

INTRODUCTION BY EDITOR

The following article, by Rémi Després, is the second on the history of X.25 systems to appear in this column. As noted by Dr. Després, the previous article focused on the Canadian Datapac system. Earlier articles on packet switching in this column have included one on the history of the Arpanet/Internet and one on early British packet switching systems. What makes this article particularly distinctive, aside, of course, from the fact that it focuses on the major contributions of French engineers to the development of packet switching as well as to X.25 standardization, is that it carefully outlines the reasons for the choice of connection-oriented virtual circuits for the Transpac network, as contrasted with datagram-based packet switching adopted for Arpanet. Interestingly, Dr. Després notes that the idea of using virtual-circuit connection-oriented packet switching in the Transpac development came from the British packet switching

activity. It is to be noted that early commercial packet switching networks in the United States, such as Tymnet and Telenet, also adopted the virtual circuit paradigm.

In a world in which competition is usually manifest throughout, it is heartening to read that engineers from a number of countries worked closely together in both developing the basic concepts of packet switching and undertaking the difficult task of harmonizing the X.25 standard. As the author points out, X.25 has had its day, to be supplanted throughout the world by the Internet. Yet some of the concepts developed for X.25 systems live on in the Internet, while engineers trained on the X.25 paradigm have been readily able to make the transition to the Internet, carrying their knowledge of packet switching with them. I commend this most interesting article to your reading and edification.

—Mischa Schwartz

X.25 VIRTUAL CIRCUITS — TRANSPAC IN FRANCE — PRE-INTERNET DATA NETWORKING

RÉMI DESPRÉS

INTRODUCTION

Two previous articles in this series, by Peter Kirstein and Tony Rybczynski, covered the early history of packet switching in the United Kingdom and Canada [1, 2]. This one is about the early history of packet switching in France. It presents the steps that led the public Transpac network, based on standards of this period, to become the largest of its generation.

Tony Rybczynski having already described in his article the contents of X.25, the dominating standard of the '80s, and the major official steps that led to its approval in 1976 by Comité Consultatif International Téléphonique et Télégraphique (CCITT), we concentrate here, concerning standardization, on the rationale behind choices made and some turning points that influenced the outcome.

HERMES PROJECT AND RCP
EXPERIMENTAL NETWORK

Studies on packet switching started in the French Telecommunication Administration when Jacques Dondoux, then directing the Centre National d'Étude des Télécommunications (CNET), launched Project Hermès in 1971. The project objective was to specify a specialized network architecture for data communications, and to do it in relationship with the international work of CCITT in its Nouveau Réseau pour Données (NRD) group.

The British Post Office (BPO), which was very active in the NRD

group, was promoting a circuit switching network with fast circuit establishment, but was also suggesting to study a new concept, the "packet mode of operation," or "packet switching." In view of the experience of my team on computer software (with Alain Bache and a few colleagues, we had developed a time-sharing system for an in-house general-purpose computer), I was proposed to investigate this subject, perceived as esoteric in the telecommunication world. The goal was to determine whether it could be useful for a public service. The most influential publications in this domain were then some from Donald Davies and Larry Roberts. Donald Davies, leading a team at the National Physical Laboratory (NPL, where Newton used to think below apple trees), had invented the packet switching concept for shared data networks [3]. Paul Baran, at the Rand Corporation, had before that worked on related ideas for the U.S. Department of Defense, but from a more speculative and futuristic perspective [4]. Larry Roberts and his colleagues had generously documented their approach for Arpanet, the U.S. network for resource sharing among academic computer centers they were developing for the Advanced Research Project Agency (ARPA) [5]. Readily convinced of the great potential of packet switching for data traffic, I enthusiastically took the job. After a theoretical study, I proposed to validate the practicability of a packet-switching-based public service on a small size network called Réseau à Commutation par Paquets (RCP).

Alain Profit, then manager of the Hermès project, agreed to finance RCP. One year later, he also agreed that the project team be moved to Centre Commun d'Études de Télévision et Télécommunications (CCETT), a newly created research center where hiring more engineers would be significantly easier than at CNET.

RCP's configuration, as planned in 1972, is shown in Fig. 1. Packet mode customer devices, typically computers and protocol converters, had access to the network via point-to-point synchronous leased lines (SLLs). On these access lines, they could support interleaved data communications with multiple other packet mode devices and with multiple character mode devices. Character mode customer devices were at that time teletypewriters and simple keyboard display terminals. They had access to RCP via the switched telephone network (STP), the switched telex network (STX), or point-to-point asynchronous leased lines (ALLs). Each one could establish a data connection with a packet mode computer, a computer supporting multiple character mode interfaces to RCP, or another character mode terminal. The three switching nodes were standard PDP-11 minicomputers from Digital Equipment. Character mode multiplexers were standard products of Société Anonyme de Télécommunications (SAT), which also developed for RCP PDP11 adapters to remotely control them.

RCP served as a testbed for the virtual circuit (VC) model. With VCs, the

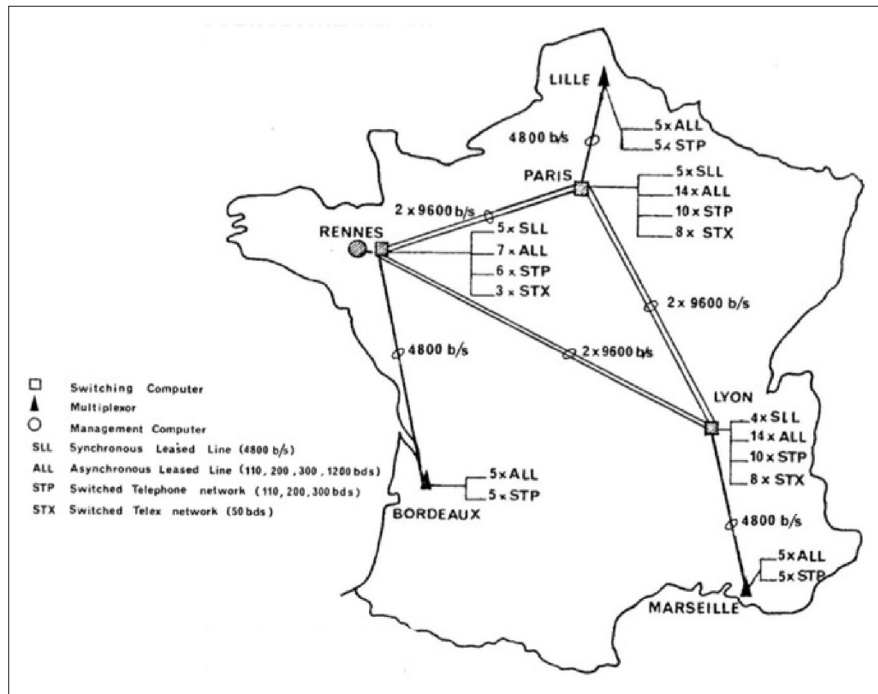


Figure 1. The RCP experimental network in 1975.

larly with character mode terminals, it was important that packet headers be kept small. As we wanted a flexible address format, capable of supporting a virtually unlimited number of customers, it would have been a great waste to repeat full addresses in every packet. With a connection-based service, once a VC is established across a transmission link, packet destinations can be implicitly coded in short labels that identify VCs to which packets belong. Thank to this, X.25 supported 60-bit-long addresses with data packet headers, including their error control and flow control fields, having 32 bits. For comparison, the current Internet supports 32-bit-long addresses with TCP/IP packet headers having 320 bits. This would not have been economically competitive at that time.

A major departure of RCP protocols from those of EPSS was the introduction of a simple and reliable link layer protocol. The EPSS link layer, between a customer device and its network packet switch, could duplicate some packets, could not sustain continuous transmission at full speed, and necessitated sophisticated specialized hardware. The RCP link layer protocol made duplications impossible, permitted continuous transmission at full speed, and was implementable in software with existing hardware. Based on an improved version of an error correction mechanism invented at the NPL, it was very simple. It was not retained for X.25 for a reason explained below, but became in 1981 part of Signaling System No. 7 for interexchange signaling in telephone networks (CCITT Recommendation Q.703).

With data integrity ensured at the link layer, specifying a protocol for reliable end-to-end service becomes easy. For flow control to be independently exercised for each VC that traverses an access link, each end of the link informs the other end, on a per VC basis, when it is ready to accept more data. For this, a classical sliding window mechanism can be used for each VC, simplified by the fact that no packet is lost at the link layer. If the network has an irrecoverable failure, established VCs are cleared (or just re-initialized if they are PVCs). Since the effect of such a failure on applications is the same as that of a temporary physical access or customer equipment failure, it must be acceptable provided network failures are made rare enough.

For a switching node to announce that it is ready to accept more packets on a VC, it must have enough memory

network is aware of connections established between packet mode devices. RCP opened service in 1974. It confirmed that a public data communication service could be offered with switching nodes based on available computer technologies, and with easily understood and implemented protocols [6]. It also proved that computer manufacturers could rather easily adapt their software to support and use VC protocols like those of RCP. IBM was first to do it, at its La Gaude research center in France, followed by Honeywell-Bull and Compagnie Internationale pour l'Informatique (CII).

RATIONALE FOR THE VC PARADIGM

During the first phase of our work on packet switching, in 1971, we had a private presentation of the Experimental Packet Switching Service (EPSS), a project of the BPO [7]. It introduced, for packet mode devices, the concept of "virtual calls" and "permanent virtual circuits." Although detailed proposed protocols were in our understanding far too complex and had severe limitations, the idea of combining packet transmission and connection-based service immediately had great appeal to us. We endeavored then to simplify and complete the concept, and to validate it on RCP. We adopted the generic term "virtual circuit" to cover both virtual calls, renamed switched VCs (SVCs),

and permanent virtual circuits (PVCs).

On one hand, packet transmission was attractive for data traffic because of its potential for multiplexing on transmission links traffic mixes having widely different characteristics. At that time, data transmission rates ranged from 50 b/s to 48 kb/s, and silence ratios on established connections were also largely variable. Another key feature of packet transmission was that it made communication between customer access links having different data rates much easier than with circuit switching. The counterpart was that some flow control would have to be exercised on high-speed sources when they transmit toward low-speed destinations; but, at least with VCs, simple solutions could be found for this.

On the other hand, connection-based services were attractive for data traffic because they allowed the specification of the quality of service (QoS) of each connection and the enforcement of differentiated committed data rates on heavily loaded shared circuits [8]. They also permitted great savings on link utilization. At that time, both these properties were important because the cost of customer access links was high, with a strong dependence on supported data rates, and international links were extremely expensive. Since the average number of octets to be sent per packet was very small in the late '70s, particu-




Figure 2. CP50 X.25 switches in a Transpac center.

to store them when they arrive. This proved to be easy to ensure at reasonable cost with available technologies, for both RCP and the full size Transpac later on.

The complete VC protocol of customer links, symmetrical at both the link and VC layers, can also be used to interconnect two operator networks. It can even be used internally between nodes of each independent network. The end result is great simplicity of the overall model. By comparison, the TCP of the Internet, having to perform error and flow control end to end over an underlying “best effort” infrastructure, is highly sophisticated. It took years to complete it, particularly with a major revision in 1988 to prevent network instability that had been observed [9].

GOING MAINSTREAM WITH TRANSPAC

In October 1973 Louis-Joseph Libois, then Directeur Général de Télécommunications, publicly launched a study to determine how a packet switching public data service might open as early as 1976. He entrusted CCETT with the task of its technical specification. This decision, first of its kind in the world, was in part influenced by external pressure: several powerful public and private organizations had formed Groupe d’Etude pour un Réseau Commuté Interprofessionnel de Paquets (GERCIP) and declared their intention to build a common packet switching network; independently, the Ministry of Industry had decided to fund its own packet switching network, Cigale, a part of the larger computer communication project Cyclades of Louis Pouzin [10].

Responsibility to supervise the work of CCETT, manage contacts with computer manufacturers, and coordinate with Cyclades was assigned to Alain Profit. Philippe Picard was given responsibility for economic and early marketing studies.

At the end of 1974, a detailed specification was available for a public network, in the meantime officially named Transpac. A draft, written by Yves Schwartz, Guy Pichon, and myself, had previously been submitted for external reactions to the BPO (then working on its EPSS), IRIA (then working on Cyclades), and Groupe pour l’Étude du Raccordement à Transpac (GERPAC, a new avatar of GERCIP after it decided not to build its own network). Upon request of the Ministry of Industry, the draft had included, besides its detailed VC specification, a datagram service specification (DG) derived from that of Cigale. This precaution having been taken, only minor comments were received, and the draft was finally approved. The DG specification was so imprecise that potential contractors would have had to complete it in their own way, but when international agreements on VCs had progressed, the request for DG service was abandoned.

Before the end of 1974, Philippe Picard had convinced the newly appointed Directeur Général des Télécommunications, Gérard Théry, that the Transpac project was ready to be launched. The necessary green light from the government was then given with three conditions: Transpac must be operated by a separate company; user representatives must have shares in this

company; and specifications of the network must be approved by the Ministry of Industry. These conditions being accepted, the call for tenders was issued in February 1975.

The winning proposal was that of the consortium led by Société d’Etude des Systèmes d’Automation (SESA). Managed by Jacques Stern and Jacques Arnould, SESA had already acquired some packet switching know-how as co-contractor for the European Informatics Network (EIN). EIN was an experimental network financed by the European Economic Community and technically derived from Cigale. High-capacity and redundant packet switches proposed by SESA, the CP50s shown in Fig. 2, had been designed by TIT, a company that sold message switching computers to the French Navy. CP50s were to be manufactured by TRT, a dynamic telecommunication equipment company already selling a large range of modems. Control units, which handled VC establishments, were Mitra-15 minicomputers of CII. Leveraging its Transpac experience, SESA later commercialized X.25 products for Euronet (a pan-European network funded by the European Economic Commission [EEC] in order to boost the scientific database market in Europe), and later for various national networks including some in Australia, Brazil, New Zealand, and China.

The contract for the first configuration, supporting up to 1500 packet mode customers, was signed in April 1976 [11]. In the meantime, the initial Transpac VC specification had been replaced by that of X.25, without cost or delay implications as the standard was very close to that originally specified. Later in 1976, when the CCITT agreement on X.3/X.28/X.29 for character mode support was finalized, its specification replaced the original one, but this time with negotiated contractual implications.

While the contract was followed through, a dedicated Transpac project team had been set up. Philippe Picard was at its head, with full responsibility on economic, marketing, technical, and operational aspects of the project.

For future customers to decide to use Transpac in advance, tariffs had been discussed with GERPAC and announced two years before service opened, with an uncertainty officially limited to 10 percent. According to an innovative decision made by Gérard Théry, tariffs had to be independent of distance in all their constituents. When Transpac opened in December 1978,

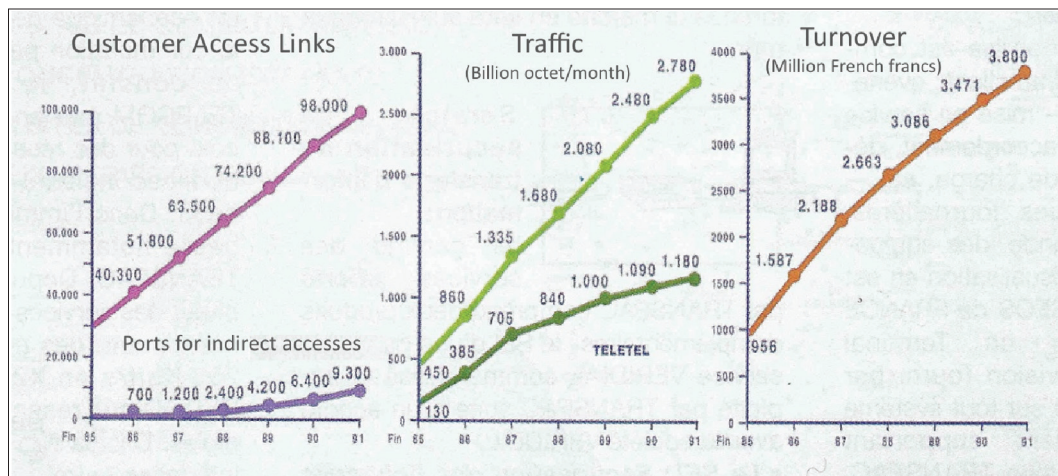


Figure 3. Transpac growth from 1985 to 1991.

the initial tariff was structured as follows (1 French franc being roughly worth US\$ 0.20). At peak times, the volume-based charge was 0.06 F/k-octet, and SVC duration charges were from 0.01 F/min, for a 1.2 kb/s committed data rate, to 0.20 F/min, for a 48 kb/s committed data rate. Important discounts were applicable to these charges at off-peak times (–80 percent during weekends and weekdays from 0:00 to 6:00 a.m.; –40 percent during weekdays from 6:00 a.m. to 8:00 a.m. and from 7:00 p.m. to midnight). Dedicated access links had monthly charges ranging from 330 F/month for character mode access at 300 baud, to 5000 F/month for packet mode accesses at 48 kb/s. PVCs had monthly charges ranging from 108 F/month for a 1.2 kb/s committed data rate, to 2160 F/month for a 48 kb/s committed data rate. These tariffs, which progressively decreased as the network grew, proved their adequacy: customer subscriptions consistently exceeded initial expectations; the financial break-even point was reached earlier than expected.

In the '70s, any device connected to a telecommunication network (e.g., a modem on a leased line) had to be certified to check that it would not disrupt network operation. But with a proper implementation of X.25 VC protocols in switching nodes, no connected device could endanger normal operation. After some hesitation, it was decided that the burden of certification, which would have been detrimental to a fast take-off of the service, could be dispensed with. Instead, a detailed, rigorous, and comprehensive description of the planned Transpac behavior was documented in the Specifications Techniques d'Utilisation du Réseau (STURS) so that manu-

facturers planning to have products connectable to Transpac could be ready in due time. In addition, a CCETT team led by Paul Guinaudeau implemented, in one year on a Mitra-15 minicomputer, a switching node that simulated Transpac's behavior, REX25. Manufacturers were allocated test sessions on REX25 to validate their implementations before Transpac could be available.

Service acceptance was encouraging from the beginning [12]. In 1980, with 2395 operational X.25 access links, banks counted for 28 percent of established VCs, service bureaus for 19 percent, industry for 15 percent, and the public sector for 14 percent. One year later, more than 5000 X.25 accesses were operational (at a time when the Internet was still in its infancy, with its first 213 hosts). Transpac's continuous growth during the 1985–1991 period is shown in Fig. 3, and its configuration in 1991 is shown in Fig. 4.

QoS, a key criterion for VC acceptance, was carefully checked, with periodical reports to UTIPAC, the Transpac user association that had replaced the GERPAC. During initial traffic buildup, a few bugs had to be eliminated, and the QoS stabilized to a satisfactory level during the first quarter of operation. Four years later, when real traffic had exceeded that which traffic generators used for acceptance tests had been able to generate, the network from time to time unduly cleared some VCs. A flaw in the design of the CP50 software was quickly identified, and corrected once and for all.

Three years later, in 1985, more than a million character-mode terminals, the Minitels, were in operation. They were made available to telephone customers to consult the national telephone direc-

tory, and were also used for other applications, including a Teletel service. The resulting increase of VC establishments per second revealed a dormant software bug that caused a serious degradation of the service. After two weeks of service interruption for Minitels, which caused much discontent, the problem was diagnosed and solved. After that, Transpac's quality of service remained satisfactory and was generally praised by its users.

THE X.25 STANDARDIZATION SAGA: VCS AND DGs

In 1972, both the CCITT and the Commission Européenne des Postes et Télécommunications (CEPT) had appointed Rapporteurs on the packet mode of operation, Halvor Bothner-By of Norway for CCITT and me for CEPT, each participating in meetings organized by the other. Having different views on how a packet-switched public service could best be offered, both Rapporteur groups decided to work in parallel on DGs and VCs. The term "datagram," so successful later on, was coined by Halvor Bothner-By and a colleague on a train between Paris and Rennes, taken to attend a CEPT Rapporteur meeting. The DG model was defined as one in which standard format packets are transmitted across a network independent of each other, and for which flow control within the network relies on a friendly cooperation of end-user devices. The network discards packets when and where its internal queues tend to grow too much; end-user sources are expected to refrain from transmitting too many packets toward destinations in directions of which packets tend to be discarded. Note that this

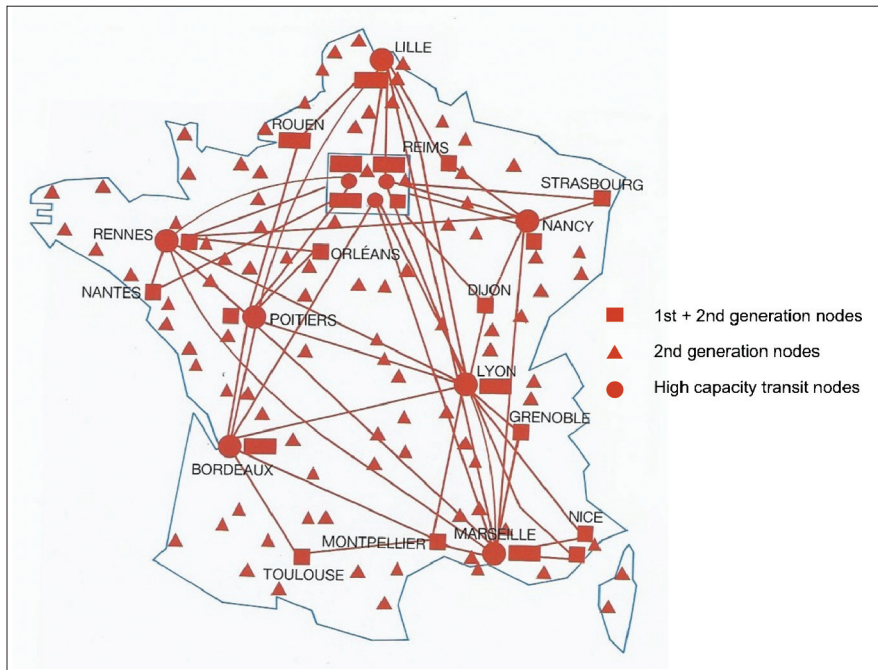


Figure 4. Transpac map in 1991.

when Philip Kelly of the BPO agreed with Philippe Picard that the Transpac technology should also be adopted for Euronet. After that, the British stand in CCITT, so far mildly in favor of datagrams by reference to the EIN project, abruptly switched to dedicated and active support for the VC multilateral agreement. Philip Kelly, highly experienced in CCITT practices, quickly helped with his colleagues to structure an appropriate set of contributions. He also involved Japan, where Dr Masao Kato of NTT also had plans for a packet-switched service.

At the CCITT Study Group VII meeting in February–March 1976, a formal contribution was submitted jointly by France and the United Kingdom, the two parties involved in the agreement that had voting status in CCITT. It contained a complete X.25 draft based on the multilateral specification [14]. Many objections were expressed by delegations that had not participated and were surprised by such rapid progress. I was then appointed as editor to try and resolve these issues during the weekend. After a full Saturday and a full Sunday of intense meetings, all points raised had been resolved among participants. By Monday morning, Tony Rybczynski and Paul Guinaudeau, who had spent all night rewriting clean handwritten versions of the modified specification, had them ready. Thus, both English and French versions were available, as required to forward a proposal to the CCITT plenary. All delegates had the photocopied documents necessary for a formal vote, and unanimously approved them. At the CCITT plenary itself, in June 1976, X.25 was unanimously endorsed, with two subjects left for future studies. One, asked for by IBM, was that a lighter variant be designed for the simplest terminals, a wish that was later found unnecessary. The other was that the specification of a DG service should be added to X.25. This addition did take place at the end of the next CCITT plenary, in 1980, but, no implementation being planned, it was deleted in 1984.

Complementary Recommendations X.3/X.28/X.29 for packet assemblers-disassemblers (PADs), necessary to support character mode terminals, were finalized in 1977, with Bernard Jamet and Chris Broomfield as the main editors. Recommendation X.75, the variant of X.25 adapted to links between X.25 networks, was finalized in 1978. The scene was thus completely set for a worldwide packet-switched service to be extensively deployed in the 1980s.

definition of DGs differs significantly from that of the Internet specified in 1981: in the Internet Protocol, each DG can be transmitted as a series of packets that share a common DG identification; the network may further fragment each of these packets into several smaller ones; final destinations are responsible for reassembling all fragments [13].

A major turning point of international discussions took place at a meeting organized by Dave Horton, head of the Canadian Datapac project, and Philippe Picard, with the participation of Tony Rybczynski and myself. Both parties had two objectives in common: opening a service as soon as possible and obtaining an international standard. Since our two projects were the two major ones at that time, chances of reaching a standard would be good only if the two of us could agree. But there was a problem: the Datapac proposal, the Standard Network Access Protocol (SNAP), was DG-based; ours was VC-based. After a long discussion between Tony and myself, walking through Paris streets, and continuing in a bar late at night, the first sketch of what might be an agreement was drawn. The point-to-point link layer protocol would be that of SNAP. It was technically more complex than that specified for Transpac but was based on the emerging high-level data link control (HDLC) promoted by IBM in the International Organization for Standards (ISO). As such, it was a much better candidate for

an international agreement. Above this link layer, we would adopt the VC layer of Transpac. It was much simpler and more efficient on customer links than the SNAP proposal, which had two layers, one for DGs and one for error and flow control, where one was sufficient. Soon after, I went to Ottawa with Paul Guinaudeau to more deeply discuss pros and cons of the new combination. Dave Horton then gave his green light. After several meetings in Canada and France, Tony Rybczynski, Claude Martel, Paul Guinaudeau, and Bernard Jamet had assembled a detailed specification. Dave Horton and Philippe Picard then made a common commitment to implement it in our networks, and to amend it only after common approval.

The next good news came soon after, when I met in Washington with Larry Roberts, Arpanet's father, and Barry Wessler. Their startup, Telenet, was known to be preparing a commercial packet-switched network, but the chosen technical approach was unknown. Discovering that their choice was a VC model was a notable confirmation that we were on the right track. They had an HDLC link layer rather similar to ours. Above it, their VC layer was different in a number of details, but there was no essential difference that would preclude compatibility. With a few agreed minor complements to the current specification, Telenet joined the agreement.

The next important step forward was

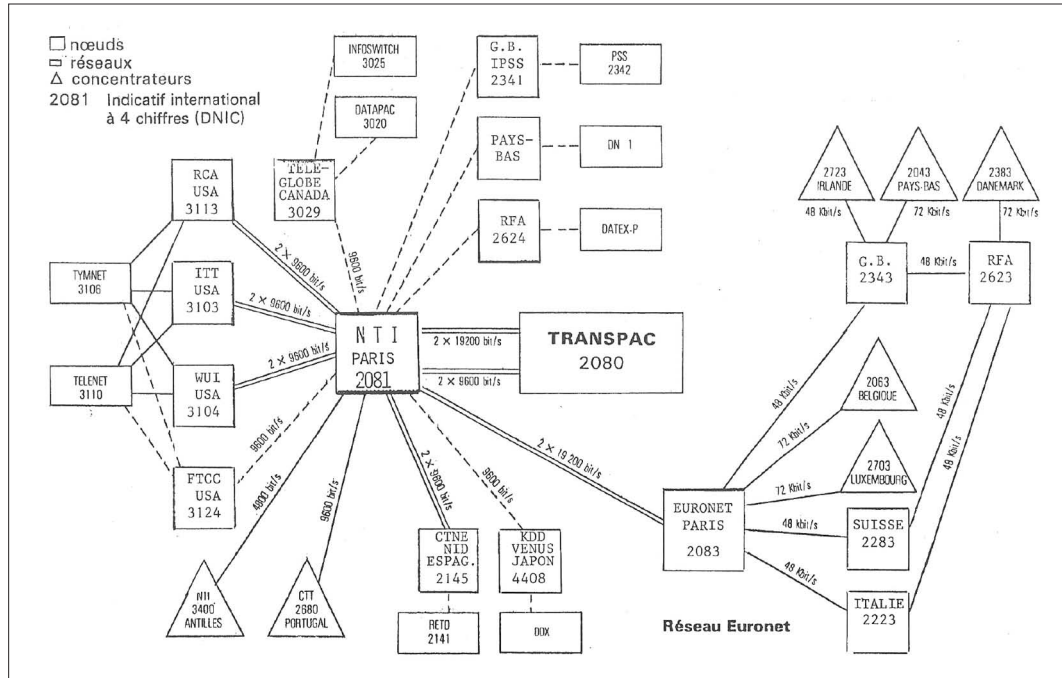


Figure 5. Transpac existing and planned international links in 1980.

INTERNATIONAL LINKS AND PRIVATE NETWORKS

For international connections of Transpac with other X.25 networks, the team that had developed REX25 implemented an international transit node (NTI).

Close cooperation then took place with U.S. international record carriers (ITT, WUI, and RCA) and with Tymnet, which implemented their equivalent of the NTI. The first transatlantic links were then opened in 1979, on pairs of 9.6 kb/s circuits. Typical international applications were then accesses to remote servers by character mode terminals and exchanges between bank message centers. Direct links followed with European countries operating their own X.25 networks (Germany, the United Kingdom, Spain, the Netherlands, Luxembourg), and with Canada and Japan (Fig. 5). The link with Euronet, operated by a consortium of telecom administrations, completed connectivity with the remainder of Europe.

Aside from Transpac, a number of private networks using X.25 started to appear for intrasite communications and also for short-distance private networks (for these, Transpac tariffs, being distance-independent, were not optimum). Products that were successful in this market included the Compac range of TRT, the Megapac range of Sagem, and the Ecom range of OST (a startup

later acquired by the Canadian Newbridge Networks). Tekelec-Airtronik developed and successfully commercialized worldwide an X.25 protocol analyzer, the TE92.

EPILOGUE

As everyone knows, the Internet has become the ubiquitous data network of the globalized planet. Its initial penetration in less developed countries used pre-existing packet switching infrastructures, with Internet IP packets transmitted on X.25 VCs, but this is now pure history.

The predominance of X.25 VCs was first shaken when frame relay PVCs were marketed as permitting data rates much higher than those of X.25. At a time when the Frame Relay Forum distributed leaflets explaining that X.25 could never exceed 64 kb/s, there were already switches supporting X.25 at 2 Mb/s, but the buzz prevailed. It was true, though, that X.25 had a limitation for links having very high data rates on links having long propagation delays. For this reason, X.45, a variant of X.25 that eliminated this limitation, was endorsed by CCITT in 1996. It was designed for transparent interworking between X.25 and X.45 customers, thus permitting incremental deployment. But it arrived much too late and was never commercially supported. The irresistible success of the Internet and its TCP/IP protocol suite was already too advanced.

For some time, asynchronous transfer mode (ATM) was presented as a panacea that could replace X.25 and frame relay, and even replace TCP/IP, but it did none of these. It was successful only as a flexible multiplexing technology for high-capacity and long-distance transmission trunks, and some asynchronous DSL (ADSL) customer access links.

The Transpac company, which had been created in 1978 to operate the Transpac network, was fully reintegrated into France Telecom in January 2006. It had in the meantime evolved to sell to its professional clients less and less X.25, more and more frame relay, and more and more TCP/IP. The number of Minitel terminals using Transpac had peaked at 6 million in 1993, and the number of X.25 customers had peaked at 105,000 in 1995. In January 2010 France Telecom announced that commercialization of X.25 would be discontinued after July 2010, and that X.25 services would be closed after November 2011.

From a purely technical point of view, applications that caused the exponential growth of the Internet, email, and the Web could have worked on a VC-based worldwide infrastructure, but other considerations did not make it possible. Some of the QoS features of connection-oriented services might some day reappear in the Internet, but this does not seem to be a priority today.

ACKNOWLEDGMENTS

Gratitude is due to all those who contributed to this captivating adventure, in particular for the great cooperative spirit that made it possible: Alain Bache, Alain Profit, Barry Wessler, Bernard Jamet, Chris Broomfield, Claude Martel, Dave Horton, Guy Pichon, Jacques Stern, Jacques Arnould, John Wedlake, Larry Roberts, Masao Kato, Paul Guinaudeau, Philip Kelly, Philippe Picard, Pierre-Yves Schwartz, Tony Rybczynski, and many others who cannot be listed here. Special thanks are, in addition, due to Guy Pichon, Jean-Michel Simon, Paul Guinaudeau, and Philippe Picard for their help in retrieving some material used for this article. Tony Rybczynski's detailed review of an early draft of this article was extremely helpful.

REFERENCES

- [1] P. Kirstein, "The Early History of Packet Switching in the UK," *IEEE Commun. Mag.*, Feb. 2009, pp. 18–26.
- [2] T. Rybczynski, "Commercialization of Packet Switching (1975–1985): A Canadian Perspective," *IEEE Commun. Mag.*, Dec. 2009, pp. 26–31.
- [3] D. W. Davies, "Packet Switching in a Public Data Network," IFIP-1971 Congress Proceedings, pp. TA4 69–73.
- [4] P. Baran, "On Distributed Communication Networks," *IEEE Trans. Commun. Sys.*, Mar. 1964, pp. 1–9.
- [5] L. G. Roberts and B. D. Wessler, "Computer Networks to Achieve Resource Sharing," SJCC-1970 Proc., pp. 543–49.
- [6] R. Després, "RCP, The Experimental Packet-Switched Data Transmission Service of the French PTT," *ICCC-1974 Proc.*, pp. 171–85.
- [7] R. D. Bright and M. A. Smith, "Experimental Packet Switching Project of the UK Post Office," NATO Advanced Study Inst. on Comp. Commun. Networks, Univ. of Sussex, 1973; *Noordhoff Proc. 1975*.
- [8] R. Després, "A Packet Switching Network with Graceful Saturated Operation," *ICCC-1972 Proc.*, pp. 345–51.
- [9] V. Jacobson, "Congestion Avoidance and Control," *ACM Sigcomm 1988 Proc.*, pp. 273–88.
- [10] L. Pouzin, "Presentation and Major Design Aspects of the Cyclades Computer Network," NATO Advanced Study Inst. on Comp. Commun. Networks, Univ. of Sussex, 1973; *Noordhoff Proc. 1975*.
- [11] A. Danet et al., "The French Public Packet Switching Service: The Transpac Network," *ICCC 1976 Proc.*, pp. 251–60.
- [12] P. Picard, "How Transpac Fits in the Data Processing Business," *Teleinformatics 1979*, IFIP, pp. 279–83.
- [13] IETF RFC 791, "Internet Protocol — Darpa Internet Program Protocol Specification," 1981.
- [14] A. Rybczynski et al., "A New Communications Protocol for Accessing Data Networks — The International Packet-Mode Interface," *NCC 1976 Proc.*, pp. 477–82.

BIOGRAPHY

RÉMI DESPRÉS (remi.despres@free.fr), after eight years of implementation experience on assemblers, compilers, and time-sharing systems, started working on packet switching in 1971. During successive first phases of the Transpac project, he was in charge of its technical aspects until 1980. After one year at Cap-Gemini-Sogeti as Telematics Division manager, he joined SESA in 1981. There, he contributed to the design of second-generation X.25 products and started a local area network activity. In 1985 he founded Réseaux de Communication d'Entreprise (RCE), a startup that sold LAN equipment at successively 1, 10, and 100 Mb/s, and provided the first frame relay concentrators of Transpac. After RCE was bought by the Compagnie des Signaux, he founded StreamCore, a startup specializing in TCP/IP bandwidth management. Leaving the company after financial investors and a new management chose a different direction for the company, he created RD-IPtech to work as an independent researcher and consultant on Internet technologies. As such he invented in 2007 the 6rd mechanism of RFC5569 whereby IPv6, the new-generation of IP, can be made available across unchanged IPv4 infrastructures. Having graduated as an engineer from Ecole Polytechnique in Paris (1961–1963), he holds M.S. and Ph.D. degrees in computer science and electrical engineering from the University of California at Berkeley (1967–1969).