

IBM France La Gaude Laboratory Contributions to Telecommunications: Part 2

Michel Bastian, Marc Boisseau, Robert Cohendet,
Alain Croisier, Claude Galand, Etienne Gorog,
Philippe Hernandez, Michel Humbert, Cuong Ngo Mai,
Pierre Secondo, and Robert F. Steen

IBM France's La Gaude Laboratory worked at the intersection of computing and electronic communications for more than 40 years. Its research led to specific technical contributions such as voice encoding, trellis-coded modulation, and many others. Part 2 of this two-part article discusses in detail the contributions, inventions, and concepts that were, at the time, landmarks in the state of the art, from the creation of the laboratory in 1959 through the early 1990s.

From its inception in 1959, the IBM La Gaude Laboratory was thrust into the middle of the convergence of computing and telecommunications. With the targeted recruitment of valuable engineers and scientists in France, and cooperation with other leading IBM laboratories, it soon became one of the world's most influential laboratories in the crucial area of applying telecommunications concepts to computer technology and vice versa.

The first installment of this two-part article dealt with the La Gaude's general contributions as well as its history, mission, projects, and products of historical significance. In addition to its general contributions to the telecommunications field, the laboratory also contributed specific products and processes in electronic switching systems, digital signal processing, and computer communications. In Part 2, we discuss in detail the innovative contributions, inventions, and concepts that were, at the time, landmarks in the state of the art.

Electronic switching systems

Engineers and scientists worldwide recognized the potential of computer-based telephone exchanges in the early 1960s. The IBM La Gaude Laboratory was one of the earliest computer company laboratories to explore the integration of these two technologies. Every part of the system posed unique challenges from the voice switch itself to the

provision of uninterrupted service in spite of computer failures. On the other hand, computer control also made new functions possible, which had to be invented. Even development process had to be reinvented to serve diverging national requirements in a common framework.

Integrated space-division switching network

Beginning in 1964, in the initial design phase of IBM's automatic branch exchanges, the design team considered all the possible switching principles for a network design. The dominant technology in the 1960s was the electromechanical matrix, or *crossbar switch*, which was gradually replacing the previous Strowger switch. Some new switching system designs used reed relays, tiny pairs of contacts enclosed in glass bulbs and operated by external induction coils. At the same time, other prototypes of PAX or private automatic branch exchanges (PABX) were built according to the principle of time-division switching whereby voice was converted to pulse-amplitude modulation (PAM). Digital voice, also known as pulse code modulation (PCM), had appeared in urban transmission in 1962, but the digital technology of that era could not yet support time-division multiplexing (TDM) switching of PCM streams.

In this new business venture, IBM believed that it could not depend on components produced by competitors. Telephone companies

produced most of their switching components, which eliminated crossbars or reed relays from IBM's available choices. Thus, La Gaude built its first design and laboratory prototype, a 30-extension-line PABX, with analog time-division switching. However, it soon became clear that the analog time-division approach (1) could not be scaled up to large switching networks and (2) could not match the stringent signal attenuation specifications imposed on the PABX by the common carrier administrations (these specifications closely matched the capabilities of metallic contact switches). Because it was a newcomer and apparent competitor, IBM could not expect the specifications to soften in its favor, so by necessity, it had to develop its own switching technology.

After investigating various electromechanical and electronic designs, the La Gaude Laboratory team arrived at an original solution based on a matrix of semiconductor cross-point switches. By the end of 1964 this solution became the team's leading option. With the help of the IBM Components Division, a physical design could be derived from the IBM transistor process with minimal changes.¹ Each cross point comprised a thyristor as a self-latching switch, plus a diode and a resistor integrated on a single piece of silicon.² These first commercially produced integrated circuit (IC) devices from IBM were manufactured on the new System/360 computers' regular production lines of logic transistors and diodes. Each integrated cross point measured 1 mm^2 (see Figure 1). The cross-point switch matrix was first used in a field-test prototype in 1966 and then commercially in 1969 with the IBM 2750 switching system.

The cross points were arranged in matrices connected in a three-stage folded network.³ The connections were established by matrix selection and held by constant current sources fed through the network (see Figure 2). Voice transmission was bidirectional, but unbalanced. That is, a single signal path was




Figure 1. IBM 2750 PABX. (a) Circuit diagram and (b) microphotograph of its integrated cross point (die size approximately 1 mm^2).

established over a common ground plane. This audacious design choice was made possible by the packaging technology La Gaude borrowed from IBM computer hardware. Attenuation was still a concern because each thyristor featured a residual impedance of about 6 ohms. This was compensated for by a negative-resistance circuit connected in series at the central junction point.

Along with the switching network development, IBM needed to develop a mass of auxiliary circuits, such as extension line and




Figure 2. Simplified schematic diagram of the IBM 2750 unbalanced bidirectional switching network.

trunk line adapters, internal junctors, status scanners, controller interfaces, and so on. Many design adaptations were also needed because of new regulatory requirements and the product's geographically wide coverage.

The new and powerful IBM space division technology from the La Gaude Laboratory remained competitive for many years, until the late 1970s when the progress in circuit integration allowed digital switching to take over and rapidly supersede all other options.

Control switch-over without service interruption

A telephone exchange or PABX provides a vital 24-hour, 7-days-a-week function, and service interruption should be a rare event. Thus, a PABX with centralized control must have redundant control components. In 1964 and 1965, the industry standard for redundancy was two identical controllers that operated synchronously. In the event of a controller failure, the other would assume the control tasks. This architecture had a serious shortcoming because it offered no protection against software failures, which could cause both controllers to fail under the same fault condition in the same operation cycle, resulting in a short service interruption and loss of all established connections. In response, La Gaude developed a different architecture, one with a design objective of a 50-year mean time between failures. This architecture had two controllers, preloaded with the same call processing and connection software, but only one of which was in charge of operations. The second controller remained in active stand-by, continuously checking itself. Both controllers sent "OK" signals periodically to a common switch-over control circuit, a device with a simple function, built of redundant—majority decision—logic circuits, and using an independent power source.

The confirmed absence of an OK signal from either processor would trigger the appropriate recovery action, sending an alert to the maintenance subsystem. If the problem originated in the active control unit, operations would be transferred to the stand-by unit by activating the physical connections to the switching network and telephone line attachments.

The switching network was self-sufficient and did not change its configuration during switchover, so that all established connections were maintained. The system would only lose the calls in the process of being established at the time of switchover. Callers would simply redial. The newly active controller

gathered information from the network to reconstruct its configuration information in memory, thus avoiding the use of possibly corrupted data from the failing unit.

Two additional features controlled this mechanism's proper operation. First, the controllers acknowledged and verified each OK signal the switch-over logic received. Also, every night during a low call period, a forced switchover was scheduled to verify the integrity of all control devices in each unit.

Advanced functions

One major additional advantage of program-controlled telephone switching was that new useful functions (beyond establishing connections between extensions, trunks, and tie-lines from A to B) were easily added and modified. Functions such as call waiting, call conferencing, and call holding were virtually impossible using current technologies. Market studies and focus groups with customers indicated that customers might accept various new functions. Program-controlled switching made so many options possible that the main challenge became one of reining in product planners and programmers who were eager to design and develop every new function imaginable. One major limiting constraint was that the new functions had to be accessible through a simple dial telephone or, at best, a touch-tone telephone handset. For cost and regulatory reasons, more sophisticated telephone instruments as we know them today could not be considered. Remarkably, many of the new functions introduced, such as double-call processing, conferencing, call-waiting signaling, and permanent personal phone numbers, have since become standard features in private and central exchanges. In addition, the IBM PABX offered specialized functions for corporate facilities. These included contact sense and activation for manufacturing control, access and attendance control, customer updating of telephone system configurations, and remote troubleshooting and maintenance.

More than a million lines of control code were developed at La Gaude between 1965 and 1970. The programming activity extended from code compilers to applications, including remote maintenance and the real-time control program—the system's heart. A major achievement was the first provision of direct inward extension dialing—that is, calling a specific extension without going through an operator, using individual public

phone numbers assigned to each extension phone. In Germany, a similar function existed in certain areas where the extension number could be dialed in addition to the public phone number. The IBM 2750 offered *German direct inward dialing*. Full direct inward dialing within the public numbering system was first introduced in Europe as a feature of IBM 3750.

Another achievement was the realization of corporate and government private networks of PABX, self-managed with signaling and routing facilities independent of the common carrier services. These private networks had advanced security features and used an integrated numbering system. In the 1970s, the French government acquired such a network from IBM for its own critical communications needs.

The IBM 3750 switching system offered remote maintenance in 1972. A support center received malfunction alarms, initiated diagnostic routines, isolated a subset of the system, and downloaded program updates. Altogether, this was a substantial innovation in the telecommunications world. Concurrently with the development of the IBM PABX, La Gaude also built an innovative computer-controlled test and qualification system known as TESA (Telephone Evaluation System Automat). TESA simulated the PABX's environment (users, trunk lines, and operators), checked its operation, and measured its performance under intensive traffic loads.

Single development process

Prior to IBM's entry in the field of telephone switching, PABXs were designed to national specifications and usually developed under supervision by the PTT laboratories. The installations including interconnection of the various component subsystems were assembled on the customer premises. IBM instead adopted a closed-box design, with a common base and preinstalled country-specific features. This approach, revolutionary in the telephony field, required extensive testing and demonstration by IBM to obtain national agreements in Europe.

In 1966, La Gaude operated a prototype system to prove IBM's closed-box feasibility and to obtain PTT's formal agreement before connecting the system to the public network in France. Beginning in 1968, one field test system was installed in a dozen major countries. Installations were scheduled according to local availability of skilled installers,

financial resources, telephone company agreements, and strategic interest. France (a second installation), Germany, the UK, Italy, the Netherlands, Denmark, Sweden, Brazil, Finland, Austria, Belgium, and the US participated in these system field tests using the IBM 2750 Switching System, which was still in development.

La Gaude's planning efforts, negotiations, field-test installations, PTT agreements, and operations experience afforded IBM, other manufacturers, and telephone administrations early awareness of these systems' requirements. Through such logistic and technical achievements, La Gaude was instrumental in accelerating and optimizing the standardization process for electronic PABXs and private networks in Europe and beyond.

Signal-processing architectures and algorithms

In the 1960s and 1970s, the early days of computer communication, IBM was a vertically integrated manufacturer. From component production to customer service, it produced nearly everything computers required, including transistors and logic circuits in large volume on automated facilities of its own design. These components were specified, designed, and intended only for computing applications. Accordingly, they fit in rigid automated design practices, from individual circuit conception to final assembly test and maintenance in the field. An independent internal organization, IBM Product Assurance, certified any purchased component not produced by IBM. Products developed with these components required additional engineering to meet the IBM manufacturing and maintenance organizations' operating requirements.

To the La Gaude Laboratory engineers developing telecommunication products, the use of in-house components was both an opportunity and a constraint. In those days, even the design of a simple amplifier was a challenge to an engineer constrained to transistors not specified for linear operation. Of course, as digital technology became cheaper and better over the next few decades, virtually all signal-processing designs moved to the digital domain. The constraints placed on designs from La Gaude in the 1960s forced the engineers to think digitally long before the practice became common, which helps explain why so much early work came from La Gaude.




Figure 3. Pulse sequence and resulting spectrum from early digital echo modulator.

The *digital echo modulation process* presented the first opportunity for a digital design in 1969.⁴ This involved the synthesis of signals from logical circuit outputs and summing networks of resistors. After generating a pulse sequence digitally and passing it through a low-pass filter, modulated signals with perfectly defined frequency and time specifications were formed (see Figure 3). Although perhaps obvious today, at the time this design was a significant innovation in generating precise amplitude and phase modulated signals. This technique was later extended to multiphase and quadrature amplitude modulation (QAM) signals.⁵

As logic circuit performance improved, digital techniques were also introduced in receivers. The underlying principles were known, but their application to voice band signals wasn't obvious considering the processors' large computation load and speed. A first step was the use of hybrid (partly digital) techniques in automatic equalization, based on delta modulation and analog transversal filtering.

The transition to purely digital processing resulted from the digital multiplex filter (DMF), an innovative signal-processing architecture developed at the La Gaude Laboratory in 1970. Other researchers independently developed a similar concept called distributed arithmetic.⁶ The DMF comprised a short transversal filter with four taps, where four successive signal samples could be multiplied with fixed coefficients and summed (see Figure 4). This circuit became a common building block in digital filtering. Several filter blocks were realized on the same




Figure 4. Digital multiplex filter (or distributed arithmetic)

hardware by time multiplexing on pre-assigned time slots.

The DMF was built in bit-serial arithmetic; it was the only affordable technology available, but relatively low in performance.

To accelerate processing, the DMF used a RAM to store the pre-computed sums of products for the 16 possible combinations of signal bits present at any instant at the shift register outputs. With this arrangement, the sum of the four products was executed within the time normally necessary for a single multiplication. Performance was further enhanced via a double accumulator where the propagation of carried-over bits was avoided—an arrangement known as carry saving accumulator.

A dynamic field-effect transistor (FET) DMF chip clocked at 2.5 MHz was fabricated in prototype quantities in 1972. At 416,000 partial products per second, its performance was remarkable for the time. The chips prototyped in an IBM manufacturing facility worked satisfactorily in laboratory applications. Unfortunately, the DMF turned out to be the only dynamic FET component in development at IBM and would have had to bear the whole cost of the manufacturing certification process; the project was therefore abandoned.

Early stored-program signal processor

Besides its IBM certification problems, the DMF was difficult to use in the design of modems because of its synchronous nature: each filtering operation had to be assigned a time slot and its output re-inserted in a different slot. This required external storage components and control that were more complex than the filter itself. To make things worse, customer requirements for modems became increasingly complex and included many different transmission modes in the same hardware depending on the type of transmission line, fall-back situations, and national regulations. In addition, the need for fast synchronization on multipoint lines led to a requirement for operating the modem in a completely different mode during the signal acquisition phase.

The only way to design for this increasing complexity was if La Gaude abandoned the DMF synchronous architecture for a fully programmable asynchronous architecture with signals stored in memory. Synchronization was invoked using interrupt signals only for I/O operations. Although this choice is obvious today, La Gaude was telecommunications, not computer oriented. For telecommunications engineers, the change from synchronous to asynchronous design was a sort of cultural revolution.

The first result of this revolution was a two-tiered processor in metal-oxide-semiconductor field-effect transistor (MOSFET) technology, one of the first signal processors in commercial use.⁷ On the signal level, a digital filtering module (DFM) working in parallel on 16 bits and fitted with a double arithmetic unit could accumulate up to 1.6 million products per second. On the control level, a general-purpose microprocessor (GPM) with a 36-bit-wide instruction word was able to effect four elementary operations in a single 600-ns cycle: data input, execution, storage, and conditional branch. This specialized architecture permitted interrupt processing without any delay.

This processor pair (see Figure 5) development started in 1976. It was commercially introduced in 1978 with the third generation of IBM modems ranging from 2,400 to 9,600 bits per second (bps).

High-performance single-chip processor

As technology and algorithms evolved, the two-tiered architecture became inadequate in cost and computing power for two reasons. First, MOSFET processor performance became insufficient to implement the emerging signal processing algorithms. For example, the demodulation of trellis-coded signals alone required 2 million products per second versus 1.6 million available in the DFM chip. Second, integrating the two tiers on a single chip was highly desirable to reach higher performance levels and to keep costs acceptable, but integrating the existing two-tier design would have exceeded the limit, approximately 5,000 logic gates, of available technologies.

To address these constraints, La Gaude engineers, in cooperation with the IBM research laboratory in Zurich, started in the late 1970s to develop a new signal processor architecture.⁸ This architecture minimized the number of logic gates necessary while achieving the performance of a complete product and accumulation in a single operation cycle. The heart of this processor was an




Figure 5. The first programmable signal processor introduced in 1976 comprised (a) separate arithmetic and (b) control elements.

innovative and simplified multiplier unit (see Figure 6). The multiplier subsystem working on 12×12 bits, with a 20-bit output, receiving no program instruction, was wired in parallel with the arithmetic logic unit (ALU) input registers, and it multiplied any data present, meaningful or not. The software decided if the results found in its output register RM were to be used or not. Two registers RA and RB alternatively accumulated partial products; with proper instruction sequencing, they allowed convolution products to be chained without the loss of a single cycle for data input. Altogether, eight data registers were used: two for data input, one for multiplier output, two for accumulation, two for indexing, and one for sequencing. A large instruction word permitted the execution of several elementary instructions in a single cycle.

Implemented in 1981 by the IBM Corbeil-Essonnes laboratory in bipolar technology with a 100-ns cycle time, this processor performed up to 10 million products and




Figure 6. An efficient signal processor design of 1981 based on a freewheeling multiplier and dual accumulator achieved 10 million multiply-add per second.

accumulations per second and 15 to 20 million elementary instructions, largely outperforming commercially available devices.⁹ Two years later the same team also implemented the processor in complementary metal-oxide-semiconductor (CMOS) circuitry at higher processing rates.

Subband coding with QMF

Encoding voice or music digitally presents unique problems, especially if it must be encoded at low-bit rates. La Gaude was working on this problem in the 1970s, long before mobile phones and digital musical CDs. To the listener, the imprecision necessarily introduced by low-bit-rate coding is perceived as noise added to the signal. For a natural-sounding reproduction, this added noise must be hidden from human perception. It is possible to hide the perceived noise if its frequency components are close to those of the coded original signal. All rate-effective coding processes try to exploit this masking effect, a psycho-acoustical phenomenon, known since the 1920s.

One way to mask the noise is to divide the signal channel into narrow frequency bands and code each band separately with a precision adapted to the actual signal level (see Figure 7). The individual subbands are then recombined after decoding to reconstruct the original signal. However, this approach

has long been unattainable due to the cost and design difficulty of the ideal narrow-band filters required for perfect separation.

If frequency band separation is imperfect, the sampling process in each subband introduces undesired signals, known as *image frequencies*, in the range of frequencies where the subbands overlap. These parasitic signals have no harmonic relation with the original signals and so are easily perceptible. Before the invention of the quadrature mirror filter (QMF), some designers attempted to use sharp filters typically crossing over at the -40 dB level, but such long filters introduce unacceptable transmission delays, distortions, or echoes.

The QMF technique made it possible to circumvent this design problem by allowing overlap of adjacent frequency subbands.¹⁰ The key was to define the filters so that they generated image frequencies that canceled each other on signal reconstruction (see Figure 8).

QMFs are applied to previously digitized signals. They are implemented as transversal digital filters and therefore have perfectly defined properties. Despite their apparent complexity, codecs operating in subbands are easily implemented in signal processors. The same algorithm can be repetitively used in each of the *N* channels, where the sampling rate is *N* times lower than that of the original signal. Therefore, the signal processor's workload and complexity increase only slightly with the number of subbands.¹¹

The original QMF concept and theory, and the first practical realizations, were developed at La Gaude in 1976.¹² Many researchers have since further explored the field, and QMF is now a basic chapter in the teaching of digital signal processing.




Figure 7. Coding audio signals in subbands.

The most popular application of QMF today is in audio coding in MP3s, where an evolution of QMF (polyphase filters) is applied. The same principle is also applied in many other audio and video coders.

GSM voice coding


For La Gaude's speech processing group, the Global System for Mobile Communications (GSM) contribution began in 1985 in cooperation with the French administration's Centre National d'Etude des Télécommunications R&D] (CNET). CNET was contributing to the specification of the pan-European mobile telephone network (which became the GSM), through the Groupe Special Mobiles [special group for mobiles] of the Conférence Européenne des Postes et Télécommunications (CEPT), an organization in which PTT and industrial partners cooperated.

The CNET laboratory in Lannion, France, focused on developing a 16-kbps speech coder, but it had no applicable research at 13 kbps, the lower bit rate required for GSM. Its specialists were aware of proposals for new designs in the 8- to 16-Kbps range by La Gaude engineers, so it prompted La Gaude to enter the selection process with a proposal of its own.

The voice coder IBM proposed at CNET's request was chosen as the official French contender after CNET conducted comparative tests at Lannion in July 1986, and it compared favorably with other European proposals at Turin, Italy, in 1986. After a technical experts' meeting in Den Haag, the Netherlands, two contenders remained: the IBM proposal based on MPE/LTP (multipulse excited with long-term prediction) and one by the Philips PKI branch in Germany based on RPE (random pulse excited). The Philips proposal had a simpler design, but it was more sensitive to transmission errors. IBM rapidly analyzed both and proposed merging the two designs to combine their respective advantages. Cooperation with Philips' PKI resulted in the 13-kbps RPE/LTP coder that was tested and approved by the experts' group and eventually received final approval in July 1987. Commercially introduced in 1991, the GSM network today connects more than 3 billion mobile phones and accounts for 89% of the world's digital mobile phone market.¹³

Data transmission

The La Gaude Laboratory contributed multiple innovations to the data transmission



The left hand diagram shows a signal component, nominally situated in the A band is in fact present in two adjacent frequency bands due to filter overlap. As each sub-channel is re-sampled at the lower rate consistent with its bandwidth, frequency spectrum folding occurs and creates image components on either side of the crossover point (center diagrams).

Then after transmission comes the adaptive decoding of each channel at the receiving end. Output (reconstruction) filters used before the recombination at the channel sampling rate eliminate out-of-band components and introduce additional frequency spectrum folding, but the image frequencies cancel each other when the original signal is reconstructed as shown in the right-hand diagram.

Figure 8. Operation of a pair of QMF filters.

field, from the application of digital techniques in modems already evoked to advanced modulation and coding techniques and digital network protocols. In addition to be a competitive development laboratory, it also actively contributed to the development of international standards.

Trellis-coded modulation

Trellis-coded modulation (TCM) is perhaps the most important contribution IBM has made to the data transmission field. The credit goes to Gottfried Ungerboeck, then a member of the Zurich IBM Research laboratory, who was on assignment at the La Gaude Laboratory in the mid 1980s.¹⁴

Introducing redundancy plays an essential role in high-performance data transmission. According to Claude Shannon's 1948 theory of information,¹⁵ redundancy permits a system design that can approach the capacity limit of transmission channels by making legitimate sequences more distinguishable from illegitimate ones in the presence of noise and distortion. This technique is widely used in space communications because noise, not available bandwidth, is the main issue. With convolutional codes redundancy can be introduced in the binary data and decoded

The principle of TCM is illustrated here using a simple example. The top two diagrams represent classical 4-phase modulation. The bottom two represent 8-phase TCM modulation.

The X, Y coordinates represent the in-phase and quadrature signal components, the numbered dots identify the possible signal states.


The numbered arcs indicate which signal state can be transmitted at any time. Either system transmits 2 bit/s per Hertz of bandwidth.

In the upper half is a classical 4 phase modulation: for each new signal element, one of the four possible states is chosen as a function of the two next bits of data, independently of past history.

Below, is an 8 phase trellis coded modulation: for each new signal element, one of the 8 phases is chosen, but this choice is limited to only four possibilities, following a precise rule which takes the previously transmitted signals into account. The rule is designed to maximize the difference between legitimate sequences.

In this example, assuming a signal-to-noise ratio of 9 dB, trellis coded modulation provides a fifty fold improvement in bit error rate, from 0.5 % to 0.01 %.

Figure 9. An example of trellis-coded modulation.



continuously, particularly by means of Viterbi's probabilistic decoding algorithm.¹⁶ In this environment, the frequency band used for transmission is increased to account for the redundant signals; the transmitted energy is spread thinner over more bandwidth.

Until 1976, however, this technique had not been successfully extended to band-limited transmission channels, such as a telephone line. For example, transmitting more than 2 bps per one Hertz (Hz) of bandwidth (for example, more than 6,000 bps over 3,000 Hz) requires complex signals: a group of 3, 4, or 5 bits are represented by a single elementary signal unit that might take 8, 16, or 32 distinct states of phase and amplitude. Therefore, errors on the transmitted bits are no longer independent, and binary redundancy schemes are ineffective.

The critical breakthrough came with TCM, combining redundant coding and modulation in a single process instead of treating them independently (see Figure 9). This invention made high-speed transmission over telephone channels a practical reality. Specifically, an enhanced error rate allowed full-duplex transmission at 9,600 bps over the

telephone switched network with roughly 3,000 Hz of available bandwidth. Later, data rates were extended to 14,400, 19,600, and 28,800 bps. These designs were adopted by CCITT recommendations between 1985 and 1991. Modems built to these specifications were produced in large quantities by many manufacturers, but by then, IBM had abandoned the modem business. Nonetheless, TCM-based modems and large-scale ICs were key factors in the massive deployment of personal communications ranging from facsimile to Internet access.

HDB3 transmission code

At the end of the 1960s, most European PTTs considered developing pulse code modulation (PCM) transmission facilities similar to the American T1 system, but this standard deployed about five years earlier was considered too restrictive for future evolution. Beyond the interconnection of local telephone exchanges, switching and higher multiplexing over long distances were needed. A standard allowing the interconnection of PCM networks was also required because of Europe's multinational environment. A Special D commission of

In the "bipolar" pseudo-ternary code of the T1 system, binary 1s are transmitted as alternatively positive and negative pulses; binary 0s are not transmitted, they are detected as an absence of pulse, either positive or negative.

In order for the line repeaters to remain synchronized, a minimum density of non-zero signals is however needed. The T1 telephone system addressed this requirement by restraining the possible values of speech samples, a clearly unacceptable constraint for data systems.

Data transparency can be regained by the introduction of "stuffing" pulse sequences when long strings of zeroes are present. In order to distinguish stuffing sequences from valid data, it is necessary to violate the rule of alternating polarities, i.e. send two consecutive pulses of the same polarity, a technique used in T1 for frame synchronization. Too many polarity violations however could lead to signal unbalance and impair the repeater operation.

The originality in the HDB3 and CHDB codes is the use of two different stuffing sequences as a function of the already transmitted signals, under a rule which minimizes signal unbalance. Thus it becomes possible to suppress any sequence of more than three zero-level signals and maintain robust repeater synchronization.

Shortly after this proposal, AT&T introduced a similar code B6ZS which is implemented in its later PCM systems, but can only remove sequences of 6 zeroes or more.




Figure 10. HDB3 operation.

CCITT was created for this purpose in 1968, and a representative of the IBM La Gaude Laboratory was invited along with the French PTT delegation.

More than 10 different designs were submitted during these meetings. They differed in bit rate, voice sampling frequency, transmission mode, and signaling channel assignment.

For data transmission, AT&T proposed an asynchronous transmission scheme that provided less than one-third of the T1 raw capacity and a channel assignment scheme incompatible with voice channels. On the other hand, IBM emphasized the value of a standardized digital channel, which could transmit either voice or data up to its full raw bit rate and could be used for international interconnection. In contrast to the T1 system, the required conditions for this capacity were that the channel data rate and voice sampling frequency be uniquely defined and that no telephone signaling be transmitted in the voice channels.

Assuming that these specified conditions could be met, an important problem remained: the line transmission signal in

current designs like T1 was not transparent to the data being transmitted—that is, not all data sequences could be sent. The pseudoternary transmission code on these lines represented logical ones with pulses of alternating polarity (for example, the first "one" occurring was a positive-level signal and the next a negative-level signal), and zeroes were coded as an absence of signal (zero-level signal). Longer transmission lines had repeaters that regenerated and reshaped the signal and depended on a certain number of signal transitions to stay synchronized. These signal repeaters would therefore lose synchronization if a long sequence of zeroes was sent, a condition obviously inadequate for transparent transmission of unrestricted binary data sequences. At the time, most European systems either in the design phase or in a field test used this type of repeater in their transmission facilities.

An IBM invention at La Gaude in 1968 and 1969¹⁷ resolved the dilemma by proposing a new family of pseudoternary codes. The codes guaranteed a minimum density of pulses regardless of the data content (see Figure 10)

**The La Gaude
Laboratory recognized
the importance of
international standards
in the early days of data
communications.**

and could be transmitted through the existing lines with repeaters needing transitions. These codes were dubbed high-density bipolar (HDB) and compatible HDB (CHDB).

IBM decided to make these codes available in the public domain. One of them, HDB order 3 (HDB3) was adopted by CEPT, and in 1970, the CCITT eventually reached a historical agreement that reduced the available number of different PCM systems from 10 to 2—T1 at 1.44 Mbps in North America and D1 at 2.048 Mbps (including HDB3) in Europe. Other countries could choose either T1 or D1. Common to both systems was voice sampling at 8,000 samples-per-second and 8-bit time slots. Data transparency was total in Europe, but limited to 7 bits per slot in North America.

A workable scheme was thus available for worldwide voice transmission, and many other types of signals could be transparently transmitted throughout Europe. Complete signal integration was achieved 14 years later with the introduction of Integrated Services Digital Network (ISDN).

Error-correcting codes

Many error-detecting codes were in use in 1960, such as the parity check, which detects a single bit error in a character, and the IBM serial transmit receive (STR) code, which consists of an alphabet of 8-bit characters, where each character is made up of four zeroes and four ones yielding 70 valid bit combinations among the 256 possible. IBM initially used the STR code with magnetic tape transmission. When the receiving equipment detected an erroneous character in a message block, the complete block was retransmitted.¹⁸

Cyclic codes were simple to implement with digital computer circuits, thanks to n -stage shift registers. The n bits remaining in the register were added to the initial message to be transmitted, and the same register

at reception, detected the potential errors. Additional logic, at the receiver, allowed automatic correction of the most probable errors.

By the end of the 1950s, several mathematicians had already proposed correcting codes for classical errors such as multiple bit errors in messages of a given fixed length. In the case of cyclic codes that matched well to data transmission, the development of error-correcting codes was essentially based on mathematical models using the finite field theory, also known as the Galois field theory in honor of the French mathematician Évariste Galois. Each code was characterized by a polynomial with binary coefficients corresponding to the shift register connections. Primitive and irreducible polynomials played an essential role in the choice of the best codes for a given job.

The La Gaude Laboratory scientists developed a new mathematical model showing that any polynomial, even if it did not belong to a finite field, had self-correcting capabilities. In this model, the error type to be corrected could be a previously studied one or a new error configuration. This discovery led to algorithms that permitted computers to generate cyclic codes able to correct a given error type, with minimum redundancy, for a given message length.¹⁹

In 1962, La Gaude published tables of the best cyclic codes able to correct two classes of errors. The first class contained error types already studied. For instance, to correct a burst of five errors maximum, wherever it occurred within a message of 1,300 bits, the best known code at the time required 17 bits of redundancy. The table generated by this new algorithm showed the existence of a cyclic code requiring only 15 bits of redundancy. As there was, and still is, no code (cyclic or not) able to accomplish this with fewer than 15 bits, this newly published code reached the absolute limit (known as an optimum code).

The second class of errors addressed new error types, such as those encountered in tests conducted at 1,200 bps over European telephone lines. As an example, there is a code using 13 bits of redundancy, capable of correcting two bursts of 2 errors within a zone of 8 bits; this zone could occur anywhere in a message of 113 bits. The results concerning the second class of error types were new and were more relevant to the errors encountered in data transmission.

In the 1960s and 1970s, when computers required a high level of data-transmission

integrity, these error-correcting and detecting code tables were often used. For example, NASA used them to secure one of the control commands of its Apollo program. These tables also exposed some of the incorrect proposals for error-protecting codes made to the Special A working group of the CCITT.

Standards development

The La Gaude Laboratory recognized the importance of international standards in the early days of the development of data communications. As a result, starting in 1965, it held worldwide responsibility (US excepted) for coordinating IBM participation and technical contributions to standards organizations at all levels, whether international (e.g., ITU/CCITT and ISO), European (e.g., ECMA and ETSI), or national (e.g., AFNOR and BSI).

La Gaude's contributions to standards development included bringing IBM's technical innovations to the attention of standards organizations. In the 1960s and 1970s, such contributions largely concerned data modems and interfaces (e.g., CCITT recommendations of the V series), protection against errors, data-link protocols, signaling (in particular by frequencies or tones, such as Q.23 bis), and packet switching (X.25). In the 1970s, IBM contributed a variant of its bit-oriented Synchronous Data Link Control (SDLC) protocol still under development in its US labs, which allowed better performance than the previous character-oriented protocols. A slightly modified version of this protocol, eventually adopted by ISO as HDLC and by CCITT as the X.25 level 2 protocol, is universally used in digital networks today.

From the end of the 1970s and during most of the 1980s, interconnection and interoperability between heterogeneous data processing systems became a crucial concern, requiring new international standards. Also, starting in 1975, public data networks began to appear. After many collaborative efforts between telecommunications operators and standards committees, circuit- (X.21) and packet-switched (X.25) public data networks were introduced, mostly in Europe. Meanwhile, IBM's Systems Network Architecture (SNA) had been independently developed and implemented worldwide since early in the 1970s.

IBM system architects at Raleigh and La Gaude contributed to the detailed analysis of interconnection and interoperability problems, thus strengthening the specialists' mutual understanding and cooperation.

Consequently, the layered network architecture of IBM's SNA, with a detailed definition of the services and protocols for each layer, a new concept at the time, provided a conceptual framework for standards development. The Open Systems Interconnect (OSI) model as well as the Internet are based on a similar structure, thus benefiting from this early standards work.

Other major contributions to standards originating at La Gaude included TCM for high-speed modems and voice coding for GSM mobile phones. La Gaude also held a major role in the standardization of network management. An expert from the laboratory was appointed "Rapporteur for Network Management" by the CCITT for its Study Group XVII. Other standards work by La Gaude applied to the development of ISDN, frame relay, ATM, data security, and OSI conformance testing standards, among others.

The significance of standards development at La Gaude was reflected by its influence in propagating new data communication and networking standards worldwide. Through its involvement in European standards organizations, the IBM La Gaude Laboratory gained recognition on both sides of the Atlantic. Its expert testimony often helped standards bodies achieve consensus where instead national prejudice might have prevailed.

Conclusion

The IBM La Gaude Laboratory's contributions run the gamut of transmission codes, modem technologies, error-protection codes, low-bit-rate voice coding, public packet network attachment, line switching and PABX technologies, distributed computer protocols, and international telecommunications standards. The products La Gaude developed or made possible include modems, communication controllers, PABXs, networking software, mobile phones, and data networks.

The history of telecommunication and computing, and especially their intersection, cannot be told without the story of the IBM La Gaude Laboratory.

References and notes

1. L.L. Rosier, C. Turrel, and W.K. Liebmann, "Semiconductor Crosspoints," *IBM J. Research and Development*, vol. 13, no. 4, 1969, pp. 439-446.
2. E.Y. Rocher and R.E. Reynier, "Response Time of Thyristors: Theoretical Study and Application to Electronic Switching Networks," *IBM J. Research and Development*, vol. 13, no. 4, 1969, pp. 447-455.

3. R.E. Reynier, "Electronic Switching Network of the IBM 2750," *IBM J. Research and Development*, vol. 13, no. 4, 1969, pp. 416-427.
4. M. Humbert, "Four-Phase Digital Echo Modulation Technique," *CCITT Communication*, special A contribution no. 150-E, 30 Oct. 1967; J.M. Pierret, "Modulation digitale à échos," *Colloque international sur la téléinformatique* [International colloquium of remote data processing], Paris, 24-28 Mar. 1969 (in French); and A. Croisier and J.-M.D. Pierret, "The Digital Echo Modulation," *IEEE Trans. Comm. Technology*, vol. COM-18, no. 4, 1970, pp. 367-376.
5. F. Choquet and H.J. Nussbaumer, "Generation of Synchronous Data Transmission Signals by Digital Echo Modulation," *IBM J. Research and Development*, vol. 15, no. 5, 1971, pp. 364-377.
6. A. Peled and B. Liu, "A New Hardware Realization of Digital Filters," *IEEE Trans. Acoustic, Speech, Signal Processing*, vol. ASSP-22, Dec. 1974, pp. 456-462.
7. P. Thirion, "Digital Signal Processing with Program Synchronization between Two Microprocessors," *IEEE Trans. Comm.*, vol. COM-26, no. 5, 1978, pp. 513-517.
8. G. Ungerboeck et al., "Architecture of a Digital Signal Processor," *IBM J. Research and Development*, vol. 29, no. 2, 1985, pp. 132-139.
9. J.P. Beraud, "Signal Processor Chip Implementation," *IBM J. Research and Development*, vol. 29, no. 2, 1985, pp. 140-146.
10. A. Croisier, D. Esteban, and C. Galand, "Perfect Channel Splitting by use of Interpolation/Decimation/Tree Decomposition Techniques," *Proc. Int'l Conf. Information Science System*, 1976.
11. C. Galand and H. Nussbaumer, "New Quadrature Mirror Filter Structures," *IEEE Trans. Acoustic, Speech, Signal Processing*, vol. ASSP-32, June 1984, pp. 522-531.
12. D. Esteban and C. Galand, "Application of Quadrature Mirror Filters to Split-Band Coding," *Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing*, IEEE Press, 1977, pp. 191-195.
13. Data from GSM, February 2009, <http://www.gsmworld.com/>; also from GSA <http://www.gsacom.com/news/statistics.php4>.
14. G. Ungerboeck, "Treillis-Coded Modulation with Redundant Signal Sets, Part I: Introduction," *IEEE Comm. Magazine*, vol. 24, no. 2, 1987, pp. 5-11, and "Treillis-Coded Modulation with Redundant Signal Sets, Part II: State of the Art," *IEEE Comm. Magazine*, vol. 24, no. 2, 1987, pp. 12-21.
15. C.E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical J.*, vol. 27, July & Oct., 1948, pp. 379-423, 623-656.
16. A.J. Viterbi, "Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm," *IEEE Trans. Information Theory*, vol. IT-13, Apr. 1967, pp. 260-269.
17. A. Croisier, "Introduction to Pseudoternary Transmission Codes," *IBM J. Research and Development*, July 1970, pp. 354-367.
18. E. Gorog, "Redundancy Alphabets with Desirable Frequency Spectrum Properties," *IBM J. Research and Development*, May 1968.
19. E. Gorog, "Les meilleurs Codes Cycliques correcteurs d'un paquet d'erreurs dans un message de longueur donnée" [The Best Cyclic Codes Correcting an Error Packet in a Given Length Message], IFIP Congress, Munich, 1962 (in French); E. Gorog, "On the Computation of Minimum Redundancy Cyclic Codes Able to Correct Given Error Pattern Types," *Inst. Statistics*, mimeo series no. 484.6, Univ. North Carolina and Univ. Paris, presented at NATO Summer School, Royan, Aug. 1965.

See Part 1 for a full list of author biographies.

For further information on this or any other computing topic, please visit our Digital Library at <http://computer.org/csdl>.

Punched-Card Systems and the Early Information Explosion, 1880-1945

Lars Heide

Early punched cards helped to process the United States census in 1890. They soon proved useful in calculating invoices and issuing pay slips. As demand for more sophisticated systems and reading machines increased in both the United States and Europe, punched cards served ever-larger data-processing purposes. Insurance companies, public utilities, businesses, and governments all used them to keep detailed records of their customers, competitors, employees, citizens, and enemies.

Lars Heide presents the first academic study of the invention that fueled today's information revolution: the punched card.

\$65.00 hardcover



THE JOHNS HOPKINS UNIVERSITY PRESS 

1-800-537-5487 • www.press.jhu.edu