The battle between Wi-Fi (IEEE) and HiperLAN (ETSI)

Philippe Jacquet,
Nokia Bell Labs France

Abstract:
The title does not do justice to the suspense. If you're not from Mars, you know what WiFi is. A riddle: 25 years ago, WiFi (under standards from the Institute of Electrical and Electronics Engineers, IEEE) and HiperLAN (under standards created by the European Telecommunications Standards Institute, ETSI) were in the same boat. Which one fell into the water?

At the origin, Wi-Fi (not yet called Wi-Fi but instead IEEE 802.11) and HiperLAN were projects of standardization for local wireless networks; the former, American, supported by the Institute of Electrical and Electronics Engineers (IEEE); the latter, supported by the European Telecommunications Standards Institute (ETSI).[1] Always the same question: did HiperLAN vanish into thin air? Isn't Wi-Fi a kind of HiperLAN?[2]

You say, "The dice were loaded: on the one side, the Americans, powerful, arrogant and united; on the other side, the small-fry, Europeans, timid and not united." In fact, this statement has it completely wrong, at least in telecommunications, the field I know. Let's take the example of mobile telephony.

In the 1980s, three standards coexisted in the United States, each with its initialism (CDMA, etc.). Their only common point was that they were fully incompatible with each other. The American strategy of standardization was as follows: take the strongest solutions that require no effort for reaching an agreement, and jealously keep their patents; then package them with a pretty ribbon, and you have a bundle of standards. In this case, the products existed before publication of the standards; and operators picked what they liked best from the package. As a consequence, many a network spanned North America. When traveling, Americans needed to take along three telephones, each corresponding to a different standard, and had to steel themselves with patience to be able to remain in communication.

The European strategy was the reverse. To force companies to reach an agreement, they were asked to invent a single, innovative system and share their patents. To this method's disadvantage, the products did not come first, before the standard; and Europe ran the risk that the market would slip out of its hands. By doubling its efforts, Europeans brought out the standard for a global system for mobile communications (GSM) within a few years, between 1982 and 1987. It was superior to its American predecessors in every way (connection speed, range, mobility) with a bonus: a single standard for a single market. Indeed, the market exploded. Within a few years, this standard had spread around the world: Africa, the Near East, then Asia… and finally, under pressure from users and operators, the United States.

[1] This article has been translated from French by Noal Mellott (Omaha Beach, France).
[2] See the articles "IEEE 802.11" (legacy mode) and "HiperLAN" in Wikipedia.

Why was it such a hit? No, the main reason was not the quality and vitality of Europeans. Engineers in Europe were not essentially stronger or weaker than engineers in the United States. Besides, the same industries took part in setting standards on both shores of the Atlantic. On the one side of the pond however, they were to keep to their own hunting grounds while, on the other side, they had to hunt for an understanding… and managed to do so. I wasn't there, and I don't want to tell secondhand stories about how fax machines were used during breaks between the sessions devoted to the standardization of GSM.

If the method worked so well, why wasn't it applied to drafting a standard for local wireless networks? This story starts in 1991. Before telling it, and parsing it to see what did not work as expected (but what did not turn out to be the opposite of what had been expected), I would like to take a break to describe the difference between a mobile cellular network and a local wireless network.

In a cellular telephone network, communications pass over thousands of kilometers of electric cables as compared with a few hundred meters (or at most one kilometer) by radio transmission — the distance that separates the subscriber's mobile phone from the base transceiver station. On the other side of the relay antenna, it's all wire; and on this side, all cells. On this side, the conductor of the orchestra is, beyond any doubt, the base transceiver station; and users' telephone receivers (normally handsets) are slaves. Receivers never emit without being asked to do so by the station. They never communicate with each other. It all takes place through the base station. After all, the receivers have a very low power of emission and are not very sensitive when receiving. Even at a distance of a few dozen meters, two handsets would not hear each other. In contrast, the base station has big ears and a big mouth. Sensitivity and power are both concentrated on one side. Cellular network technology is asymmetric. For this reason, the telephone receivers that use less technology are cheaper and last longer.

During this "prehistoric" era, cellular networks were devoted to routine telephone communications — no streaming, no Facebook, no Web. In contrast, wireless networks had to bear the brunt of sporadic, violent transfers of big chunks of data like those used on the Internet. The base station antenna, which routinely collected and distributed a hundred times per second tiny bits of telephone conversations, was overwhelmed. The conductor could no longer beat time since it did not know the score in advance, and since the chunks now had to be transmitted as fast as possible.

Besides, the computers connected in the network have equal rights. Local wireless networks are a fully symmetrical technology. Imagine a cloud of radio waves in a room or building, where all terminals are competing to communicate with each other. The daunting task is to avoid collisions or at least to solve the problem of collisions without any conductor leading the orchestra. Of course, one of these terminals might be a gateway toward a cable on a traditional network; and in this case, it could be equipped with a little green light (to make it look fancy) and with two antennas to actually cover the room or building. But even in this case, the technology inside the box was still the same as in the other terminals.

The key factors in this setup are the flow rate of bits to be transmitted and the protocol for handling collisions. Let us take the time to examine each of these. We shall see how the first was fatal to GMS and how the second was a probable cause of its resuscitation.

When I mention connection speed, some readers start squirming, especially if they are using Wi-Fi to browse this text. Since we want speeds from one hundred to one thousand times higher than what sufficed for a telephone, we might imagine that it would also be necessary to increase the power of the emitted signal… and, in an additional step of the imagination, that the cell phone might become a toaster. Let me reassure you right away: we are under the protection of information theory. May I have the pleasure of slipping a little equation into this text, if only to challenge readers?

In 1948, Claude Shannon, an engineer at Bell Laboratories, proved that the expression $W \log_2(1+S/B)$ yields the capacity of a radio communication, where W is the radio bandwidth used by the transmission, S the intensity of the received signal, and B the interfering noise picked up by the receiver. The base 2 of the logarithm comes from the fact that the result is yielded in bits per second (bps). Quite clearly, if we want to multiply the transmission capacity by 100, it would be necessary to multiply the power of the signal by 2^{100} — unaffordable. But we're in business if, instead, we multiply the bandwidth frequency by 100! In fact, Shannon's law is intuitive. If you play a piece of music that has more notes in it, there's no use turning up the volume. To come back to our local network, the power of the emitted signal can even be lowered compared with a portable telephone since emissions are over a short distance.

The American standard based on existing products in the 1980s was part of a technology that had proven its mettle since WW II: the spread spectrum technique, a rather technical subject that I shall not touch. It allowed for reaching a very decent speed of one million bits per second (Mbps). The European standard was immoderately chasing after a speed of 23 Mbps. To realize what this means in terms of speed, imagine that if HiperLAN had been the Concorde flying across the Atlantic in 1969, then Wi-Fi would be the Blériot XI crossing the English Channel in 1909! Clearly, the one was too far in advance; or else the other, too far behind. To reach 23 Mbps, the ETSI work group wagered on an equalization technique, which had worked for GSM. When applied to HiperLAN however, it required huge calculations at a speed between 100 and 1000 times higher than on GSM. That was possible with Fourier transforms but probably not suitable for technological products for the general public in the 1990s! I recall that the first (rather rudimentary) prototype presented to the European Commission in 1995 had an inconspicuous electric wire connecting the emitter to the receiver in order to ensure synchronization, which we could not yet do by radio. This would turn out to be a tripwire that, along with the lack of enthusiasm among industry representatives, caused HiperLAN to stumble and explode.

What about the protocol for handling collisions? The situation turned out in favor of HiperLAN. An efficient algorithm was, since the 1970s, being used in cabled networks of an Ethernet type. Ethernet (not to be confused with the Internet) is a standard for local wireless networks with multiple types of access. It has become the dominant standard for local networks. We might say that Ethernet networks did as much using cable as Wi-Fi and HiperLAN were to do using radio bands. The principle is simple: when several data packets risk colliding, they are retransmitted after short intervals chosen at random. If a new collision takes place, the algorithm is repeated until, by insisting, the packets end up passing. It's a little like encountering someone in a hallway and choosing randomly whether to pass to the left or the right… until you pass. The solution for settling the collision problem is randomized. It can take a very short time or very long, especially when other packets are involved. The basic point is that this procedure calls for no centralized supervisory authority.

Wi-Fi's ancestor opted for the Ethernet algorithm, a wise and economical solution since it did not imply too many modifications in the computers made at the time. However the first IEEE committee omitted a key point: Ethernet required a rapid detection of collisions in order to immediately interrupt transmission before the situation would get out of hand — a little like rivalry for a place in a parking lot. If the fenders of the two cars just barely brush, the drivers separate on good terms and look for another parking place. If, however, the cars snag each other, the parking lot will fill up with wrecks once the admission rate of cars rises above a certain threshold. Well, it is impossible to detect a collision on a radio band. On a cable, you need but compare what you receive with what was emitted. If you observe a difference, you know there was a collision. This observation is possible because the signal passing through a cable is barely attenuated. What enters comes out with nearly as much power. Radio emissions are different. The signal is emitted in a three-dimensional space; it is broadcast — cast broadly all around — over a large distance. The ratio of the signal emitted by an antenna and the signal received at a distance of three meters is more than one billion — no sense comparing the two methods!

So, we proposed an algorithm for HiperLAN that we called the “comb method”. It alternated very short phases of emission and reception before emitting another data packet. Situations of collision could thus be detected and avoided. Furthermore, this comb could be calculated so as to filter the packets as a function of their priority or urgency. This option made it possible to superpose the data exchanged for streaming a video by slightly delaying the exchanges of files that were not urgent in order to let pass video packets during peaks of congestion.

Now, the main question: how to choose the right standard? A particularity of ETSI is that it receives European funds for its role as a scientific expert on standards. This is not the case of the IEEE. The particularity of European standards is that the EU finances programs for building their prototypes. It might seem like a truism to say that the United States does not have any European programs, but this is false, since the major industrial players in this technology have subsidiaries in Europe that are fully eligible for funding. What is true is that no equivalent arrangements exist in the United States for forming an alliance between industry and science in pursuit of an ambitious program. American academics receive funding from the National Science Foundation (NSF) for the papers they write and their work with students, and that’s it. The government contracts received by American industries are oriented toward finished products. An industry may bill thrice its cost price to cover the burden of research upstream in the production process — on condition that the B52s are delivered on time. Furthermore, US manufacturers are reluctant to work with American academics, who tend to see every little intellectual property right as a money-maker.

There was a European program called LAURA. I’ve forgotten what this acronym meant, but acronyms are usually just an excuse for making a pretty sounding name. Thanks to LAURA, we were funded as a public research center (INRIA) to take part in drafting the HiperLAN standard. With my colleagues Pascale Minet and Paul Mühlethaler, we attended all ETSI’s sessions. At the start, the meetings were itinerant before settling down at ETSI’s headquarters in Sophia-Antipolis on the Riviera. There could have been worse places to be, especially in the winter… These headquarters are a sort of big inn, a potluck for all nationalities from the Old World. The point everyone disagreed on? The food. But a consensus list of local restaurants was in circulation. It took us a while before realizing that the restaurants in this standard were in fact the ones where servers understood English, *o tempora, o mores*… And Wi-Fi and HiperLAN in all that? I’ve wandered off the subject!

As expected, the first generation of Wi-Fi — proposed simultaneously to the IEEE and ETSI — had to pick itself up off the floor after the session for simulations during the spring of 1994. It came in far behind the comb method. IEEE’s proposal was modified in an emergency. For short messages, the time before retransmission would be counted in small empty channels and suspended in case of a signal, a sort of comb with packets as teeth. For longer packets, the foreseen procedure was a sort of comb with two teeth: the sender announces the packet (“request to send”), the receptor acknowledges the announcement (“clear to send”), and the other devices shut up for the necessary period of time. This pretty much amounts to standard Wi-Fi, which has not changed since the 1990s. HiperLAN had a comb with one tooth but of variable size. So, we find the comb in both standards, a boost for the morale.

Over the years, the stakeholders in radio technology have evolved considerably, and the family of standards (IEEE 802.11, 802.11b, 802.11a, 802.11g and 802.11p) has grown along with constantly rising speeds (up to 54 Mbps for short ranges). At the protocol level however, the only significant change in the Wi-Fi standard has been the algorithms used for security. The engineers who designed Wi-Fi knew nothing about security. They designed WEP, an encryption algorithm with the vexatious habit of sending its key for anyone with enough patience to read. Since then, WAP, which is bulletproof, and its variants have replaced the WEP protocol. Nor did those of us who designed HiperLAN know much about security. We were about to design something idiotic but were stopped in our tracks by ETSI’s security experts (SAG), who concocted a protocol of a WAP sort. Cooperation among technicians worked better at ETSI than at the IEEE.

So, HiperLAN had more than one ace up its sleeve. To make up for short distances for radio transmission, plans were even made for equipping networks with internal routers to reach deaf spots. As odd as this might seem, this arrangement had been foreseen in the GSM standard but was not implemented because of the overbearing asymmetry between the technologies used. Internal routing (intraforwarding) would crop up later, once again, in the Internet and *ad hoc* mobile networks under the name OLSR, but optimized link state routing is another tale I shall not wander off to telling.

The moral of this story? Although the HiperLAN standard was published in 1995 (and ratified in 1996) at the same time as the Wi-Fi standard, no product came out that used it. WI-FI, improved thanks to HiperLAN, ultimately won out and was incorporated in European standards (sometimes called HiperLAN 2). Why? Lack of European solidarity? No, probably and quite simply: realism. At the time, the market for local wireless networks was small. Such networks did not need a license and were of little interest to players with "deep pockets". So there was no money to develop HiperLAN, a Concorde for wireless networks. The market was oriented toward less ambitious, less expensive products. Although HiperLAN did not repeat GSM's success from a business viewpoint, we can conclude that, from a technological viewpoint, HiperLAN survived in the Wi-Fi standard and was probably one of the reasons for its success.

Bibliography

IEEE 802.11 Working Group (1999) "Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Higher-speed physical layer extension in the 2.4 GHz band", *ANSI/IEEE Std 802.11*.

ETS 300 652 Radio Equipment en System (RES) High Performance Radio Local Area Network (HIPERLAN); Type 1 Functional Specification, 1995.

JACQUET P., MINET P., MÜHLETHALER P. & RIVIERRE N. (1997) "Increasing reliability in cable-free radio LANs low level forwarding in HIPERLAN", *Wireless Personal Communications*, *4*(1), pp. 51-63.

JACQUET P., MINET P., MÜHLETHALER P. & RIVIERRE N. (1997) "Priority and collision detection with active signaling: The channel access mechanism of HiperLAN", *Wireless Personal Communications*, *4*(1), pp. 11-25.

JACQUET P., MINET P., MÜHLETHALER P. & RIVIERRE N. (1997) "Data transfer for HiperLAN", *Wireless Personal Communications*, *4*(1), pp. 65-80.

JACQUET P., MÜHLETHALER P., CLAUSEN T., LAOUITI A., QAYYUM A. & VIENNOT L. (2001) "Optimized link state routing protocol for *ad hoc* networks" in *Multi Topic Conference, 2001. IEEE INMIC 2001. Technology for the 21st Century. Proceedings. IEEE International*, pp. 62-68.