

< INC-PROJECT, MAP-CRITIQUE.NLS.10, >, 12-Aug-83 11:46 AMW ;;;;

RFC 874

September 1982
M82-50

A CRITIQUE OF X.25

M.A. PADLIPSKY
THE MITRE CORPORATION
Bedford, Massachusetts

ABSTRACT

The widely touted network interface protocol, "X.25", and its attendant conceptual framework, the International Standards Organization's Reference Model for Open System Interconnection (ISORM), are analyzed and found wanting. The paper is a companion piece to M82-48, and M82-51.

A CRITIQUE OF X.25

M. A. Padlipsky

Introduction

According to some sources, the International Standards Organization's (ISO) "Open System Interconnection" (OSI) effort has adopted the International Consultative Committee on Telephony and Telegraphy (CCITT) developed X.25 protocol(s) as its Levels 1-3. ("Loose constructionists" of the ISORM would hold that X.25 is a mechanization of L1-L3 rather than the mechanization, and at least one British source holds that "we in the U.K. don't believe that ISO have adopted X.25.") In the U.S. Government arena, where the author spends much of his time, the Government Accounting Office (GAO) has suggested that the Department of Defense (DoD) ought to consider adopting "X.25 networks," apparently in preference to networks based on protocols developed by the DoD-sponsored intercomputer networking research community. That intercomputer networking research community in turn has, with a few recent exceptions, adhered to its commitment to the Oral Tradition and not taken up the cudgels against X.25 in the open literature, even though X.25 is an object of considerable scorn in personal communications.

Although the DoD Protocol Standards Technical Panel has begun to evolve a "Reference Model" different from ISO's for reasons which will be touched on below, there seems to be a need to address the deficiencies of X.25 on their own demerits as soon as possible. Without pretending to completeness*, this paper will attempt to do just that.

The overall intent is to deal with X.25 in the abstract; because of who pays the bills, though, a necessary preliminary is to at least sketch the broad reasons why the DoD in particular should not

* Various versions of X.25 and ISO documentation were employed; one incompleteness of note, however, is that no attempt has been made to do proper bibliographic citation. Another incompleteness lies in the area of "tutoriality"; that is, appropriate prior knowledge is assumed on the part of the reader. (The author apologizes for the omissions but hasn't the time or the energy to be overly scholarly. Reference [3] might be of use to the reader who feels slighted.)

employ intercomputer networks which base their protocol suites on the ISO Reference Model (ISORM) with X.25 as Levels 1-3. (Note that this is a different formulation from "use communications subnetworks which present an X.25 interface.") Very briefly, the DoD has concerns with "survivability," reliability, security, investment in prior art (i.e., its research community has a working protocol suite in place on many different operating systems), procurability (i.e., ISORM-related protocol suites do not as yet fully exist even on paper and the international standardization process is acknowledged even by its advocates to require several years to arrive at full suite specification, much less offer available interoperable implementation), and interoperability with a much wider range of systems than are ever likely to receive vendor-supplied implementations of ISORM protocol suites. Regardless of which particular concerns are considered to dominate, the DoD cannot be expected to await events in the ISO arena. (Particularly striking is the fact that DoD representatives are not even permitted under current doctrine to present their specific concerns in the area of security in the sort of unclassified environment the ISO arena constitutes.)

Some zealous ISORM advocates have suggested that the DoD research community suffers from a "Not Invented Here" syndrome with respect to ISORM-related protocols, though, so even if the various reasons just cited were to prevail, there would still be an open issue at some level. At least one or two zealous members of the research community have asserted that the problem is not Not Invented Here, but Not Invented Right, so an assessment of the apparent keystone of the ISORM suite, X.25, from the perspective of whether it's "good art" ought to be appropriate. That's what we're up to here.

Problems With the Conceptual Model*

There is confusion even amongst its advocates as to the real conceptual model of X.25-based ISO networking. Some draw their Reference Model as two "highrises," others draw "parking garages" beside each highrise. That is, some draw the seven ISORM layers in large rectangles (representing Hosts) next to one another, showing each layer in communication with its "peer" in the other Host/Open System; this implies an "end-to-end" view of X.25. Others draw smaller rectangles between the larger ones, with Levels 1-3 having peer relationships from the Host-OS ("Data Terminal Equipment") to the Comm Subnet Node ("Data Circuit Terminating Equipment"); this implies a "link-by-link" view of X.25. This ambiguity does not engender confidence in the architects, but perhaps the real problem is with the spectators. Yet it is indisputable that when internetting with X.75, the model becomes "hop-by-hop" (and it is likely it's meant to be link-by-link even on a single comm subnet).

A major problem with such a model is that the designers have chosen to construe it as requiring them to break the "virtual circuit" it is supposed to be supporting whenever there is difficulty with any one of the links. Thus, if internetting, and on some interpretations even on one's proximate net, rerouting of messages will not occur when needed, and all the upper levels of protocols will have to expend space-time resources on reconstituting their own connections with their counterparts. (Note that the success of the reconstitution under DCE failure appears to assume a certain flexibility in routing which is not guaranteed by the Model.) This can scarcely be deemed sound design practice for an intercomputer networking environment, although many have conjectured that it probably makes sense to telephonists.

* Note that we are assuming an ISO-oriented model rather than a CCITT-oriented one (X.25/X.28/X.29) because the latter appears to offer only "remote access" functionality whereas the sort of intercomputer networking we are interested in is concerned with the full "resource-sharing" functionality the former is striving for. This might be somewhat unfair to X.25, in that it is taking the protocol(s) somewhat out of context; however, it is what ISO has done before us, and if what we're really accomplishing is a demonstration that ISO has erred in so doing, so be it. As a matter of fact, it can also be argued that X.25 is itself somewhat unfair--to its users, who are expecting real networking and getting only communication; cf. Padlipsky, M. A., "The Elements of Networking Style", M81-41, The MITRE Corporation, October 1981, for more on the extremely important topic of resource sharing vs. remote access.

Indeed, it appears the virtual circuit metaphor is in some sense being taken almost literally (with the emphasis on the "circuit" aspect), in that what should be an environment that confers the benefits of packet-switching is, at the X.25 level, reduced to one with the limitations of circuit-switching. On the other hand, the metaphor is not being taken literally enough in some other sense (with the emphasis on the "virtual" aspect), for many construe it to imply that the logical connection it represents is "only as strong as a wire." Whether the whole problem stems from the desire to "save bits" by not making addresses explicitly available on a per-transmission basis is conjectural, but if such be the case it is also unfortunate.

(As an aside, it should be noted that there is some evidence that bit saving reaches fetish--if not pathological--proportions in X.25: For instance, there does not even appear to be a Packet Type field in data packets; rather--as best we can determine--for data packets the low order bit of the "P(R)" field, which overlaps/stands in the place of the Packet Type is always 0, whereas in "real" Packet Type fields it's always 1. [That may, by the way, not even be the way they do it--it's hard to tell ... or care.]

There is also confusion even amongst its advocates as to what implications, if any, the protocol(s) has (have) for comm subnet node to comm subnet node (CSN) processing. Those who draw just two highrises seem to be implying that from their perspective the CSN (or "DCE") is invisible. This might make a certain amount of sense if they did not assert that each floor of a highrise has a "peer-relationship" with the corresponding floor of the other highrise--for to do so implies excessively long wires, well beyond the state of the wire-drawing art, when one notices that the first floor is the physical level. (It also appears to disallow the existence of concatenated comm subnets into an internet, or "catenet," unless the CSN's are all identically constituted. And those who hold that the ISORM dictates single protocols at each level will have a hard time making an HDLC interface into a Packet Radio Net, in all probability.)

Those who, on the other hand, "draw parking garages," seem to be dictating that the internal structure of the CSN also adhere to X.25 link and physical protocols. This implies that Packet Radio or satellite CSNs, for example, cannot "be X.25." Now that might be heartening news to the designers of such comm subnets, but it presumably wasn't intended by those who claim universality for X.25--or even for the ISO Reference Model.

Even granting that ambiguities in the conceptual model do not constitute prima facie grounds for rejecting the protocol(s), it is important to note that they almost assuredly will lead to vendor implementations based on differing interpretations that will not interoperate properly. And the unambiguous position that virtual circuits are broken whenever X.25 says so constitutes a flaw at least as grave as any of the ambiguities.

Another, in our view extremely severe, shortcoming of the X.25 conceptual model is that it fails to address how programs that interpret its protocol(s) are to be integrated into their containing operating systems. (This goes beyond the shortcoming of the X.25 specifications in this area, for even the advocates of the ISORM--who, by hypothesis at least, have adopted X.25 for their Levels 1-3--are reticent on the topic in their literature.) Yet, if higher level protocols are to be based on X.25, there must be commonality of integration of X.25 modules with operating systems at least in certain aspects. The most important example that comes to mind is the necessity for "out-of-band signals" to take place. Yet if there is no awareness of that sort of use reflected in the X.25 protocol's specification, implementers need not insert X.25 modules into their operating systems in such a fashion as to let the higher level protocols function properly when/if an X.25 Interrupt packet arrives.

Yet much of the problem with the conceptual model might turn out to stem from our own misunderstandings, or the misunderstandings of others. After all, it's not easy to infer a philosophy from a specification. (Nor, when it comes to recognizing data packets, is it easy even to infer the specification--but it might well say something somewhere on that particular point which we simply overlooked in our desire to get the spec back on the shelf rapidly.) What other aspects of X.25 appear to be "bad art"?

"Personality Problems"

When viewed from a functionality perspective, X.25 appears to be rather schizophrenic, in the sense that sometimes it presents a deceptively end-to-end "personality" (indeed, there are many who think it is usable as an integral Host-Host, or Transport, and network interface protocol, despite the fact that its specification itself--at least in the CCITT "Fascicle" version--points out several functional omissions where a higher-level protocol is expected--and we have even spoken to one or two people who say they actually do -- use it as an end-to-end protocol, regardless); sometimes it presents a comm subnet network interface personality (which all would agree it must); and sometimes (according to some observers) it presents a

"Host-Front End Protocol" personality. Not to push the "bad art" metaphor too hard, but this sort of violation of "the Unities" is, if demonstrable, grounds for censure not only to literary critics but also to those who believe in Layering. Let's look at the evidence for the split-personality claim:

X.25 is not (and should not be) an "end-to-end" protocol in the sense of a Transport or Host-to-Host protocol. Yet it has several end-to-end features. These add to the space-time expense of implementation (i.e., consume "core" and CPU cycles) and reflect badly on the skill of its designers if one believes in the design principles of Layering and Least Mechanism. (Examples of end-to-end mechanisms are cited below, as mechanisms superfluous to the network interface role.) The absence of a datagram mode which is both required and "proper" (e.g., not Flow Controlled, not Delivery Confirmed, not Non-delivery mechanized) may also be taken as evidence that the end-to-end view is very strong in X.25. That is, in ISO Reference Model (ISORM) terms, even though X.25 "is" L1-3, it has delusions of L4-ness; in ARPANET Reference Model (ARM) terms, even though X.25 could "be" L I, it has delusions of L II-ness.*

X.25 is at least meant to specify an interface between a Host (or "DTE") and a comm subnet processor (or "DCE"), regardless of the ambiguity of the conceptual model about whether it constrains the CSNP "on the network side." (Aside: that ambiguity probably reflects even more badly on certain X.25 advocates than it does on the designers, for there is a strong sense in which "of course it can't" is the only appropriate answer to the question of whether it is meant to constrain generic CSN processors (CSNP's) in the general case. Note, though, that it might well be meant to constrain specific DCE's; that is, it started life as a protocol for PTT's--or Postal, Telephone, and Telegraph monopolies--and they are presumably entitled to constrain themselves all they want.) Yet the end-to-end features alluded to above are redundant to the interfacing role, and, as noted, extraneous features have space-time consequences. There are also several features which, though not end-to-end, seem superfluous to a "tight" interface protocol. Further, the reluctance of the designers to incorporate a proper "datagram" capability in the protocol (what they've got doesn't seem to be

* For more on the ARM, see Padlipsky, M. A., "A Perspective on the ARPANET Reference Model", M82-47, The MITRE Corporation, September 1982; also available in Proc. INFOCOM '83. (Some light may also be cast by the paper on the earlier-mentioned topic of Who Invented What.)

usable as a "pure"--i.e., uncontrolled at L3 but usable without superfluous overheard by L4--datagram, but instead entails delivery confirmation traffic like it or not; note that "seem" is used advisedly: as usual, it's not easy to interpret the Fascicle) suggests at least that they were confused about what higher-level protocols need from interfaces to CSNP's, and at worst that there is some merit to the suggestion that, to paraphrase Louis Pouzin, "the PTT's are just trying to drum up more business for themselves by forcing you to take more service than you need."

Examples of mechanisms superfluous to the interface role:

1. The presence of a DTE-DTE Flow Control mechanism.
2. The presence of an "interrupt procedure" involving the remote DTE.
3. The presence of "Call user data" as an end-to-end item (i.e., as "more" than IP's Protocol field).
4. The "D bit" (unless construed strictly as a "RFNM" from the remote DCE).
5. The "Q bit" (which we find nearly incomprehensible, but which is stated to have meaning of some sort to X.29--i.e., to at least violate Layering by having a higher-level protocol depend on a lower level mechanism--and hence can't be strictly a network interface mechanism).

The final "personality problem" of X.25 is that some of its advocates claim it can and should be used as if it were a Host-Front End protocol.* Yet if such use were intended, surely its designers would have offered a means of differentiating between control information destined for the outboard implementation of the relevant protocols and data to be transmitted through X.25, but there is no evidence of such mechanisms in the protocol. "Borrowing" a Packet Type id for H-FP would be risky, as the spec is subject to arbitrary alteration. Using some fictitious DTE address to indicate the proximate DCE is also risky, for the same reason. Further, using "Call user data" to "talk to" the counterpart H-FP module allows only 15 octets (plus, presumably, the 6 spare bits in the 16th octet) for the conversation, whereas various TCP and IP options might require many more octets than that. Granted that with sufficient ingenuity--or even by the simple expedient of conveying the entire H-FP as data (i.e., using X.25 only to get channels to demultiplex on, and DTE-DCE flow control, with the "DCE" actually being an Outboard Processing Environment that gets its commands in the data fields of X.25 data packets)--X.25 might be used to "get at" outboard protocol interpreters, but its failure to address the issue explicitly again reflects badly on its designers' grasp of intercomputer networking issues. (Another possibility is that the whole H-FP notion stems from the use of X.25 as a Host-Host

* That is, as a distributed processing mechanism which allows Host operating systems to be relieved of the burden of interpreting higher level protocols "inboard" of themselves by virtue of allowing Host processes to manipulate "outboard" interpreters of the protocols on their behalf. Note that the outboarding may be to a separate Front-End processor or to the CSNP itself. (The latter is likely to be found in microprocessor-based LAN "BIU's.") Note also that when dealing with "process-level" protocols (ARM L III; approximately ISORM L5-7), only part of the functionality is outboarded (e.g., there must be some Host-resident code to interface with the native File System for a File Transfer Protocol) and even when outboarding Host-Host protocols (ARM L II; approximately ISORM L4 plus some of 5) the association of logical connections (or "sockets") with processes must be performed inboard--which is why, by the way, it's annoying to find ISO L5 below ISO L6: because, that is, you'd like to outboard "Presentation" functionality but its protocol expects to interact with the "Session" protocol, the functionality of which can't be outboarded. (Although this approach, not the proper context for a full treatment of the H-FP approach, it is also of interest that the approach can effectively insulate the Host from changes in the protocol suite, which can be a major advantage in some environments.)

protocol so that some might think of it in its Host aspect as "simply" a way of getting at the H-HP. This interpretation does give rise to the interesting observation that DCE's seem to need a protocol as strong as TCP amongst themselves, but doesn't strike the author as particularly convincing evidence for viewing X.25 as anything like a proper H-FP--if for no other reason than that a central premise of Outboard Processing is that the Host-side H-FP module must be compact relative to an inboard generic Network Control Program.)

X.25, then, is rather schizophrenic: It exceeds its brief as an interface protocol by pretending to be end-to-end (Host-Host) in some respects; it is by no means a full end-to-end protocol (its spec very properly insists on that point on several occasions); it's at once too full and too shallow to be a good interface; and it's poorly structured to be treated as if it were "just" an H-FP. (Some would phrase the foregoing as "It's extremely ill layered"; we wouldn't argue.)

A Note on "Gateways"*

Although it was at least implied in the discussion of conceptual model problems, one aspect of X.25/X.75 internetting is sufficiently significant to deserve a section of its own: Not only does the link-by-link approach taken by CCITT make it unlikely that alternate routing can take place, but it is also the case that ARPANET Internet Protocol (IP) based internetting not only permits alternate routing but also could alt-route over an "X.25 Subnet." That is, in IP's conceptual model, Gateways attach to two or more comm subnets "as if they (the Gateways) were Hosts." This means that they interpret the appropriate Host-comm subnet processor protocol of whatever comm subnets they're attached to, giving as the "proximate net address" of a given transmission either the ultimate (internet addressed) destination or the address of another Gateway "in the right direction." And an implementation of IP can certainly employ an implementation of ("DTE") X.25 to get a proximate net, so ... at least "in an emergency" X.25 interface presenting Public Data Networks can indeed carry IP traffic. (Note also that only the proximate net's header has to be readable by the nodal processor of/on the proximate net, so if some appropriate steps were taken to render the data portion of such transmissions unintelligible to the nodal processors, so much the better.)

* This section was added to address the ill-founded concerns of several ISORMites that "TCP/IP won't let you use Public Data Nets in emergencies."

(Further evidence that X.75 internetting is undesirable is found in the fact that the U.S. National Bureau of Standards has, despite its nominal adoption of the ISORM, inserted IP at approximately L3.5 in its version of the Reference Model.)

The Off-Blue Blanket

Although touched on earlier, and not treatable at much length in the present context, the topic of security deserves separate mention. We are familiar with one reference in the open literature [1] which appears to make a rather striking point about the utility of X.25 in a secure network. Dr. Kent's point that the very field sizes of X.25 are not acceptable from the point of view of encryption devices would, if correct (and we are neither competent to assess that, nor in a position to even if we were), almost disqualify X.25 a priori for use in many arenas. Clearly, uncertified "DCE's" cannot be permitted to read classified (or even "private") data and so must be "encrypted around," after all.

It would probably be the case, if we understand Dr. Kent's point, that X.25 could be changed appropriately--if its specifiers were willing to go along. But this is only one problem out of a potentially large number of problems, and, returning briefly to our concern with the interplay of X.25 and the DoD, those persons in the DoD who know best what the problems are and/or could be are debarred from discussing them with the specifiers of X.25. Perhaps a sufficiently zealous ISORM advocate would be willing to suggest that Professor Kuo's publisher be subsidized to come out with a new edition whenever a problem arises so that if Dr. Kent happens to spot it advantage can continue to be taken of his ability to write for the open literature--but we certainly hope and trust that no ISORMite would be so tone-deaf as to fail to recognize the facetiousness of that suggestion.

In short, it appears to be difficult to dispute the assertion that whatever sort of security blanket X.25 could represent would at best be an off shade of blue.

Space-Time Considerations

Another topic touched on earlier which deserves separate mention, if only to collect the scattered data in a single section, is that of what have been called space-time considerations. That is, we are concerned about how well X.25 in particular and the ISORM-derived protocols in general will implement, both in terms of size of protocol interpreters (PI's) and in terms of execution and delay times.

On the space heading, certainly the fact that X.25 offers more functionality in its end-to-end guise than is required to fulfill its network interface role suggests that X.25 PI's will be bigger than they need be. As an aside--but a striking one--it should be noted that X.25's end-to-end functions are at variance with the ISORM itself, for the "peer entity" of a DTE X.25 entity must surely be the local DCE X.25. Perhaps a later version of the ISORM will introduce the polypeer and give rise to a whole new round of Layering-Theologic controversy.* Speaking of the ISORM itself, those who hold that each layer must be traversed on each transmission are implicitly requiring that space (and time) be expended in the Session and Presentation Levels even for applications that have no need of their services. The Well-Known Socket concept of the ARM's primary Host-Host protocol, the Transmission Control Protocol (TCP), lets Session functionality be avoided for many applications, on the other hand--unless ISORM L5 is to usurp the Host's user identification/authentication role at some point. (Yes, we've heard the rumors that "null layers" might be introduced into the ISORM; no, we don't want to get into the theology of that either.)

On the time heading, X.25's virtual circuit view can be debilitating--or even crippling--to applications such as Packetized Speech where prompt delivery is preferred over ordered or even reliable delivery. (Some hold that the X.25 datagram option will remedy that; others hold that it's not "really datagrams"; we note the concern, agree with the others, and pass on.) Speaking of reliable delivery, as noted earlier some observers hold that in order to present an acceptable virtual circuit X.25 must have a protocol as strong as TCP "beneath" itself; again, we're in sympathy with them. Shifting focus again to the ISORM itself, it must be noted that the principle that "N-entities" must communicate with one another even in the same Host via "N-1 entities" even in the same Host is an over-zealous application of the Principle of Layering that must consume more time in the interpreting of the N-1 protocol than would a direct interface between N-level PI's or such process-level protocols as FTP and Telnet, as is done in the ARPANET-derived model.

Other space-time deficiencies could be adduced, but perhaps a shortcut will suffice. There is a Law of Programming (attributed to Sutherland) to the effect that "Programs are like waffles: you should always throw the first one out." Its relevance should become

* And perhaps we now know why some just draw the highrises.

clear when it is realized that (with the possible exception of X.25) ISORM PI's are in general either first implementations or not even implemented yet (thus, the batter, as it were, is still being mixed). Contrast this with the iterations the ARPANET-derived PI's--and, for that matter, protocols--have gone through over the years and the grounds for our concern over X.25/ISORM space-time inefficiency become clear irrespective of corroborative detail. Factor in the consideration that space-time efficiency may be viewed as contrary to the corporate interests of the progenitors of X.25 ("the PTT's") and at least the current favorite for ISORM Level 4 (ECMA--the European Computer Manufacturers' Association), and it should become clear why we insist that space-time considerations be given separate mention even though touched upon elsewhere.*

Getting Physical

Still another area of concern over X.25 is that it dictates only one means of attaching a "DTE" to a "DCE." That is, earlier references to "the X.25 protocol(s)" were not typographical errors. Most of the time, "X.25" refers to ISORM Level 3; actually, though, the term subsumes L2 and L1 as well. Indeed, the lowest levels constitute particular bit serial interfaces. This is all very well for interfacing to "Public Data Nets" (again, it must be recalled that X.25's roots are in CCITT), but is scarcely appropriate to environments where the communications subnetwork may consist of geosynchronous communications satellite channels, "Packet Radios," or whatever. Indeed, even for conventional Local Area Networks it is often the case that a Direct Memory Access arrangement is desired so as to avoid bottlenecks--but DMA isn't HDLC, and the "vendor supported X.25 interface" so prized by some won't be DMA either, one imagines. (Speaking of LAN's, at least the evolving standard in that arena--"IEEE 802"--apparently will offer multiple physical interfaces depending on comm subnet style [although there is some disagreement on this point amongst readers of their draft specs]; we understand, however, that their Level 2 shares X.25's end-end aspirations--and we haven't checked up on DMA capability.) X.25, then, imposes constraints upon its users with regard to interface technology that are inappropriate.

* The broad issue of design team composition is amplified in Padlipsky, M. A., "The Illusion of Vendor Support", M82-49, The MITRE Corporation, September 1982.

Other Observers' Concerns

This paper owes much to conversations with a number of people, although the interpretations of their concerns are the author's responsibility. Mention should be made, however, of a few recent documents in the area: The Defense Communications Agency (DCA Code J110) has sent a coordinated DoD position [2] to NBS holding that X.25 cannot be the DoD's sole network interface standard; Dr. Vinton Cerf of the ARPA Information Processing Technology Office made a contribution to the former which contains a particularly lucid exposition of the desirability of proper "datagram" capability in DoD comm subnets [3]; Mr. Ray McFarland of the DoD Computer Security Evaluation Center has also explored the limitations of X.25 [4]. Whether because these authors are inherently more tactful than the present author, or whether their positions are more constraining, or even whether they have been more insulated from and hence less provoked by uninformed ISORMite zealots, none has seen fit to address the "quality" of X.25. That this paper chooses to do so may be attributed to any one of a number of reasons, but the author believes the key reason is contained in the following:

Conclusion

X.25 is not a good thing.

References

- [1] Kent, S. T., "Security in Computer Networks," in Kuo, F., Ed., Protocols and Techniques for Data Communications Networks, Prentice-Hall, 1981, pp. 369-432.
- [2] Letter to NBS from P. S. Selvaggi, Chief, Interoperability and Standards Office, 7 April 1982.
- [3] Cerf, V. G., "Draft DoD Position Regarding X.25" in undated letter to P. S. Selvaggi.
- [4] Personal communications.