and the Ethernet board to shrink size in an attempt to target the PC market. At about the same time a new version of Ethernet for thinner coaxial cables was developed.<sup>9</sup> Finally, the end of the decade witnessed the standardization and commercialization of a version of Ethernet (10Base-T) running over twisted pair copper cables.

The diffusion of Ethernet occurred in parallel with the development of another communication technology: Token Ring. Token Ring was based on the early work of David Farber who, in the early 1970s at University of California Irvine, had invented a network based on an access method which was different from the one underlying the functioning of Ethernet. In a Token Ring network, transmission was controlled by a bit pattern (the token) which was sent continuously and uni-directionally around the ring. Stations willing to transmit seized the token and passed it by with a destination address. When data were received, the token returned to the sender node which had to pass it around free before transmitting again, thereby avoiding monopolizing the communication. Contrary to CSMA/CD, this access method was deterministic because it made it possible to calculate the time interval between successive appearances of the free token. In a Token Ring network, each node could act as a repeater of the signal so there were no limits on speed or distance as in the bus topology. Moreover, it could be deployed either in a ring or in a star topology as opposed to the bus topology of Ethernet. A ring topology gave Token Ring a competitive advantage over Ethernet whose bus topology was ideal for connecting a small number of stations within a building, but could prove to be ineffective for connecting stations between buildings (Sirbu and Hughes 1986; von Burg 2001, 62–7).

Variants of the ‘Farber ring’ flourished during the 1970s such as the Cambridge Ring developed in 1974 or the ‘star shaped ring’ developed at MIT by Saltzer, Pogran and Clark. In parallel, IBM had been experimenting with Token Ring over its own computer systems.<sup>10</sup> At the beginning of the 1980s, other manufacturers had started shipping Token Ring products targeting several kinds of environment. In 1981, Proteon, a start-up company, had shipped a proprietary 10 Mbps Token Ring to be used by universities and research institutes. LANs, mainly based on proprietary ‘token based’ topologies were also diffused in factory environments as opposed to offices (Bresnahan and Chopra 1990). In 1984, the Token Ring standard was first drafted by the IEEE.

Since most of its characteristics were already known, several manufacturers, besides IBM, were supporting Token Ring in the competition against Ethernet. Two elements were crucial in the evaluation of which standard to adopt. The first element was the kind of media supported. From this viewpoint, Token Ring held a cost advantage since Ethernet required coaxial cables which were more expensive than the Shielded Twisted Pair (STP) wires chosen for Token Ring. The second element was the extension of the configuration. Again Token Ring had a cost advantage. Because of its bus topology, Ethernet usually required a higher (and more expensive) total length of cabling than the ring topology. One technical advantage of Ethernet over Token Ring was its higher speed (10 Mbps vs. 4 Mbps). The combination of a more expensive wiring system and a wider area to connect in most cases represented a disadvantage on the demand side for Ethernet compared with Token Ring. Nevertheless, the demand for LAN connections turned out to be relatively inelastic with respect to the connection cost.<sup>11</sup> Moreover, costs were deemed to decrease quite fast as a consequence of technical change coupled with the strategy of open licensing *deliberately* pursued from the start by the DIX Alliance and culminated in the standardization effort by the IEEE. Towards the end of the 1980s, the Ethernet was able to attract more support from manufacturers than Token Ring.<sup>12</sup> A mix of technological characteristics, standardization

strategy and temporal advantage had given decisive momentum to Ethernet when compared with the competing alternative.

### 3.2. Further developments and the arrival of the fast Ethernet standard

The standardization of both Ethernet and Token Ring accelerated network expansion and represented a turning point in the diffusion of LANs. From a technical viewpoint, expanding LANs and developments such as the integration of voice and video transmission with data transmission also brought an increase in the quantity of data and information that would be transmitted which increased congestion and reduced performance. The most straightforward way to reduce congestion was by increasing the aggregate LAN bandwidth by developing new standards. The definition, development and commercialization of new LAN standards, characterized the 1990s when, as the result of the efforts of manufacturers co-ordinated by official standardization bodies, three high speed standards for LANs were developed: Fast Ethernet, 100VG-AnyLAN and Asynchronous Transfer Mode (ATM).<sup>13</sup>

Fast Ethernet and 100VG-AnyLAN were both developed by the IEEE as a result of a long and complex process of standardization during which proponents could not find an agreement on just one of two alternative proposals.<sup>14</sup> ATM instead was a technology originally developed in the fields of telephony with the ambition of becoming viable for the LAN environment.<sup>15</sup> These standards followed divergent trajectories of development and only Fast Ethernet eventually became dominant for high-speed communication.

Pre-standard Fast Ethernet products were introduced by Grand Junction at the end of 1993. Early adapter cards for Fast Ethernet were instead shipped in 1994.<sup>16</sup> A significant milestone was the implementation in the hardware of adapter cards of 'auto-negotiation',

Table 1. Milestones in the evolution of LANs.

Year	Ethernet	Other LAN standards and equipment
1970–71		Farber's UC Irvine Ring
1973	Experimental Ethernet at Xerox PARC	
1974	Design of early Ethernet transceiver	
1980	Bluebook v. 1 by DIX consortium	Star shaped Token Ring at MIT
1981	Ethernet transceivers and controllers for minicomputers	
1981	Thick Ethernet (10Base-5)	
1981	AMD and Mostek deliver ICs for Ethernet	
1981	3Com delivers Thin Ethernet (10Base-2) and transceivers and boards for PCs	Proteon delivers a proprietary 10Mbps Token Ring
1982	Bluebook v. 2 by DIX consortium	
1985	Ethernet standard ratified (IEEE 802.3)	SynOptics delivers first Multiport Ethernet Hub
1985	Thin Ethernet (10Base2)	
1987	IEEE begins official standardization of 10Base-T	
1988		16Mbps Token Ring for STP wire
1990	10Base-T standard is ratified (IEEE 802.3i)	Kalpana delivers the first Ethernet LAN switch
1991		16Mbps Token Ring for UTP wire
1992	IEEE begins official standardization of Fast Ethernet	
1993	Grand Junction delivers first Fast Ethernet product	
1995	Fast Ethernet standard ratified (IEEE 802.3u)	

a mechanism capable of auto-detecting both network speed and configuration of the equipment they were attached to. This mechanism, while letting different segments of an existing network operate at different speeds, enabled network managers to know all the basic information about the features of the adapter, provided that it was implemented in all the equipment within the same segment. Moreover, a further capability of supporting different cabling types (coaxial as well as twisted pair) on the same Ethernet environment was offered by the presence of a media independent interface (MII) mechanism standardized by the IEEE (Melatti 1994).

Table 1 summarizes the milestones in the evolution of Ethernet that have been reviewed in this section together with other main events that have characterized the development of LAN technology.

#### 4. Data and summary statistics

In this paper, we use US patent data as our main data source. The technological field of LAN technology is delimited using the International patent classification (IPC) which seems better suited for our purposes than the classification provided by the US patent office. We consider patents contained in the subclasses of class IPC H04L12 which refers to ‘data switching networks’. A recent analysis by Grebel, Krafft, and Saviotti (2006) has shown that this class has been the most dynamic within the field of telecommunications, experiencing, especially during the 1990s, the highest rate of growth in terms of patents. To check whether the evolution of LAN standards and technologies is in part responsible for this growth, we perform a detailed analysis of patenting activity within this class.

We proceed as follows. First, we take a general look at the patents for the LAN technology within IPC H04L12. Working at a high level of dis-aggregation (seven digits) and on the basis of expert knowledge, we select 10 subclasses as relevant for studying the evolution of LAN technical systems. Table 2 reports the technological subclasses that we have considered relevant for the domain of LAN technology together with their share in the total number of patents. Patent numbers belonging to these classes have been retrieved using the search engine available on the USPTO website.<sup>17</sup> The source for patent information (granted year, assignee and citations) is the NBER patent data set (Jaffe and Trajtenberg 2002).

Table 2. Patent classes and definition.

	IPC patent class	Definition	No. of patents
(1)	H04L12/28	Data switching networks characterized by path configuration (i.e. LAN)	1767
(2)	H04L12/413	Bus networks with random access (i.e. CSMA/CD)	581
(3)	H04L12/417	Bus networks with deterministic access (i.e. token system)	170
(4)	H04L12/43	Loop networks with synchronous transmission (i.e. Time Division Multiplexing method)	101
(5)	H04L12/433	Loop networks with asynchronous transmission (i.e. token ring method)	319
(6)	H04L12/44	Star or tree networks	369
(7)	H04L12/46	Interconnection of networks	809
(8)	H04L12/50	Circuit switching systems	129
(9)	H04L12/52	Circuit switching systems using time division techniques	102
(10)	H04L12/56	Packet switching systems	6137

Source: International Patent Classification, 8th edition, version 2007.1.

The time span covered by the original data set is 1963–1999. For this paper we have extended the original NBER data set adding the patents granted up to the period of February 2002. Our final data set has been constructed by matching the patent numbers in our technological classes of interest with the corresponding number in the extended NBER patent data-set.

It can be noted that the distribution of patents across subclasses is uneven with four subclasses, (1), (2), (7) and (10), accounting for more than 80% of the total patents in the sample. Figure 2 shows the share in the total number of patents over time for each subclass. For dating the patents, throughout this paper we have considered the year in which the patent was granted. In the figure we have highlighted the shares for the four subclasses mentioned above.

The growth in the shares is particularly evident in the case of H04L12/28 (i.e. subclass (1)) and H04L12/56 (i.e. subclass (10)). Subclass H04L12/413 (i.e. subclass (2)) experiences the highest share between 1983 and 1989 and then we witness a further increase between 1995 and 1998 following the standardization of Ethernet and Fast Ethernet, respectively. The share of H04L12/56 (i.e. subclass (10)) is always high and increases considerably from 1987 onwards (particularly during the 1990s). Given the time lag between the application and the granting date, the expansion seems attributable to the arrival of the switch as well as to the implementation of high speed standards in use.

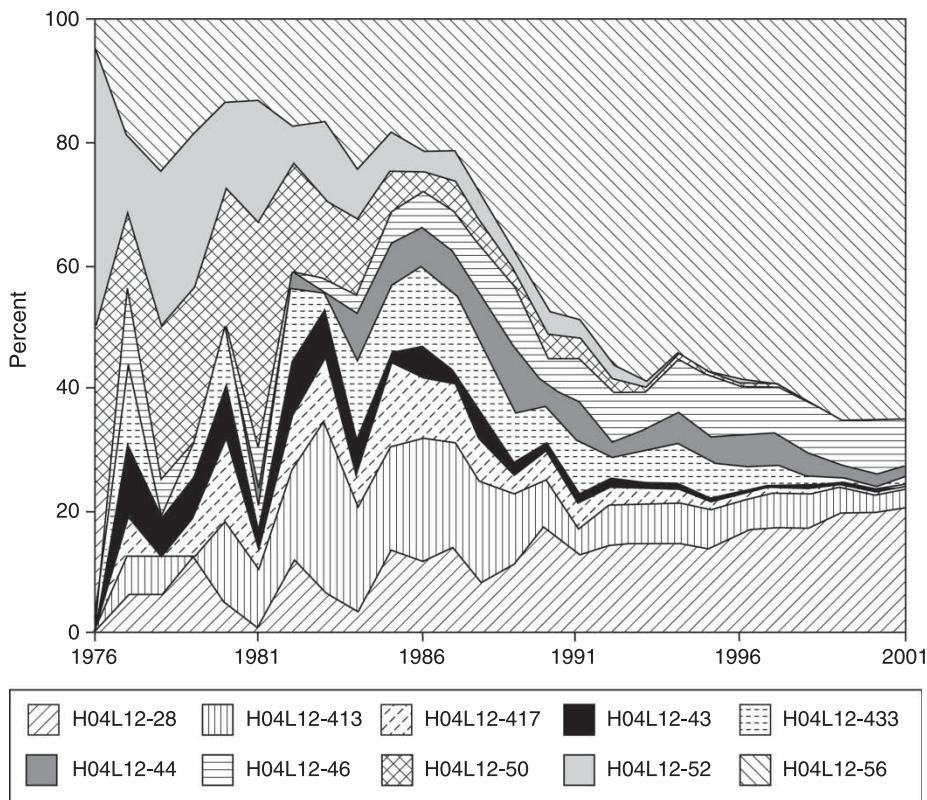


Figure 2. Share of LAN patents for selected subclasses.

Table 3 presents the pattern of cross-citations for all the chosen technological classes together with a summary statistics of the number of citations made and received.

Each cell  $(i, j)$  reports the number of times a patent is cited or cites another patent in our sample. Reading the table by column gives us the number of times patents in class  $j$  have been cited by patents in class  $i$ . Reading the table by row gives us the number of times patents in class  $i$  have cited patents in class  $j$ . The total number of citations received by each class is reported at the bottom of the table. The last column reports instead the total number of citations made by each class.

The following observations can be made. First, for each class the total by column is generally higher than the total by row, thus indicating that the chosen technological classes generally tend to be ‘recipients’ rather than ‘sources’ of citations. Important exceptions are represented by columns (1), (7) and (10) which, however, correspond to classes with relatively higher number of patents. Second, cells along the main diagonal (i.e. those highlighted in bold) generally contain relatively high values when compared with the rest. This indicates that patents of each class tend to cite and be cited mostly by patents of the same class. This suggests that the chosen classes represent relatively autonomous and coherent technological fields in our selection of the LAN technology domain.

It is worth remarking the peculiarities of patents in class (10). Around 78% of citations received by this class are self-citations. Moreover, patents in this class cite a relatively high number of patents belonging to other technological classes. Class (10) contains patents relative to packet switching which is a relative broad definition encompassing a variety of functions and aspects of network technology, ranging from the definition and design of equipment to the implementation of specific standards in use. It is understandable that many patents in this class refer to other more specific classes. An example is provided by class (2) which tends to receive most citations by classes (6) and (10). Class (2) contains patents referring to the Ethernet standard while class (6) and class (10) patents can be broadly related to the design of LAN equipment such as hubs and switches, respectively. Being LANs technical systems, most of these citations refer to the implementation of the standard specifications in use.

Table 3. Citations made and received by IPC class.

IPC Class	Cited										TOT citing (j)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)		
Citing												
(1)	<b>2723</b>	412	185	71	127	162	596	24	20	2,849	7,169	
(2)	316	<b>1116</b>	151	20	82	315	189	9	6	466	2,670	
(3)	75	124	<b>266</b>	16	121	25	32	2	2	97	760	
(4)	50	9	16	<b>84</b>	66	9	15	13	12	136	410	
(5)	66	51	121	97	<b>580</b>	56	189	9	7	253	1,429	
(6)	179	371	47	41	127	<b>552</b>	379	40	38	548	2,322	
(7)	579	250	86	81	294	445	<b>3008</b>	42	24	6,217	11,026	
(8)	5	4	5	4	9	5	9	<b>57</b>	52	63	213	
(9)	4	2	4	3	5	2	3	53	<b>50</b>	24	150	
(10)	2732	854	371	254	886	1200	5259	163	105	<b>37,211</b>	49,035	
TOT cited (i)	6729	3193	1252	671	2297	2771	9679	412	316	47,864	75,184	

(1): H04L12/28; (2): H04L12/413; (3): H04L12/417; (4):H04L12/43; (5): H04L12/433; (6): H04L12/44; (7):H04L12/46; (8): H04L12/50; (9): H04L12/52; (10): H04L12/56.

To check whether the results of Table 3 are affected by the relative differences in the total number of patents in each chosen class, we normalize the observations by the number of total patents in the chosen class. In particular, we construct the following symmetric relative specialization index (RSI).

The RSI is given by

$$RSI = \frac{AI - 1}{AI + 1}, \quad (1)$$

where AI is the activity index calculated as the share of the citations received by each technological class and the share of the technological class in the total number of citations (received and made). It takes values in the  $[-1, 1]$  range.

$$AI = \frac{c_{ij} / \sum_i c_{ij}}{\sum_j c_{ij} / \sum_i \sum_j c_{ij}}, \quad (2)$$

where  $i$  refers to citations received,  $j$  refers to citations made and  $c_{ij}$  to the number of citations. Since the RSI is the ratio of AI minus 1 and AI plus 1, it indicates whether a class has a higher than average citation activity in a given class ( $RSI > 0$ ), or a lower than average ( $RSI < 0$ ).

Table 4 reports the RSI for each technological class.

Again numbers on the main diagonal refers to self citations and are always positive, thus signalling a higher than average intra-class citation activity. This result confirms the previous results in Table 4. Focusing on specific subclasses, some interesting patterns emerge. First, subclass (2), which contains Ethernet related patents, receives a higher than average number of citations from subclasses (3) and (6) that refer to Token Ring and LAN equipment related patents, respectively. As already mentioned, most of the citations from class (6) constitute attempts to implement the standards specification in hubs. Second, some subclasses tend to ‘specialize’ in the sense that they tend to be either sources or recipients of citations. Cells in row (6) for instance are almost always positive indicating that this subclass tends to cite more than average the other classes. Off diagonal cells for subclass (10) are instead always negative indicating that this subclasses displays a lower than average citation activity both as a recipient and as a ‘source’ of citations. Finally, it is interesting to note how some subclasses display a similar pattern of citations. This is particularly true for subclasses

Table 4. Relative specialization index.

IPC Class	Cited									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Citing										
(1)	<b>0.619</b>	0.150	0.216	0.052	-0.266	-0.240	-0.215	-0.242	-0.202	-0.231
(2)	0.139	<b>0.816</b>	0.545	-0.087	0.003	0.524	-0.290	-0.238	-0.303	-0.570
(3)	0.049	0.587	<b>0.909</b>	0.405	0.678	-0.057	-0.507	-0.351	-0.230	-0.666
(4)	0.153	-0.319	0.402	<b>0.917</b>	0.681	-0.253	-0.557	0.705	0.749	-0.315
(5)	-0.319	-0.087	0.671	0.768	<b>0.860</b>	0.031	0.013	0.069	0.076	-0.565
(6)	-0.075	0.580	0.097	0.329	0.283	<b>0.732</b>	0.118	0.517	0.591	-0.459
(7)	-0.260	-0.304	-0.362	-0.097	-0.068	0.045	<b>0.359</b>	-0.180	-0.318	-0.061
(8)	-0.584	-0.387	0.170	0.356	0.161	-0.222	-0.506	<b>0.960</b>	0.966	-0.366
(9)	-0.541	-0.522	0.231	0.383	0.044	-0.469	-0.731	0.969	<b>0.975</b>	-0.598
(10)	-0.233	-0.418	-0.375	-0.265	-0.257	-0.202	-0.091	-0.245	-0.325	<b>0.088</b>

(1): H04L12/28; (2): H04L12/413; (3): H04L12/417; (4):H04L12/43; (5): H04L12/433; (6): H04L12/44; (7):H04L12/46; (8): H04L12/50; (9): H04L12/52; (10): H04L12/56.

(8) and (9) which, as we saw in Table 3 display a low number of citations and tend, on average, to make and receive citations from the same subclasses. Interestingly, patents in these subclasses refer to ‘circuit switching’, a type of technology linked to telephony and not so widespread in the realm of data communications.

## 5. Patent citations networks in the LAN technology

In this section we focus on the structure of connectivity of a specific subclass in our overall patent population for LAN technology. We perform the structure of connectivity study on subclass H04L12/413 which pertains to the Ethernet standard. Figure 3 provides a general overview of the development of the network of patent citations for this specific subclass over time. There are six series in the graph: ‘PATENTS’ represents the cumulative number of patents of this subclass, the series ‘START’ refers to the number of start points, the series ‘END’ represents the number of endpoints, and the series ‘ISOLATES’, ‘INTERMEDIATES’ and ‘COMPONENTS’ represents the number of isolate vertices, the number of intermediate points and the number of network components,<sup>18</sup> respectively. The number of patents and components are measured on the vertical left and right axes, respectively.

Two observations are in order. First, Figure 3 suggests that all the series display a behaviour characterized by an almost constant growth rate which tends to slightly level off toward the end of the time period. Our interpretation is that the pattern reflects a sustained phase of exploration of the space of technological opportunities, corresponding to the subsequent generations of the standard (i.e. Ethernet, and Fast Ethernet) which constantly fuel technical progress in this subclass. Second, the number of components is also increasing. This trend would suggest that the development of the technology occurred in different environments without displaying a strong tendency towards technological convergence (if this was the case, we would instead see a decrease in the number of components, as they become increasingly interconnected). However, the overall network, since the early 1990s,

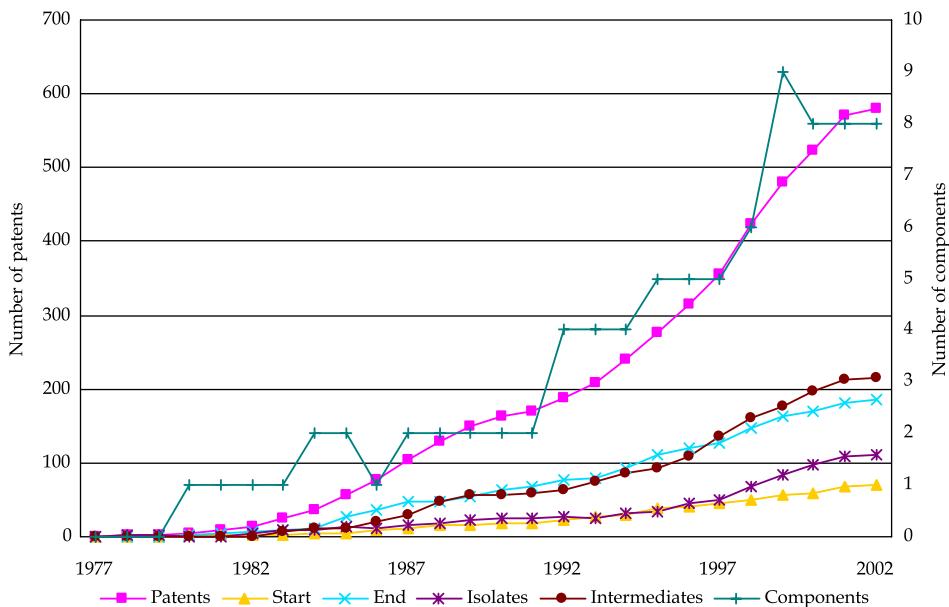


Figure 3. Development of the Ethernet patent citation network over time.

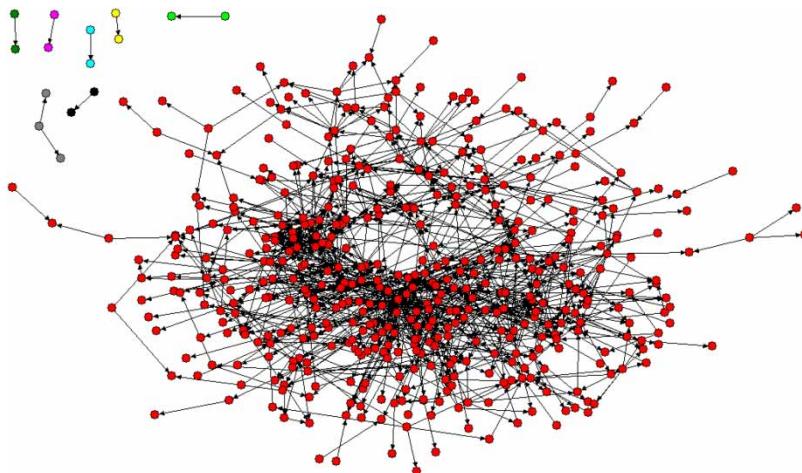


Figure 4. Patent citation network for the Ethernet class, 1976–2002.  
Note: Isolates have been deleted.

is characterized by the presence of a largest component which contains the majority of patents and by a number of smaller components and isolated nodes. In 2002 the largest component, which is displayed in Figure 4, contains 456 patents out of a total of 581. This corroborates the idea of considering the patents in the Ethernet class as belonging to one integrated ‘technical system’.

As a next step in our analysis, we calculate the value of SPNP for each citation pair of the complete network spanning from 1977 to 2002.<sup>19</sup> Then, following the procedure proposed by Hummon and Doreian (1989), we construct the network of main paths as explained in Section 2.

Figure 5 displays the largest component of this network of main paths. This network is characterized by a rather long sequence of patents contained in the bottom right part of the graph. The sequence from patent 4096355 to patents 6078591, 6118787 and 6178177 represents what we may call the ‘top path’ of the network (i.e. the sequence of patent citations going from a start point to an end point whose sum of SPNP values is highest). The second highest ‘top path’ is represented by the sequence from patent 4332027 to again patents 6078591, 6118787 and 6178177. The third highest ‘top path’ contains the sequence from patent 4063220 to again patents 6078591, 6118787 and 6178177. These paths are represented in Figure 5 with white dots. Intuitively, these sequences of patents should represent the fundamental flow of knowledge (technological trajectory) in the overall network of patent citations. These patents are listed in more detail in Table 5. The following points merits attention.

First and foremost, the path actually originates from two rather distinct trajectories. Patent (No. 4063220) lies at the origin of the first trajectory. This is the patent by Metcalfe et al. which sets the basis of the CSMA/CD contention mechanism that characterizes the functioning of Ethernet. The next three patents on this trajectory (No. 4271523, 4345250, and 4560984) represent attempts at dealing with the design of the collision mechanism within a *baseband network environment*. The second and the third trajectories originate instead from Patent (No. 4096355) and Patent (No. 4332027), respectively, and continue through three subsequent patents (No. 4210780, 4412326, and 4598285), all concerned with the design of a collision detection mechanism in *broadband network environments*.<sup>20</sup> The

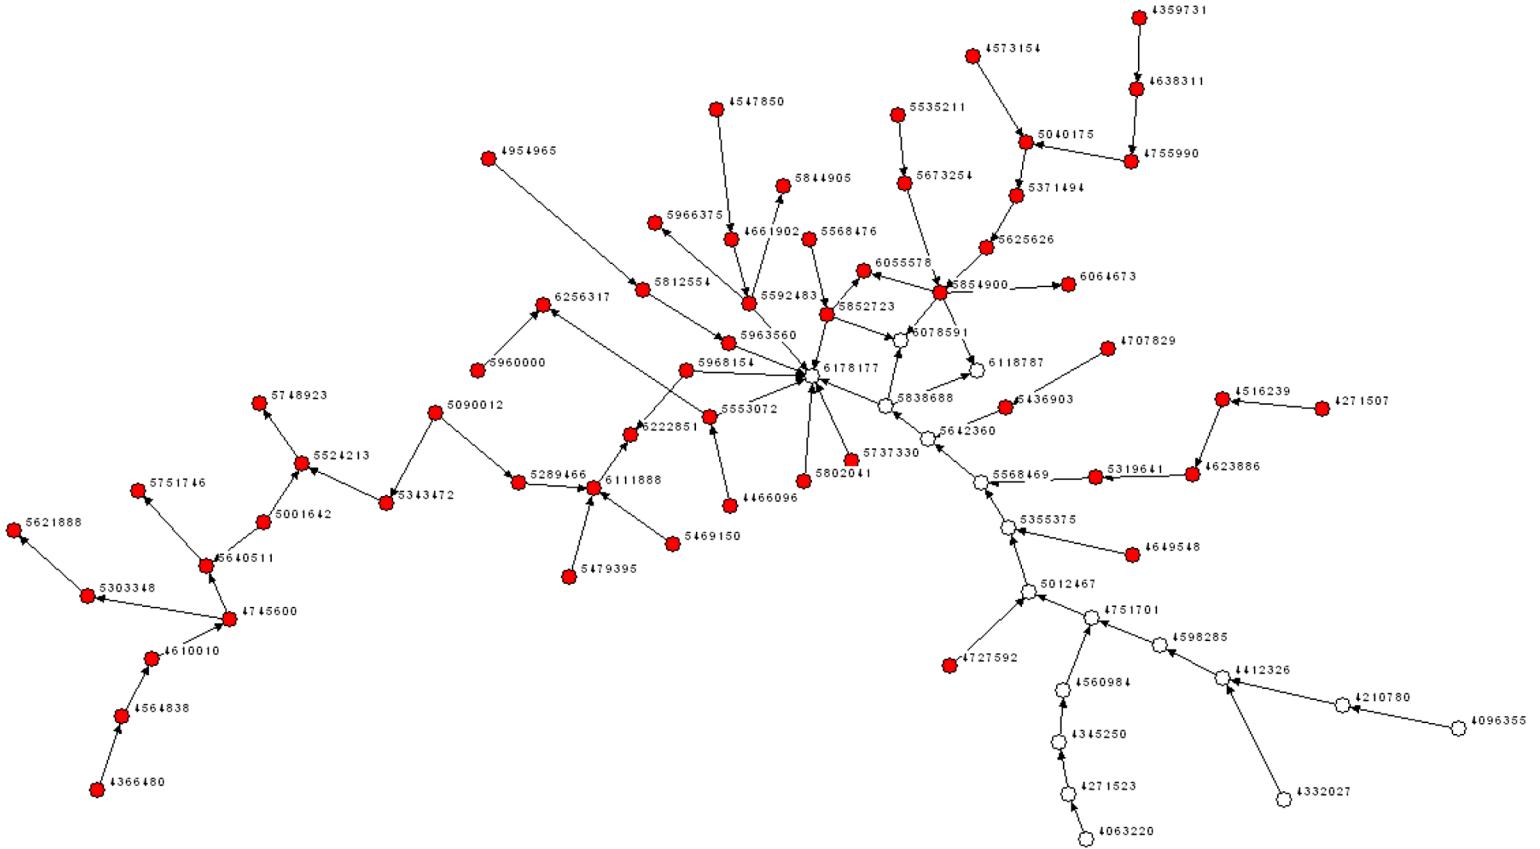


Figure 5. The largest component in the network of ‘main paths’ for the Ethernet network, 1977–2002.

Table 5. The longest sequence of patents on the main path.

Pat. No.	App. Year	Assignee	Patent content
4063220	1975	XEROX Corp.	Multipoint data communication system with collision detection
4096355	1977	IBM	Common channel access method for a plurality of data stations in a data transmission system and circuit for implementing the method
4210780	1978	The Mitre Corp.	Multiple access digital communications system
4271523	1979	Ford Motor Company	Contention interference detection in data communication receiver
4332027	1980	Burroughs Corp.	Local area contention network data communication system
4345250	1980	Honeywell Information Sys.	Information communication system with collision avoidance
4412326	1981	Bell Telephone Labs	Collision avoiding system, apparatus and protocol for a multiple access digital communications system including variable length packets
4598285	1983	AT&T Bell Labs	Scheme for reducing transmission delay following collision of transmissions in communication networks
4560984	1984	Codenoll Technology Corp.	Method and apparatus for detecting the collision of data packets
4751701	1985	Hughes Network Systems	TDM collision detector
5012467	1989	3Com Corp.	Method and apparatus for collision detection in a local area network transceiver
5355375	1993	Network Systems Corp.	Hub controller for providing deterministic access to CSMA local area network
5568469	1994	3Com Corp.	Method and apparatus for controlling latency and jitter in a local area network which uses a CSMA/CD protocol
5642360	1995		System and method for improving network performance through inter frame spacing adaptation
5838688	1996	Advanced Micro Devices, Inc.	Determining the number of active nodes on an Ethernet network by counting a number of packet receptions
6078591	1997	Advanced Micro Devices, Inc.	Apparatus and method for selectively modifying collision delay intervals
6118787	1997	Advanced Micro Devices, Inc.	Apparatus and method for regulating assigned bandwidth in high speed packet switched networks
6178177	1997	Alcatel	Data-processing network having non-deterministic access, but having deterministic access time

two trajectories converge on patent (No. 4751701) owned by Hughes Network Systems. This patent deals with the design of a transmission protocol for a new detection mechanism based on detection of a signal gap in the transmission medium following collision rather than detection of the presence of simultaneous energy emission from transmitters. Interestingly, this is a solution that seems to allow solving contention problems in both baseband and

broadband network environments and justifies the position of this patent at the ‘junction’ of the two trajectories. All these patents, both baseband and broadband related, are quite specific and revolve around the particular issue of the design of the contention mechanism, and are relatively old. The next six patents (No. 5012467, 5355375, 5568469, 5642360, 5838688, and 6078591) are less specific and more recent. They deal with various aspects of the implementation of the standard in use related to the design of various components of the LAN technical system such as repeaters and hubs, as well as the performance of these types of equipment. The nature of these patents will be analysed in more details below.

Second, patents belong to companies that have been the major players in the birth and growth of the networking industry such as Xerox, IBM, and 3Com. However, as can be seen from Table 6, no company seems to be ‘dominant’ in the sense of claiming ownership of the majority of the patents that lies on this fundamental path that we have identified in the network. This suggests that no company seems to be strategically placed along the main path of knowledge flow and that the clustering reflects a general engineering logic rather than the logic based on the internal search strategy of a specific company. This is also consistent with the literature which has emphasized how Ethernet achieved market dominance as a consequence of licensing and the pursuit of an open standard strategy (von Burg 2001).

Finally, a comment on the relationship between the number of citations and the positioning of the patent along the top path is in order. Table 6 reports the list of the 20 most cited patents in the Ethernet subclass.

While some of the patents that lie on the path we have identified are also included in the list of the most cited patents, it can be noted that some important patents, such as patent

Table 6. Top 20 Ethernet patents by internal citations.

Patent No.	Internal citations	Total citations
<b>4063220</b>	<b>74</b>	<b>131</b>
<b>4332027</b>	<b>25</b>	<b>54</b>
5319641	21	29
<b>4412326</b>	<b>19</b>	<b>38</b>
5418784	19	24
5353287	18	25
5436903	18	23
5526355	16	21
4715031	15	20
<b>4210780</b>	<b>14</b>	<b>36</b>
4234952	14	26
5012467	14	22
5121382	14	20
4429384	13	14
<b>4598285</b>	<b>12</b>	<b>18</b>
4623886	12	22
<b>4751701</b>	<b>12</b>	<b>18</b>
5249183	12	23
<b>5642360</b>	<b>12</b>	<b>14</b>
5446735	11	24

Note: Internal citations refer to citations within class H04L12/413 (Ethernet). Total citations refer to the total number of citations within the 10 patent classes covering Data Communication Equipment. Patents marked in bold belong to the top path sequence identified in Figure 5.

No. 4751701 discussed earlier, do receive a relatively lower number of citations and/or are not included in the list. In this case, it would seem that the Hummon and Doreian method has allowed the identification of a patent that, although not characterized by many citations, is positioned at a strategic 'junction' along the trajectory of technological development, linking the broadband and baseband environment. However, a qualification is in order. There are other patents, or groups of patents, that in the underlying network of citations represented in Figure 4 could serve the same critical role.<sup>21</sup> However, if one accepts that the Hummon and Doreian approach is able to pin down the most significant knowledge flows, it is likely that such alternative patents or group of patents will not embody the extent of knowledge flows passing through patent 4751701. The reason for this is that the knowledge flow measure we use does not only depend on direct citations, but especially on indirect citation links.<sup>22</sup>

Values of SPNP can also be calculated for any time interval. Thus, by repeating this exercise for a sequence of time periods, it is possible to produce a series of 'snapshots' that portrays the evolution of the patent citation network of Ethernet over time. This exercise provides an interesting perspective on the dynamics of innovation in large technical systems. We may expect that when inventive activities are guided by a consolidated search heuristics patents of the temporal sequence of top paths will show a high degree of overlap. Vice versa, in phases of high uncertainty we could expect a relatively heterogeneous temporal sequence of paths. The evolution of the top path over time is portrayed in the sequence of Figure 6. Note that the years in the figure indicate the intervals for which the networks were calculated rather than the years in which the patents were granted.

In our case, we observe that the patents at the origins of the two trajectories of the top path of Figure 5 appear relatively early in our sequence of 'snapshots'. However, the baseband trajectory seemed to evolve relatively more rapidly than the broadband trajectory which took off only between 1984 and 1988. It is also interesting to note that in the later stage of development, the two trajectories first converge and then branch off into separated technological environments. This evolution is clearly represented in the last snapshot of our sequence (1977–2003), in which nodes of different shapes and colours identify patents belonging to distinct environments. In the remaining part of this section, we analyse the specificities of these clusters and the extent that they are related to particular events in the development of the Ethernet standard.

Shaded circles identify the first technological environment characterized by 17 patents. These patents are similar in the sense that they all are variations on the Metcalfe et al. 1977 patent (no. 4063220). However, each sequence is characterized by a specific research heuristic. Baseband related patents lie in the lower sequence (starting from patent no. 4063220). Each of them embodies some sort of modification of the mechanism aimed at improving collision detection.<sup>23</sup> Broadband related patents instead lie in the upper sequence (starting from patent no. 4096355). While all these patents generally seek to improve the collision detection mechanism as in the baseband case, this improvement is mainly seen as instrumental to the achievement of a higher transmission bit rate in principle obtainable in broadband networks. Patent (No. 4751701), owned by Hughes Networks and discussed above, is the prelude to the convergence of the two sequences. Interestingly, convergence occurs right after 1985 the year in which the IEEE officially completed the standardization of the Ethernet, as if the official ratification reduced uncertainty and put an end to a phase of exploration of different alternatives.

The main core of the network branches off in four different environments. The environment identified by white circles was the earliest to develop. Innovations belonging to this cluster aim at simplifying the design of the network and at improving reliability

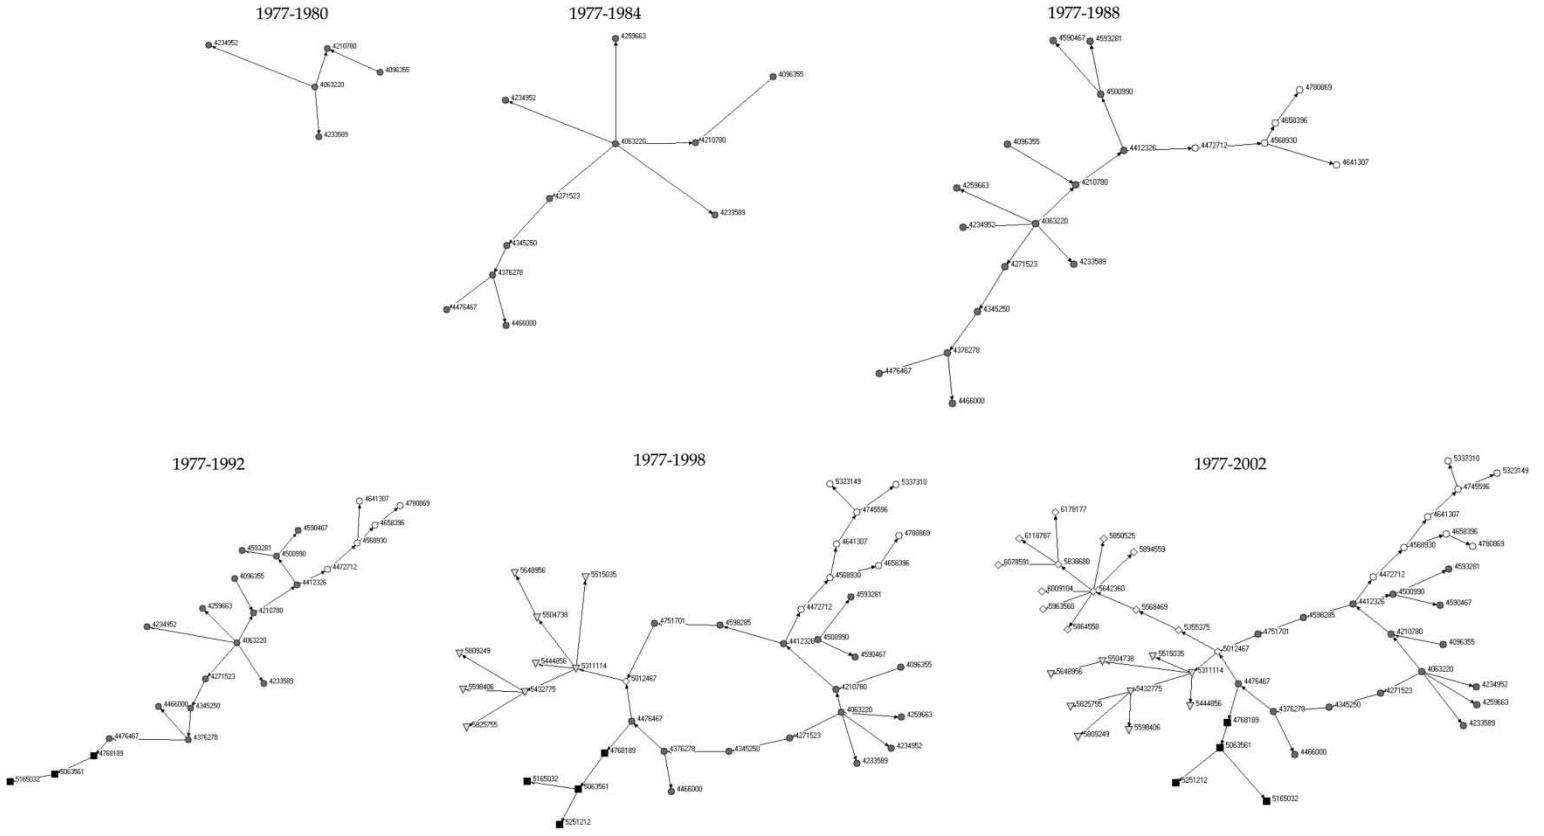


Figure 6. The evolution of the top path over time 1977–1980.

of transmission. Simplification was mainly sought through the reduction of the number of transceivers along the network (see patent nos. 472172 and 4568930 for instance). Reliability instead tended to be achieved by improving cabling redundancy as in patents (Nos. 4658396 and 4780869). Black squares characterize a second cluster of four patents. These patents pertained mainly to applicability of contention detection mechanisms to context characterized by a plurality of buses. The cluster does not grow after 1991 and it does not seem to be particularly important for our analysis of the subsequence evolution of the Ethernet trajectory.

The two environments that branched off from patent no. 5012467 deserve instead to be analysed in more detail since they both mark a discontinuity with the early phase of exploration. While early improvements were mainly characterized by variation upon 'the collision detection mechanism theme', these two clusters are mainly related to applications of Ethernet in use and therefore to the design of the components of the overall technical system. The cluster of diamonds refers to equipment such as transceivers and hubs that needed to be deployed in LANs as they started growing in extension. The patent that lies at the origin of this cluster is a design of a method and apparatus for collision detection in a transceiver (patent no. 5012467). This patent is owned by 3Com and the inventor is Ronald Crane who, as we saw in Section 3.1, in 1981 adapted the early transceiver designed by Metcalfe and Boggs to shrink equipment size and open for Ethernet the way to the PC market. Other patents include the design of a controller for hub (patent no. 5355375) and a method for improving hub performance by reducing latency and jitter (patent no. 5568469).<sup>24</sup> The rest of the patents within this cluster deal with data packets control flow. Interestingly, five out of seven of these patents are owned by Advanced Micro Devices. This evidence suggests that in the most recent years the search trajectory has become more and more company specific and moved away from the 'open' engineering logic that seemed to prevail in the early years of the evolution of the trajectory. Moreover, this is also consistent with recent developments in which semiconductor companies are increasingly challenging traditional networking companies as innovators, a pattern that may lead to a change in leadership in the LAN industry.

Similar considerations apply to the cluster indicated by the triangles at the bottom left-hand side of Figure 6. All these patents refer to the Full Duplex implementation of Ethernet which occurred in parallel with the arrival of switching. In Full Duplex, data can be sent and received simultaneously over a link. One advantage of this approach is that the full duplex link can theoretically provide twice the bandwidth of normal (half duplex) Ethernet. The full duplex mode of operation requires that each end of the link only connects to a single device, such as a workstation or a switched hub port. This represented a paradigmatic change with respect to the functioning of early networks based on the principle that the total bandwidth available for communication had to be shared among all the users (Spurgeon 2000). Within this cluster, an important sequence is constituted by the three patents (Nos. 5311114, 5504738, and 5648956). All these patents pertain to the design of a transceiver for Ethernet 10Base-T and are owned by Seeq Technology Inc., a spin-off from Intel, which was among the first companies to develop a Full Duplex version of Ethernet and to market it in 1993 through its subsidiary Kodiak Technology. The other patents in this cluster are ancillary to this main trajectory. Indeed, the implementation of Full Duplex features required further changes in the hardware specification of the equipment such as the development of 'auto-negotiation' which has the capability of switching from half duplex to full duplex (patent no. 5432775) or between different media specification of Ethernet (patent no. 5444856). Also these patents are owned by traditional semiconductor firms (Advanced

Micro Devices and Intel, respectively) somewhat confirming the increasing importance of these companies in the LAN industry.

## 6. Concluding remarks

This paper has analysed the pattern of technical change in large technical systems as captured by patent citations networks. Given the characteristics of innovation in these systems (i.e. distributed inventive activity, ‘systemic nature’, cumulativity, and speed of technical change), we have argued that it is possible to use measures of structure of connectivity such as the *search path link count* or the *search path node pair* proposed by Hummon and Doreian (1989) to identify the major paths and the critical bottlenecks characterizing the evolution of the system.

In the case of Ethernet, our approach was capable of producing fruitful insights. The main trajectories we have identified using patent citation networks displayed a coherent economic and engineering logic, consistent with more qualitative accounts of the evolution of the Ethernet standard. Indeed, our evidence has revealed that inventive activity occurred in distinct phases. During the first phase, the main focus was on the definition of the standard and inventors were mainly concerned with changing and refining the original contention mechanism proposed by Metcalfe and Boggs. This ‘exploration’ phase ended with the official standardization of Ethernet that consolidated the specifications of the technology. The second phase was concerned with the application of Ethernet in use and the design of the components of the overall technical system (i.e. transceivers first, then hubs and switches). The third phase dealt with further improving the functioning of the individual components.

Further, the approach sheds light on critical patents which have punctuated the development of this technology. In the literature, the assessment of the economic and technological significance of individual innovations is based on the weighting of individual patents by means of received citations. The approach based on the structure of connectivity analysis opens another interesting perspective. As noted, some of the key patents calculated using citations as weights do appear in our mapped trajectories. This suggests that the approach used in this paper could fruitfully integrate the simple counting of citations. However, the study of the structure of connectivity of patent citation networks can also provide additional insights, by pointing out patents with low number of citations that occupy strategic positions in the unfolding of a technological trajectory.<sup>25</sup>

Although the method is of relatively straightforward implementation we would like to suggest some word of caution. First, it is necessary to adopt particular care in the identification of the relevant patent classes that constitute the object of analysis. Our position is that given the uncertainty surrounding the criteria adopted by patent offices for assigning a patent to a specific class, when performing this type of exercise, it is important to try out different possible selections of the patent universe considered, and devote some effort to investigate the patterns of citations existing between the various patent classes. Finally, as already noted by Verspagen (2007), the tools for the analysis of the structure of connectivity should not be employed ‘blindly’ but in combination with a detailed study of the engineering trends in the field in question.

## Acknowledgements

We thank the editor, two anonymous referees, Jackie Krafft, Arianna Martinelli, Koen Frenken, and Onder Nomaler for their insightful comments on previous drafts of this paper. We are grateful to Roberto I. Ruggiero for providing excellent research assistance. The research of Alessandro Nuvolari

has been funded by the Netherlands Organization of Scientific Research (Veni grant, 'Inventive activities, patents and the industrial revolution'). Roberto Fontana acknowledges support from the Italian Ministry for Education, Universities and Research (FIRB, Project RISC - RBNE039XKA).

## Notes

1. The original study of Hummon and Doreian (1989) deals with scientific research (they reconstruct the main path in the network of citations between scientific papers leading to the discovery of DNA).
2. The implicit assumption is that the significance of a citation link is adequately captured by its SPLC or SPNP weight.
3. The first data communication network was built in 1958 by the US Air Force as part of the semi-automated ground environment (SAGE) air-defence system. The first civilian application of a data network built around the principles of SAGE was realised in 1964 with the development of the semi-automatic business related environment (SABRE) airline reservation system by American Airlines. One characteristic of both the SAGE and SABRE communication systems was that the terminals were connected to the central mainframe through telephone lines. Given the characteristics of the medium, these networks operated at lower speed than the later LANs. Both their low speed and the fact that these networks were nationwide made these early application wide area networks (WANs) rather than local networks (Hellige 1994; von Burg 2001: 56).
4. These developments were among the major achievements of the ARPANET project. ARPANET was a computer network, built up at the end of the 1960s, to connect research centres engaged with the Defence Advanced Research Project Agency (DARPA) funded by the US Department of Defence. Originally only those centres working on military related fields were part of the network. Later on, other centres, government agencies and eventually countries joined it, opening the way to what later on would constitute the Internet (O'Neill 1995; von Burg 2001: 58)
5. The layer system was subsequently adopted by the International Standards Organization (ISO) which developed a reference model, the Open System Interconnection (OSI) model, based on seven functional layers: physical layer, data link layer, network layer, transport layer, session layer, presentation layer and application layer. LAN technologies were mainly concerned with the first and the second layers, and therefore making data transmission completely independent from the definition of higher level protocols.
6. According to CSMA/CD, all stations on the network could achieve access to network communication capacity without specifying a priority or hierarchy to any of them. Stations requesting access to the network channel had to be able to check that the channel (or 'carrier') was idle before transmitting and defer transmission in the case a transmission signal was detected. However, when two stations found the channel idle and started to transmit simultaneously frames might collide on the channel. This access method was able to register collisions, notify the stations that this had occurred and reschedule the transmission very rapidly so that, following a collision, stations could choose a random delay in retransmission.
7. The first Ethernet network connected hundreds of nodes and had a speed of 2.94 Mbps which was useful for the PCs and the laser printers installed at the Xerox PARC (Metcalfe 1994). However, it far exceeded the needs of most of LAN users at that time. Up until the mid 1980s, the memory of most of available PCs was less than 2 Mbps. For these computers, simpler technologies seemed to be more appropriate and cost-effective.
8. Indeed, the position of DEC, Intel and Xerox with respect to the IEEE initiative ranged from opposition to support conditional on the acceptance of their proposal as the main basis for the standard. Before official standardization occurred, the DIX alliance published their own proprietary set of Ethernet specification (known as the 'bluebook' v.1 and v.2). Its position, combined with pressures from other manufacturers sponsoring solutions based on different technologies such as Token Ring, contributed to delay the standardization process and ended up undermining the efforts of creating a unique LAN standard (Sirbu and Hughes 1986).
9. This version became known as 'thin Ethernet' (10Base-2) as opposed to 'thick Ethernet' (10Base-5) which was the original Ethernet standard.
10. In the 1970s, IBM were using the ring topology to link controllers to looped workstations in the IBM 3600 on-line computer system mainly adopted by banks (Abrahams 1991, 2).
11. In 1989, a user survey by the journal Data Communication revealed that, although coaxial cables were more expensive than copper based ones, 90.4% of LAN users employed coaxial cables for

- their networks and 63.5% had a network which spanned between buildings. Moreover, for most of the users LAN functionality was the most important factor behind their LAN purchase. Price was decisive only for 6% of them (Data Communications, June 1989 Issue).
12. The number of manufacturers supporting Ethernet rose from 34 in 1985 to 99 in 1987 and topped 131 in 1989. In 1985, only five manufacturers were supplying Token Ring products in 1985 although their number rose to 53 and 77 in 1987 and 1989, respectively (Fontana 1999).
  13. A fourth standard, Fiber Distributed Data Interface (FDDI) was also developed (Lehr 1996).
  14. As high speed upgrades for Ethernet, fast Ethernet and 100VG-AnyLAN were both proposed to the IEEE in November 1992 by 3Com Corp. and Hewlett Packard, respectively. The 3Com proposal, backed by SynOptics, Sun Microsystems and LAN Media Corp., was modelled on the recent work of a start-up company, Grand Junction Networks, which studied the possibility of upgrading the ET speed to 100 Mbps while maintaining the CSMA/CD access method adopted by Metcalfe's Ethernet. The solution proposed by Hewlett Packard and supported by AT&T, named 100Base-VoiceGrade (100Base-VG), adopted a so-called 'demand priority scheme' as the access method which would enable it to handle two different kinds of traffic by prioritising the access to the network (Rauch and Lawrence 1995). The adoption of this access method would turn out to be the major shortcoming of this standard relative to Fast Ethernet.
  15. The roots of the development of ATM extended back to early researches in cell switching conducted at the Bell Laboratories in the late 1960s when the system of cell relay, which would constitute the basis of ATM, had been conceived under the name of Asynchronous Time Division Multiplexing (ATDM). Some work on its principles had continued during the 1970s. However, ATDM was never commercialized (Gould 1994, 122).
  16. Data Communications (May 21, 1994: 16).
  17. Patents have been retrieved by means of queries in the international patent class (ICL) field.
  18. A component is a set of patents which are all connected. For a more detailed discussion, see Wasserman and Faust (1994, 109–110).
  19. Similar results were obtained using SPLC as measure of the connectivity of the links.
  20. A baseband network, such as Ethernet, is one that provides a single channel for communications across the physical medium (e.g., cable), so only one device can transmit at a time. With broadband, the physical cabling is virtually divided into several different channels, each with its own unique frequency. These different frequencies are transferred onto the network cabling in such a way to allow multiple simultaneous transmissions to occur (see <http://www.ethermanage.com/ethernet/enet-faqs/ethernet-faq.html> (accessed January 22, 2008)).
  21. We thank an anonymous referee for raising this point and stimulating us to explore more thoroughly this issue.
  22. To probe further into the role of patent 4751701, we have removed this patent and its citation links from the sample and calculated again the network of main path. It is interesting to note that, in this counterfactual world, the top path of the network includes almost the same sequence of patents as found in Figure 5. In other words, in the counterfactual world, we see a similar dynamics of convergence between the broadband and baseband environments. However, this converge occurs much later in time, with patent 5568469 (filed in 1994), instead than with patent 4751701 (filed in 1985). Thus, in our judgment, it is not far-fetched to ascribe to patent 4751701 a significant role in bridging the two environments *at that point in time*.
  23. Within information communication systems there are several methods for detecting and avoiding collision. One method consists in detecting when a collision occurs by measuring phase shift in the information signal received. This is the method embodied in the original 1977 Ethernet patent by Metcalfe and Boggs. Patent no. 4271523 represents an improvement upon this patent in the sense of reducing the susceptibility of phase shift to noise in the medium. Patent nos. 4345250 and 4376278 instead, both owned by Honeywell, aim at reducing collision of signals by charging the medium with a predetermined voltage. Another method consists of using a time slot to check when transmission may occur. Other methods entail the use of a controller to ask devices permission to transmit or the division of the channel into different frequencies with different bandwidth (Hopkins 1979).
  24. Latency is the delay existing from the source sending a packet to the destination receiving it. Jitter instead is an abrupt and unwanted variation of one or more signal characteristics (i.e. the interval between successive pulses), the amplitude of successive cycles, or the frequency or phase of successive cycles of transmission.

25. In this respect, of course, it is necessary to take into account the qualifications made above in Section 5.

## References

- Abrahams, J. 1991. *Token ring networks*. Oxford: NCC Blackwell Limited.
- Antonelli, C. 1992. The Economic theory of information networks. In *The economics of information networks*, ed. C. Antonelli. New York: Elsevier Science Publishers.
- Antonelli, C. 1993. The dynamics of technological interrelatedness. The case of information and communication technologies. In *Technology and the Wealth of Nations*, ed. D. Foray and C. Freeman, London: Francis Pinter.
- Bresnahan, T.F., and A. Chopra. 1990. The development of the Local Area Network market as determined by user needs. *Economics of Innovation and New Technology* 1: 97–110.
- Bunn, J.A., and P.A. David. 1988. The economics of gateway technologies and the network evolution: Lessons from electricity supply history. *Information Economics and Policy* 3, no. 2: 165–202.
- von Burg, U. 2001. *The triumph of Ethernet*. Stanford: Stanford University Press.
- Christensen, K.J., L.C. Haas, F.E. Noel, and N.C. Strole. 1995. Local Area Networks. Evolving from shared to switched access. *IBM Systems Journal* 34: 347–74.
- David, P.A., and S. Greenstein. 1990. The economics of compatibility standards: an introduction to recent research. *Economics of Innovation and New Technology* 1: 3–41.
- Dosi, G. 1982. Technological paradigms and technological trajectories. *Research Policy* 11, no. 3: 147–62.
- Fontana, R. (1999), Twinning networks: co-evolution and competition of component system technologies in the Local Area Network industry. SPRU Electronic Working Paper Series 31, SPRU – University of Sussex, UK.
- Fontana, R., and L. Nesta. 2006. Product entry in a fast growing industry: The LAN switch market. *Journal of Evolutionary Economics* 16, nos. 1–2: 45–64.
- Gould, J. 1994. ATM's long, strange trip to the mainstream. *Data Communications* 23: 120–30.
- Grebel, T., J. Krafft, and P.P. Saviotti. 2006. On the life cycle of knowledge intensive sectors. *Revue de L'OFCE, Special Issue* June: 63–85.
- Hellige, H. 1994. From SAGE via ARPANET to Ethernet. Stages in computer communications concepts between 1950 and 1980. *History and Technology* 11: 49–75.
- Herman, J. 1991. Smart LAN hubs take control. *Data Communications* 20: 62–73.
- Hopkins, G.T. 1979. Multimode communications on the MITRENET. Paper presented at the LACN Symposium, May, in Boston, MA.
- Hughes, T.P. 1987. The evolution of large technological systems. In *The social construction of technological systems*, ed. W.E. Bijker, T.P. Hughes, and T.J. Pinch, 51–82. Cambridge, MA: MIT Press.
- Hummon, N.P., and P. Doreain, 1989. Connectivity in a citation network: The development of DNA theory. *Social Networks* 11: 39–63.
- Jaffe, A., and M. Trajtenberg. 2002. *Patents, citations and innovations*. Cambridge: MIT Press.
- Katz, M.L., and C. Shapiro. 1994. System competition and network effects. *Journal of Economic Perspectives* 8, no. 2: 93–115.
- Lehr, W. 1996. Compatibility standards and industry competition: Two case studies. *Economics of Innovation and New Technology* 4: 97–112.
- Melatti, L. 1994. Fast Ethernet: 100 Mbit/s made easy. *Data Communications* 23: 111–16.
- Metcalfe, R. 1994. How Ethernet was invented. *IEEE Annals of the History of Computing* 16, no. 4: 81–8.
- Mina, A., R. Ramlogan, G. Tampubolon, and J. S. Metcalfe. 2007. Mapping evolutionary trajectories: Applications to the growth and transformation of medical knowledge. *Research Policy* 36, no. 5: 789–806.
- Rauch, P., and S. Lawrence. 1995. 100VG-AnyLAN: The other Fast Ethernet. *Data Communications* 24: 129–134.
- Rosenberg, N. 1976. *Perspectives on technology*. Cambridge: Cambridge University Press.
- Sirbu, M., and K. Hughes. 1986. Standardisation of Local Area Networks. Paper presented at the 14th Annual Telecommunications Policy Research Conference, April, in Virginia.
- Spurgeon, C.E. 2000. *Ethernet: The definitive guide*. Sebastopol: O'Reilly and Associates.

- Verspagen, B. 2007. Mapping technological trajectories as patent citations networks. A study on the history of fuel cell research. *Advances in Complex Systems* 10: 93–115.
- von Wartburg, I., T. Teichert, and K. Rost. 2005. Inventive progress measured by multi-stage patent analysis. *Research Policy* 34: 1591–607.
- Wasserman, S. and K. Faust. 1994. *Social network analysis. Methods and application*. Cambridge: Cambridge University Press.