Compatibility Standards and the Internet

William Lehr

Columbia Institute for Tele-Information
Graduate School of Business
809 Uris Hall
Columbia University
New York, New York 10027
(212) 854-4222

596

Compatibility Standards and the Internet

Paper presented to The Twentieth Annual Telecommunications Policy Research Conference, Solomons Island, MD., September 12-14, 1992

William Lehr
723 Uris Hall
Graduate School of Business
Columbia University
New York, NY 10027

email: wlehr@research.gsb.columbia.edu
phone: 212-854-4426

This is a preliminary draft. Comments are encouraged and welcome, but please do not cite or distribute without contacting the author.

Compatibility Standards and the Internet

Abstract

Economists have started to examine how different institutional arrangements affect the setting of voluntary compatibility standards. Many participants in the computer and telecommunications industries have complained that traditional Standards Development Organizations (SDOs) are too slow, and have pointed to the approach followed by the Internet community as more effective. This paper critically examines the procedures used in the Internet community to determine how it differs from more traditional SDOs such as ANSI, the CCITT or AUSSIE, and to explain why the Internet approach may have been perceived as more effective. I attribute most of the Internet's prior success to special factors of the decision-making environment which are currently changing and requiring the Internet to become more like traditional SDOs. One important difference which merits special attention is the Internet's approach to anticipatory standards. A game-theoretic model of anticipatory standard setting is introduced to better understand the general problem of standardization in the face of uncertainty, when future decisions may be partially or fully reversible.

Introduction

Economists have started to examine how differences in standardization processes might influence the setting of voluntary compatibility standards[1]. Numerous participants in the computer and telecommunications industries have complained that the formal de jure approach of traditional Standards Development Organizations (SDOs) such as ISO, the CCITT and ANSI is too slow[2]. Some have argued that less formal bodies such as the Internet Engineering Task Force (IETF) are more effective[3].

This paper explains how the IETF's approach differs from other SDOs and why it may have been perceived to be more effective. My results may be summarized by two principal conclusions. First, most of the Internet's past success appears due to special environmental circumstances which are currently changing. As the Internet has evolved from a relatively small community of researchers and academics into a much larger, more heterogeneous community with more important commercial implications, standard setting has become more contentious. Consequently, the IETF's approach will need to become more structured and formal -- more like the traditional SDOs[4]. These trends are already occurring. An unfortunate side-effect is that standards are likely to take longer and will be more expensive to develop.

---

1 See for example, Besen (1990), Besen and Farrell (1991), Lehr (1992a,1992b,1992c) or Farrell and Shapiro (1992). For a good review of the economics literature on standardization see David and Greenstein (1990).

2 See Vasquez (1990) or Moad (1990). All of these organizations are involved in the development of technical compatibility standards for the computer and telecommunications industries. In structure and procedure these organizations are all more similar to each other than to the IETF. ISO is the International Standards Organization; the CCITT is the International Telegraph and Telephone Consultative Committee; ANSI is the American National Standards Institute. See Cargill (1989) for a description of these institutions or Lehr (1992a) for analysis of structure.

3 See McQuillan (1990) for claim IETF process works better.

4 This prediction is based on my earlier attempt to interpret the structure of traditional SDOs as a rational response to the problem of standardization (Lehr, 1992a).

Second, and more interesting for its broader theoretical implications, the Internet community approaches anticipatory standards differently than other SDOs. Anticipatory standards address technologies which are still emerging and hence involve great collective and private uncertainty. We know that what we do today will require further actions and decisions later. In the Internet community, there is a strong bias in favor of what I shall call "incremental standardization". When faced with a need to develop standards for technologies which are still emerging, they proceed step-by-step. In contrast, ISO (most notably in the OSI effort) is willing to attempt true "anticipatory standards", wherein the committee trys to anticipate future developments[5]. We may think of the former as sequential short-term contracting; whereas the latter approach requires long-term contracting.

In both cases, standardization is partial, since the final outcome depends on the resolution of uncertainty. In ISO, they appear more likely to include multiple options, leaving it to subsequent standards bodies to specify which options are required to guarantee compatibility (e.g., conformance tests and standards implementation profiles). In the Internet, they are more likely to leave out features about which there is great uncertainty or limited consensus (which are often related events). The IETF practices a form of standard setting by "lowest common denominator".

The main body of the paper is divided into two parts, wherein I address each of my

5 OSI stands for Open Systems Interconnection, which is most often asociated with the 7-layered OSI Reference Model for protocol layering which was first proposed by ISO in 1979. See Tannenbaum (1988) or Rose (1990) for a discussion of OSI.

conclusions in turn. In the second part, I introduce a preliminary version of a game theoretic model with which to analyze anticipatory standard setting. I hope to refine this model and use it to generate general propositions in future research, and apologise in advance, if it seems incomplete in its current form.

Part I: Standard Setting in the Internet

i. What is the Internet and how is it distinct

The Internet is a network of autonomous data networks which electronically links host computers at universities, research centers, government offices and corporations around the globe. Today's Internet evolved from the ARPANET, which was created in 1969 under the aegis of the U.S. Department of Defense as the first wide-area packet-switched data network linking a few prominent research universities and defense contractors within the United States[6]. The Internet permits users on its sub-networks to transfer files, exchange electronic mail ("email") and access remote host computers. These functions require a collection of networking protocol standards which guarantee that communications between hosts on disparate networks will be compatible. During the 1970's, the family of TCP/IP protocols were developed for use in the

6 See Cook (1992) or Kahin (1992) for histories of the Internet. The ARPANET was finally phased out in 1988 and its traffic was shifted to NSFnet, the network funded by the National Science Foundation (NSF) to provide backbone services to the Internet.

Internet[7].

The continuing evolution of the TCP/IP suite of protocols is managed by the Internet Engineering Task Force (IETF), which is one of the two task forces under the supervision of the Internet Activities Board (IAB)[8]. The IETF is organized into collections of working groups in 9 functional areas, each of which is under the supervision of an area director. Each area director supervises the work of from 5 to 9 working groups which are actively engaged in developing new or revising existing protocol standards. The IETF is itself managed by the Internet Engineering Steering Group (IESG), consisting of the IESG Chairman and the nine area directors. The members of the IESG, IAB and the chairpersons and participants in the working groups are all volunteers, who participate as individuals (rather than as representatives of their employers).

The actual standards are developed in the IETF working groups, where anyone who is interested may participate, although typically, draft documents are produced by a few dedicated engineers. Each working group must have a charter approved by the IESG and IAB. These charters limit the technical scope of a working group and working groups are usually disbanded

7 TCP/IP stands for Transport Control Protocol and Internet Protocol. See Rose (1990, 1991) or Tannenbaum (1988) for an introduction to networking protocols. Vinton Cerf and Bob Kahn co-authored the first paper describing the Internet architecture in 1974. The TCP/IP protocols were formally adopted for use in the ARPANET in 1983.

8 The IAB was formed in 1983 from the re-organized Internet Configuration Control Board (ICCB), which was originally formed in 1979 by Vinton Cerf, who at the time was the Defense Advanced Research Projects Agency (DARPA) program director. The IETF was formed in January 1986. The other important task force is the Internet Research Task Force (IRTF) which supervises experimental networking research in the Internet. See Cerf (1989).

once their chartered task is either completed or abandoned[9].

In addition to working group meetings and IETF plenary meetings[10], where all of the working groups gather to discuss their progress, draft standards are debated on electronic bulletin boards. To further facilitate communication within the Internet community, the IAB maintains the "Request for Comments" (RFCs) document series. The RFCs report important IAB positions, official standards documents and statements of procedures[11]. The~RFCs, meeting minutes for the IAB, IESG and the various working groups and all of the working drafts are electronically accessible at no cost via the Internet[12].

The ability to communicate electronically dramatically lowers the costs of remotely tracking and monitoring the activities of the IETF. This distinguishes the IETF from more traditional SDOs like ISO or ANSI. Other SDOs usually sell standards and procedural documentation, and seldom, if ever, use email and electronic file access to facilitate geographically dispersed communication. Although the cost of standards documentation from traditional SDOs may be unlikely to deter those who are directly involved in implementation of the standards, these costs make it more difficult for peripherally interested parties to keep

9 The ability to constrain the scope and lives of the working groups adds to the power of the IESG and IAB. In other SDOs, working groups are often longer lived.

10 IETF Plenaries ocur three times a year.

11 See Chapin's (1992) RFC1310 or Postel's (1992) RFC1280 for descriptions of the Internet standard setting process.

12 Users may access Internet documentation via FTP to the host NIC.DDN.MIL, using the account ANONYMOUS with password GUEST.

informed. The IETF is more "open" than traditional SDOs in the sense that it is much easier for outsiders to keep abreast of the IETF's activities and to contribute comments[13].

When it comes to real decision-making power, however, the IETF process appears less open than the more formal SDOs. First, the IETF does not have formal voting rules. Voting rules are a mechanism for establishing property rights to decision-making power. Other SDOs usually require supramajority voting before major decisions are approved. Also, participants usually vote by firm or organization to prevent vote-packing[14].

Although the IETF claims to implement "consensus decision-making", the process is far less formally constrained than in other SDOs. For example, in the ANSI Accredited X3 Committee for Information Processing Systems (X3), standards must progress through a series of 18 "mile stones". Supramajority voting (usually 2/3rds) is used for most decisions and there must be a documented attempt to reconcile all negative comments. In practice, ANSI approval procedures require multiple ballots and time-consuming comment and review cycles before achievement of consensus can be demonstrated.

In the IETF, working group chairs have much more discretion in how they manage the

13 During the course of this research, I have reviewed the email archives from several working groups and have seen many comments from obviously technically unsophisticated would-be participants.

14 The SDOs differ by voting rules. In the IEEE, participants vote as individuals; whereas in the ANSI X3 Committee, ISO, CCITT, ETSI votes are by organization or firm. In the CCITT, votes are by country. Besen (1990) discusses why ETSI decided to relax its supramajority voting requirements to make it easier to pass decisions. Although vote-packing might be perceived as a threat when participants vote as individuals, this does not seem to have been a major problem for the IEEE. See Lehr (1992a, 1992c) for further discussion of the procedures employed by other SDOs.

development process[15]. The working group submits standards to the IESG which decides what action to recommend to the IAB. Both the IESG and IAB proceedings are closed, and both committees have been dominated by a relatively small group who have worked together for a long time. In the past, the IAB decided who its eleven members would be[16]. They use their discretion to determine whether consensus has been reached, but have not been constrained by formal decision-making procedures.

It is probably easier for the IAB and IETF to exercise their discretion to block decisions they oppose than it is to enforce decisions which are opposed by members of the Internet community[17]. Rather than demanding that consensus be convincingly demonstrated, as under the ANSI procedures, the IETF process presumes consensus exists in the absence of vocal opposition. If it is costly to participate and if standards produce positive externalities for non-participants then there may be a free-rider problem. Assigning the burden of evidence to the process to conclusively demonstrate that consensus has been achieved makes it relatively more

---

15 In Rose (1991), Marshall Rose, who chaired the SNMP working group and is one of the more out-spoken proponents of the IETF process, described his role as follows: "The role of chair of the SNMP Working Group is little more than political arbiter. Technical work is performed by a few key members of the group, and the chair tries to provide an environment to allow this. As current chair, there are a couple of simple-minded tricks I use, For example, meetings are scheduled in locations off the beaten track, simply to reduce attendance by hangers-on and standards go-ers - - who attend to make seminal contributions and enjoy many fine lunches and dinners...." (p248). Later, Rose comments: "Thus, the job of the chair is to make the process appear to be 'open and friendly,' whilst making sure that none of the 'doggie biscuits' tossed about have any bad effects on the working aspects of the technology." (p250)

16 See Desmond (1990). The new charter for the IAB calls for a 15 member IAB, under the control of the Internet Society's Trustees.

17 It is much easier to ignore calls for action or disregard comments for additional features than it is to approve features which are opposed by a vocal opposition.

---

difficult to over-turn the status quo[18]. By shifting the burden of evidence, the IETF makes it easier to produce new standards.

The openness of the electronic communications are refreshing and make it abundantly clear when the Internet community opposes a new standard. This was demonstrated recently in a decision by the IAB regarding an important routing protocol. A number of participants felt that the IAB process left little time to debate alternatives and appeared autocratic. According to Robert Hinden, the IETF's area director for routing, the "IETF is not a democracy, but if it is to be successful, there needs to be consensus"[19]. In a recent editorial, David Buerger writes that "the IAB episode ... presages a new headache for standards-making bodies: a potentially longer review process.....the day of the benevolent dictatorship is past".[20]

ii. Why is the Internet process changing?

In the past, the decisions of the IAB attracted little controversy. Many of the important positions in the IAB, IETF and various working groups are still controlled by the same academics, network researchers and consultants who participated in the development of the

18 This is the problem of voter participation, wherein it is privately irrational for a citizen to incur the cost of voting unless she beleives her vote will be pivotal.

19 Quoted in Gibson and Girishankar (1992).

20 See Buerger (1992), p42.

original ARPANET technology[21]. All of the key participants are experts in networking technology with long working relationships. In a small, relatively homogeneous community, an informal decision-making process can work quite well. It affords the flexibility to adjust to changing circumstances and the common level of expertise reduces problems associated with asymmetric information. Long-term relationships (a la Axelrod (1984)) reduce the threat from short-term opportunistic behavior and create opportunities for inter-temporal side-payments (i.e., I'll yield today and you'll yield tomorrow). Small groups are often better able to achieve consensus (Olson, 1965)[22].

Over the years, however, the Internet has grown dramatically: from 213 hosts in 1981 to 80,000 hosts in 1989 to over 700,000 hosts today[23]. The TCP/IP protocols are the most-widely used, vendor-independent family of networking standards publicly available[24]. Analysts have estimated that the worldwide commercial market for TCP/IP products and services was $607 million in 1990 and expect it to grow to over $1.8 billion by 1995[25]. In recent years, both

21 See Malkin (1992), which provides the curriculum vitae for 24 of the IAB and IESG members. Most of the members have long backgrounds either in academics, research centers or university computer centers. Most are Americans. A number of prominent participants were associated with SRI or Bolt, Beranek and Newman, one of the original ARPANET contractors.

22 With a fixed total benefit, private gains decrease as the size of the group which must share the benefits increases.

23 See Lottor (1992).

24 IBM's proprietary System Network Architecture (SNA) is the most widely used protocol suite. Digital Equipment Corporations DECnet is another vendor-specific architecture which is in wide-use. The only significant public standards competition for TCP/IP comes from the ISO-sponsored OSI standards (see further discussion below).

25 See Cook (1992).

the U.S. Congress (lead by Senator and Democratic Vice Presidential candidate Albert Gore) and the Executive Office have proposed a number of initiatives which would dramatically affect the future evolution of the Internet[26]. In December 1991, President Bush signed a $1Billion authorization for establishment of a National Research Education Network (NREN), which will be based on the Internet. In light of these events and the growing commercial importance of the TCP/IP suite of protocols for corporate wide-area "enterprise" networks[27], the Internet community is in the process of rapid change.

Participation in IETF meetings keeps growing and many of the new participants come from heterogeneous backgrounds[28]. They do not have long term relationships with either each other or with the present leaders of the IAB and IESG. Many are much less well informed regarding the engineering technicalities. According to Lyman Chapin, the new chairman of the IAB,

> "In size and extent, it [the Internet] has grown beyond what any of us envisioned 10 years ago. It's easily the largest telecommunications network in the world. At last count, 110 countries were connected...... There is also the difficulty of management. There is a much broader

26 In 1989, the White House Office of Science and Technology Policy published it's High Performance Computing Plan which called for an upgrading of the National Science Foundation's (NSF) NFSnet, which provides backbone services for the Internet. In 1991, the High Performance Computing Act (S.272) passed Congress which calls for the creation of a National Research Education Network (NREN). See Cook (1992) and Kahin (1992).

27 The personal computer became prominent in the first half of the 1980s. These were interconnected into Local Area Network (LANs) during the latter half of the decade. Corporations are now in the process of inter-connecting LANs into "enterprise-wide" networks spanning both the country and the globe.

28 According to Clark (1991), "over the past few years, there have been increasing signs of strains on the fundamental architecture, mostly stemming from continued Internet growth. Discussions of these problems reverberate constantly on many major mailing lists".

group of interests using the network and the issue of commercialization. The scope of interest in the Internet community is broader now and you can no longer manage it as you used to, by calling up a few graduate students....The biggest long-range problem is how the internet will be managed. Before, it was sort of an insiders' club and there was no formal structure." (see Messmer, 1992)

In recognition of these problems, the IAB and Internet are changing. In order to create a new institutional home for the IAB which is less beholden to the U.S. government and in recognition of the increased internationalization of the Internet, the Internet Society was created and held its first meeting this summer in Kobe, Japan[29]. In addition, a new IETF Working Group is being formed to discuss proposals on how (and if) the IAB and IETF should reform their decision-making rules[30].

The increased formalism which will perforce accompany these changes is necessary in light of the changing environment. The beginnings of this trend are apparent in the new IAB charter which specifies election voting rules and term limits for IAB and IETF officers[31]. In the future, the Internet Society Trustees will oversee IAB membership and nominations may be made by the membership of the Internet Society.

29 According to Lyman Chapin: "The effort now is to try to promote the Internet Society (a nonprofit organization formed this January) as an overall umbrella to include the IAB as well as the IETF and the IRTF, both chartered by the IAB." (Messmer, 1992). According to Kozel (1992), "the Internet Society is an attempt to impose structure on the Internet's loose confederation of interconnected networks. In a formal way, this global professional society will be an Internet institution answerable to those outside the Internet. It will set broad policy and standards, sponsor conferences."

30 A an August 13, 1992 email announcement to the IETF mail list, Phil Gross, the chairman of the IETF, reported that "....the IAB has recommended that the most effective means available to the Internet Community for gathering recommendations for refining our existing procedures is to request that an IETF Working Group be formed."

31 See Chapin (1992). The IAB is being renamed the Internet Architecture Board.

Much of this formalism is required in order to educate new participants and convince them that their longer term interests will be protected. In the absence of long-term relationships, the new participants seek assurance that IAB and IETF decisions will not be capricious. As the rents which are at stake have grown, so has the threat of opportunistic behavior and conflict. Increased bureaucracy and procedural rules are likely to slow the standards process, but this slowing is partially intentional.

When the potential for conflict and strategic rent-seeking are large, bureaucratic red-tape creates "early-warning" systems, reducing ex ante expected gains from opportunistic behavior. The increased difficulty in creating new standards provides greater weight to the technological "status quo". Since the adoption of a standard often requires specialized investments, the willingness-to-adopt is reduced if adopters fear future attempts to expropriate their quasi-rents. If it's too easy to overturn the status quo with a new standard, then there's little incentive to cooperate in today's adoption decision. In Lehr (1992a, 1992c), I argued that the procedures employed by traditional SDOs should be understood, in part, as a rational response to the collective choice problem of standardization when private and collective interests conflict. Based on recent trends in the Internet, I would expect the IETF to become more like traditional SDOs both in terms of structure and in terms of the difficulties faced in developing standards quickly.

Probably the most important change which may occur is the introduction of formal voting rules. Introducing voting into the IETF process will be problematic. First, if electronic voting is used (i.e., over the Internet), it will be difficult to supervise the allocation of voting rights.

Vote-packing and other abuses will be hard to detect and prevent. Any restrictions on who is allowed to vote would endanger the IETF's claim to generating consensus. On the other hand, using another mechanism for voting (e.g., requiring physical attendance at standards meetings) would increase participation costs and would diminish the central role which electronic communication has played in standards development. If the IETF decides to introduce voting to make access to decision-making power more open, they may need to become less open with respect to access to their proceedings and documentation (in the sense discussed earlier).

If the Internet Society is able to guarantee that the officers of the IAB and IETF fairly reflect the interests of the larger community, introducing voting more widely may be unnecessary. The recent moves to formalize the IAB and IETF election procedures appear to be steps in the right direction.

Part II: Anticipatory Standards and the Internet

i. How the IETF approaches anticipatory standards

In addition to the differences cited above, the IETF's approach differs from the approach followed by traditional SDOs in its attitude towards "anticipatory standards". Anticipatory standards are standards which are developed in advance of the frontiers of current engineering expertise. A classic example of anticipatory standard setting is provided by ISO's Open Systems

Interconnection (OSI) standards[32]. The OSI protocols offer the only public-standards, vendor-independent alternative to TCP/IP and are usually viewed as being in competition with TCP/IP[33].

The OSI effort began about the same time as the TCP/IP standards, but was much more ambitious. The developers of the OSI standards attempted to include the latest theoretical advances, whereas the developers of TCP/IP were careful to never attempt to standardize anything they could not implement[34]. The IETF's careful bias against standardizing in advance of the current state of market-operational technology is embedded in the standards process.

The essence of anticipatory standardization is the need to make decisions in the face of significant market uncertainty. All SDOs involved in de jure standard setting must face this problem, but the IETF chooses to address the problem in a special way. Whereas ISO, and many other standards institutions, seek to develop standards which (it is hoped) will be able to anticipate future technological and market trends and so will not have to be re-designed in the

32 See Tannenbaum (1988) or Rose (1990) for an introduction to OSI and it's 7-layered reference model.

33 For example, according to Desmond (1990), "the proposed standardization of four TCP/IP protocols has garnered widespread support among users in the Internet community but is drawing the ire of Open Systems Interconnection advocates who say the effort will slow the growth of OSI".

34 Rose (1991) compared the TCP/IP and ISO OSI development efforts thus: "...whilst work on both protocol suites began roughly at the same time, the Internet suite of protocols, by focusing on specific technical objectives, has by-and-large achieved the promise of Open Systems, and is rapidly becoming a commodity market. In contrast, by 1990, over ten years after the official "birth" of the OSI suite, there were only OSI pilot projects, nothing more." (page 6) Later Rose argues that "....by striving for generality and theoretical completeness in all things, proponents of OSI network management are producing a system which will likely be unable to respond to critical failures in OSI-based internets." (page 251). Although many analysts continue to argue that OSI standards will dominate the future, TCP/IP continues to dominate the present (see Ringling, 1992).

future; the IETF relies on "incremental standardization", developing standards step-by-step as the technology evolves[35].

The difference in approach is clearest when one compares the OSI and TCP/IP standards for network management. The OSI standard is CMIP, while the TCP/IP standard is called SNMP[36]. The effort to develop SNMP began in 1987 and the IAB approved the standard in May 1990. By 1991, over 200 vendors were shipping SNMP products, while ISO was still trying to complete the specification of the CMIP standard[37]. SNMP is much less complicated than CMIP[38]. Supporters of CMIP claim the added complexity is essential to meet the needs of managing largescale networks which require both accounting and security capabilities lacking in the original SNMP specification. The willingness to leave these features out of the initial standard as items to be dealt with in the future provides an example of the IETF's incremental approach.

The IETF's incremental approach is embedded in its lifecycle model for standardization. An Internet standard is supposed to progress through each of three states: proposed standard,

---

35 According to Perkins (1992) :"The (IETF) process encourages new ideas and innovation, yet restricts the progress of standards without independent implementation and interoperable deployment. This contrasts with the process used by other standards bodies which finalize a standard before it has been implemented or deployed."

36 CMIP stands for Common Management Information Protocol; SNMP stands for Simple Network Management Protocol.

37 See King (1991).

38 According to Satz (1992), SNMP's "primary goal was to be simple so nothing would stand in the way of its ubiquitous deployment....however, like many achieved goals, the primary strength can also become a weakness....an extreme example of simplicity versus power can be realized by comparing SNMP against CMIP....however, to implement CMIP, a subset of the protocol must be selected."

draft standard and Internet standard[39]. Associated with each standard is a status which may be either required, recommended, elective or not recommended. Although proposed standards are seldom (if ever) recommended, draft standards may have recommended status. Given the increasing returns which usually accompany growth in the installed base[40], this should accelerate the pace of standardization and make it more difficult to reverse a standard's progress.

The minimal time which must elapse before a proposed standard may be elevated to draft status is four months; and the minimal time before a draft standard may be approved as an Internet standard is six months. Draft standards which are not reviewed within 6 months are removed from the draft directory; and, the IESG reviews the charters of working groups which have not produced a standard in 24 months. Either you "move it, or lose it".

These tight timing constraints are complemented by a requirement that proposed standards demonstrate their practical feasibility. The IESG will not recommend a proposed standard to the IAB for promotion to draft status until there are at least two independent and interoperable implementations communicating in the Internet. This prevents the standards process from progressing faster than the current frontiers of operational engineering expertise. This too is distinctive since other SDOs like ISO do not require working implementations of the proposed

---

39 There are also experimental and historical protocols, which are not yet or are no longer considered on track to become approved Internet standards. See Postel's (1992) RFC1280 or Chapin's (1992) RFC1310 for a lengthier discussion of standardization procedures.

40 A supporter of the ANSI X3 approach to standard setting, commented to me that the IEEE's willingness to publish interim standards was "like being almost pregnant". Letting draft standards achieve recommended status appears as an even more aggressive mechanism for speeding up the process.

standard as a pre-condition for approval[41].

ii. A Model of Anticipatory Standards

One of the most important problems faced in de jure standard setting processes is the need to make decisions in the face of significant technical uncertainty. In a de facto process, firms rely on market competition to determine which technology or technologies emerge as de facto standards. During the competition, practical experience is gained, making it clearer which technology offers superior benefits. Unfortunately, this is not always the technology which triumphs. If network externalities are important enough, the firm with the largest installed base may win with an inferior technology. Perfect 20-20 hindsight may allow us to identify which technology we should have adopted yet it may no longer be collectively worthwhile to reverse the market's verdict. Since firm-specific investments (e.g., sunk investments in learning, advertising, specialized capital equipment, etc.) are likely to accumulate during the market competition, firms may become unwilling to compromise on a common design.

The economics literature usually distinguishes between de jure and de facto standardization based on whether market introduction follows or precedes the emergence of a standard. Farrell and Saloner's (1988) path-breaking analysis of standards processes used the "Battle of the Sexes" game to show why committee-based de jure standardization may be superior to market-based de facto standardization. Even when delay is costly, the more time-

41 When comparing the ISO and IETF approaches, Rose (1990) comments that "*all* Internet technology is based on things which are known to work before standardization occurs." (p591, italics in original)

consuming de jure process is unambiguously preferred to the de facto process because of the greater likelihood of coordination failure under the latter[42]. Their analysis highlights the difficulty of achieving coordination when participants' payoffs depend on which of the coordination equilibria are selected.

Let us ignore for the moment the difficulties raised by strategic competition and consider the more general problem of making decisions in the face of uncertainty which will be revealed over time. If there is no cost to waiting and it is costless to reverse earlier decisions, there is no difference between making a choice today (which we can change if we find out we made a mistake) and waiting until the uncertainty is resolved before making our decision. On the other hand, if waiting is sufficiently costly, it will be preferable to make decisions before the uncertainty is revealed, or if possible, to invest in resolving the uncertainty. Farrell and Shapiro (1992) explore this question in their analysis of High Definition Television (HDTV) standardization. Using a Mechanism Design approach, they explore the optimal timing for standardization. The planner balances the gains from waiting until a better technology (i.e., lower cost) becomes available and the opportunity costs of postponing the benefits from early adoption[43]. Their analysis is useful in evaluating alternative policies for government intervention in the standards process or where participants can commit to abide by the decisions

42 With costly delay, both the committee and the market processes are resolved more quickly, yet the market process results in coordination failures more often. The committee process is modelled as a "War of Attrition" game and the de facto process as a "Grab the Dollar" game. There game assumes symmetric, perfect information.

43 The larger the installed base, the greater the network externalities. Earlier realization of these benefits increases the present value of total benefits from adoption.

of a central authority.

If we regard the IAB (or IESG) as a central authority, then Farrell and Shapiro's approach may be appropriate. Since the goal of this paper, however, is to see what lessons can be learned which are useful to voluntary standard setting more generally, I would like to consider a slightly different model inspired by Farrell and Saloner's (1988) earlier work to examine how uncertainty influences the choice of standard setting process.

Consider a situation where two firms are deciding whether to adopt compatible designs. Assume that it is clearly beneficial for them to choose compatibility, but that the joint benefits which will be created are either A or B, and 0 if non-compatibility is chosen[44]. Let $A > B > 0$ so that total benefits are maximized if A is adopted, yet adopting B is preferred to no agreement. The firm whose design is adopted (i.e., "the winner") will receive a share s of the total benefits, while "the loser" will receive (1-s). Assuming that a firm whose technology is adopted as an industry standard expects to capture a larger share[45] to the total joint-benefits created, let $s \geq ½$.

44 In its present form, the model described here is timeless. Implicitly, I am assuming all participants are risk neutral with quasi-linear utilities in monetary transfers. Think of all amounts as being measured in present value terms.

45 This is not always true. Pioneers often end up with arrows in there back. However, this is the usual assumption in the literature. All that matters for the present model, however, is that participants prefer different coordination equilibria which produce different levels of total surplus.

We can represent this as a "Battle of the Sexes" game with the following payoff matrix:

		Sponsor A	
		Adopt B	**Adopt A**
Sponsor B	**Adopt B**	**sB,(1-s)B**	**0,0**
	Adopt A	**0,0**	**(1-s)A,sA**

This game has three Nash equilibria: two pure strategy equilibria AA and BB, and a mixed strategy equilibrium. The mixed strategy equilibrium with the sponsors of A and B choosing A and B with probabilities a and b, respectively, is given by the following probabilities:

$$a = \text{Prob. sponsor of A chooses A} = sB/[(1-s)A + sB]$$

$$b = \text{Prob. sponsor of B chooses B} = sA/[(1-s)B + sA]$$

The expected joint benefit under the mixed-strategy equilibrium is,

$$a(1-b)A + ((1-a)bB = \frac{s(1-s)AB(A+B)}{((1-s)A + sB)(sA + (1-s)B))}$$

which can be shown to be less than B, the joint benefit realized under the less attractive of the two pure strategy equilibria. This occurs because under the mixed equilibrium there is a significant probability that the two firms fail to agree on a common design. The probability they fail to agree (i.e., play AB or BA) is:

$$\text{Prob. no agreement} = ab + (1-a)(1-b) > 0$$

As s approaches 1, both a and b approach 1 and the probability that the two firms fail to adopt compatible designs goes to 1, yielding an expected mixed strategy net benefit of 0. This confirms our intuition that bargaining failure should become more probable as the distribution of the joint-benefits becomes more asymmetric[46].

Now, consider what happens if neither firm knows whether it's technology is A or B. The firms can agree to adopt one of their technologies today or they can wait until they know which firm owns which technology, and then decide which to adopt. Even if they decide to adopt today, they may decide to reverse their decision tomorrow once they know which is the best technology. Let $W \leq B$ represent the cost of waiting in terms of the total reduction in net benefits available from adopting a compatible solution tomorrow. Thus, if they wait to decide, their joint benefit will be between 0 and (A-W) depending on whether they fail to adopt compatible designs or if they adopt the better of the two technologies.

Similarly, let $R \leq B$ represent the cost of reversing today's decision tomorrow. If the decision is reversed[47], then the joint benefit will be between 0 and (A-R). Letting p be the probability that B is selected if an agreement is reached at Stage 1, then we can represent the

46 It is also worth noting that b>a so in the mixed strategy equilibrium, the probability that B insists on B is larger.

47 If one but not both firms reverse the adoption decision, then assume this is a coordination failure resulting in total surplus of zero.

possible outcomes as in Figure #1.

Assuming for the moment that an efficient agreement is reached at each stage, the expected joint benefit from either agreeing today and never changing; agreeing today and changing if a bad choice was made; or waiting to agree tomorrow are:

	Expected Joint Benefit
Agree and never change	$(1-p)A + pB$
Agree and maybe change	$(1-p)A + p(A-R) = A - pR$
Wait and Agree	$A - W$

If $R < (A-B)$ and $pR < W$, then the first best optimum is to agree today and reverse our decision tomorrow if necessary. Let us now consider four possible cases.

Case 1: W small, so waiting is good idea (i.e., $W < pR$ and $W < B$)

Within the context of standard setting, we can think of the "wait and decide" approach as either representative of a situation like the one examined by Farrell and Saloner (1992) involving HDTV standardization, or as representative of de facto market standardization. Some of the participants in formal SDOs appear to believe that SDOs can never produce useful standards in advance of a market-based decision. SDOs still may play a valuable role in de facto standardization. For example, formal endorsement of a de facto standard by an SDO may legitimize a technology, which in turn may reduce consumer search costs (who are likely to be less well-informed than direct participants in the standards competition); facilitate entry by new suppliers; and reduce transactions costs (e.g., purchase orders and performance contracts can reference the standard). Although these functions may be important, they are not the primary focus of this paper. I am interested in the role of standards in shaping the choice of technology and so will consider situations where the SDO acts before significant market experience is obtained.

Even though the IETF requires working implementations before a standard is approved, the IETF will approve standards before the market's verdict is clear and will balk at attempts to provide "rubber stamp" approval to a firm's de facto technology.

Case 2: R Large (i.e., R > A) and W large (i.e., W > A)

As noted above, it is often quite costly to wait before attempting to adopt a standard. We may fail to adopt A or B if vested interests grow over time. Also, we prefer not to postpone the realization of the benefits of compatibility. Committee-based standard setting is most important and interesting when the costs of waiting, W, are sufficiently large that we always prefer to try and reach an agreement today.

Even when waiting is not desirable, it may be extremely costly to reverse an earlier decision[48]. Just as with waiting, there may be large switching costs associated with converting the installed base associated with the earlier adopted, albeit inferior technology. If it is sufficiently costly to reverse a bad decision, then the problem at stage 1 is to solve the "Battle

48 Most of the earlier literature on standard setting assumes adoption decisions are irreversible. See Cowans (??).

of the Sexes" game with uncertain payoffs. Let q be the commonly expected prior probability that Firm # 1 is sponsoring A and Firm #2 is sponsoring B; and without loss of generality, let $q \geq ½$. With these assumptions, the expected payoffs under different agreements are[49]:

		Firm #1	
		Adopt 2's	**Adopt 1's**
Firm #2	**Adopt 2's**	**s((1-q)A+qB),(1-s)((1-q)A+qB)**	**0,0**
	Adopt 1's	**0,0**	**(1-s)(qA+(1-q)B),s(qA+(1-q)B)**

If the firms bargain efficiently such that the Coase Theorem applies (e.g., they can arrange side-payments), then we should expect Firm #1's technology to be adopted[50], producing an expected joint benefit of qA+(1-q)B. If side-payments are impossible, however, then bargaining may break down (a la Farrell and Saloner's (1988) analysis) and the Pareto-inferior mixed strategy equilibrium may prevail. Although each firm prefers to adopt the other's technology rather than fail to coordinate, they disagree regarding which technology is preferred whenever the distributional effects, as measured by s, are sufficiently strong. Firm 2 prefers

49 I am assuming that if the participants are uncertain regarding market shares,s, or the relative difference in values between the superior and inferior technologies, (A-B), or R and W, this uncertainty is not resolved in the course of the standards process. My intuition suggests that increased uncertainty in s may make it easier to reach an agreement if the expected value of s=.5. Similarly, increased uncertainty which leads to ex ante expectation (A-B)=0 may make it easier to reach agreement. Increased uncertainty in R and W, however, may either improve or aggravate the problem of reaching an early agreement.

50 Under the assumption that it is more likely that firm #1 is sponsoring the superior technology, the expected value of winning for firm #1 always exceeds the value to firm #2 by an amount (2q-1)(A-B), thus there are gains from trade to be realized by firm #2 conceding.

winning whenever[51]:

$$s > s_2 \equiv (qA+(1-q)B)/(A+B)$$

even though everyone believes Firm #2 is more likely to be sponsoring the inferior technology. Even if q=1, Firm #2 would sponsor it's technology for sufficiently high s.

This suggests one obvious way to reduce conflict is to reduce the distributional effects of standards. Standards committees try to select technologies which do not offer too great an advantage to any participant[52]. Unfortunately, it is seldom possible to select competitively neutral standards.

Let us imagine for a moment that q=½ so that ex ante the firms are completely uninformed as to which technology will produce the larger total benefits. Even if each prefers to win, they would prefer to submit to a decision by an unbiased mediator who agrees to pick the "best" technology. If q=½, then the mediator may as well flip a coin and the expected pay-off to each firm is symmetric and equal to ¼(A+B) which is greater than the expected mixed strategy pay-out[53]. Suppose that after the firm's agree to submit to the mediator, they privately learn the true value of q. Their incentive to misrepresent the true value of q would make their information suspect. Thus, collective learning in the committee may be biased by attempts to use

51 Firm 1 prefers winning if $s > s_1 \equiv ((1-q)A+qB)/(A+B) \geq s_1$. Note also that $s_2 \geq ½$, but $s_1 < ½$ is possible.

52 In the development of the Ethernet standard, the IEEE P802 committee modified a number of technical details in the original Ethernet proposal from DEC, Xerox and Intel in order, partially at least, to reduce the competitive advantage which would accrue to the sponsors if their proposal were adopted without modification (see Lehr (1992b)).

53 The expected total benefits are ½(A+B) and each firms expected market share is ½ so each firm's expected payout is ¼(A+B).

information strategically to alter the committee's adoption decision.

Case 3: R small ($0 \leq R < \min(B, A-B)$) and W large, but only efficient reversal

If the cost of reversing decisions is low, then it may be possible to reverse poor first stage decisions. The current approach to software standardization is based on layering of protocols. This modular approach permits the replacement of individual layers and is one mechanism by which the costs of reversing (or, equivalently, enhancing) an inferior earlier decision are reduced. Gateways and other converter technologies which lower switching costs also help to reduce the cost of reversing decisions, R.

The simpler, more modular the standard, the more easily it may be replaced (and hence the lower R). This suggests that the IETF's minimalist approach may face lower R than the more ambitious, heavily-optioned approach taken by ISO. Moreover, it seems likely that R should be lower for software than for hardware standards because of the (generally) higher costs associated with replacing hardware. Also, because it is more difficult to predict the performance of software without trying it out, we might expect that software developers would invest more effort in lowering R, and thus might expect R to be lower for software standards[54]. This may help explain why the CCITT and ISO (i.e., high R processes) and the IETF (i.e., low R process) differ.

54 Since you have to run software to find the bugs, you know you'll have to fix bugs after you distribute the software so you design the software to make it easier to fix (e.g., modular sub-routines, etc.) Software is often distributed via upgrades (although this is also a mechanism for price discrimination). Relatively recently, CPU chip vendors have begun designing personal computer CPUs to facilitate modular chip set upgrades.

When R is low, however, there is a danger of opportunistic behavior. The winner at the first stage may prefer to resist a jointly profitable switch to a superior technology; while the loser at the first stage may prefer to switch to an inferior technology. The threat of future opportunistic behavior will influence bargaining in the first stage.

Let us consider what would happen if the SDO were able to assure that the first stage decision would be reversed if and only if it would increase joint-benefits. In this case, the expected joint benefit would be (A-pR), and the payoff matrices at the first stage would be:

		Firm #1: Adopt 2's	Firm #1: Adopt 1's
Firm #2	Adopt 2's	**s((1-q)A+q(A-R),(1-s)((1-q)A+q(A-R)**	**0,0**
	Adopt 1's	**0,0**	**(1-s)(qA+(1-q)(A-R)),s(qA+(1-q)(A-R))**

The private gain which each firm expects to capture from having its technology adopted instead of its competitor is:

Firm #1's net value for winning = (s+q-1)R

Firm #2's net value for winning = (s-q)R

Notice that in this case, were either firm to privately learn which technology were superior, they would have an incentive to report it truthfully. Since the best technology will be adopted in the end, winning with the wrong technology at stage 1 only decreases their expected payments. Not surprisingly, the most desirable decision-making framework is one which can guarantee ex ante that information will be used solely to maximize total joint benefits.

Case 4: R small, W large, No restrictions on reversal.

Suppose there is no way that the participants can pre-commit to not behaving opportunistically after the uncertainty is resolved. This is analogous to the case where only short term contracts are possible at stage 1. Both Cases 2 and 3 correspond to special cases of long term contracts, where a long term contract is defined as a contract which seeks to constrain future behavior. If we think of the two firms in the model as actually representing two opposing industry coalitions, it becomes clear why long term contracting may be difficult. As uncertainty is revealed over time, it is likely to affect coalition members asymmetrically, making it difficult to sustain the sponsoring coalition[55].

If we could presume efficient bargaining at both stages then the outcome would be the same under both short-term and long-term contracting (Milgrom and Roberts (1988))[56]. In the model presented here, I implicitly assume that bargaining is not efficient since the only distributional agreement which is allowed is pre-determined by the market, $s \geq ½$. In this case,

55 This is one consideration which motivates the design of regulatory agencies, where procedures are selected to assure that the Agency will mirror the interests of the original sponsoring coalition (see Noll, McCubbins and Weingast (1987)). The recent history of the Advanced Computing Environment consortium demonstrates the difficulties of sustaining industry coalitions over time.

56 In the long-term contracting regime, the firms could explicitly prohibit opportunistic behavior; while in the perfect short-term contracting regime they could use first stage side-payments to protect themselves against opportunistic behavior in subsequent stages. [Need to check if this is really true!!!]

there is a risk that bargaining will break-down at the second stage. If both firms expect the mixed strategy equilibrium to prevail at the second stage then this reduces the expected payoff from bargaining strenuously at stage 1, and may make it easier to achieve an agreement (relative to case 2)[57]. Unfortunately, it also reduces the participants expected gains from participating in standards processes and may lead to insufficient investment in developing new standards[58].

A more open standard setting process may make it more difficult to reverse a good decision or to resist reversing an inefficient decision. If customers care only that the best technology is adopted then their participation may offer a mechanism for discouraging opportunistic behavior at the second stage. In the limit, an expectation of such an effect may be sufficient to make this case observationally equivalent to Case 3. On the other hand, if efficient enforcement of the reversal rule is not certain, there are likely to be bargaining position (i.e., distributional) benefits associated with winning in the first stage, even with an inferior technology. However, it seems reasonable that the reduction of the threat against opportunistic behavior would tend to enhance the efficiency of the overall decision process.

57 This suggests that if we modelled first stage bargaining using a "War of Attrition" game as in Farrell and Saloner (1988), there might be a reduction in the delay (or bargaining cost) required to reach an agreement.

58 Since it is costly to participate in standards committees, the lower the expected private benefits, the less likely a firm will be to participate. Since standards are public goods, there is already a free-rider problem. In the model here, I am focusing on the benefits from compatibility which accrue to the participants; however, there are also significant benefits which are distributed more widely (e.g., to new suppliers and customers).

iii. Relating the Model to the Internet Process

I believe the "incremental" approach followed by the IETF is most closely described by Case 4; whereas the more traditional SDO approach is more like Case 2. In the past, when the Internet community could rely on long-term relationships to discourage opportunistic behavior and the community was sufficiently small that bargaining costs were relatively low, the incremental approach worked quite successfully, approaching the ideal of Case 3. It is unlikely that this approach will work as well in the future; however, it is unclear whether the "Case 2" approach employed by the SDOs is viable.

In the models present elementary form, it is not yet possible to adequately address these issues. However, the analysis thus far does suggest some interesting avenues for future research. First, does a low R (which I have argued is likely to be more typical of software standardization) make it easier to implement feedback loops in the standards process? Are there circumstances when the increased ex ante uncertainty actually improves the decision process (as hypothesized earlier)? Are there real world examples where user participation helped reduce the threat from opportunistic behavior and improved the efficiency of the decision process? These and other interesting questions must await the fuller specification of the model introduced above.

Conclusions

This paper examined the processes used to develop technical compatibility standards within the Internet community in order to account for the perception among industry participants that the process was superior to the approach followed by more traditional SDOs such as ISO or ANSI. It seems that most of the Internet's prior success may be accounted for by specialized circumstances which are currently changing. In the more heterogeneous and contentious environment wherein the Internet finds itself, a more traditional and formal approach to standard setting is called for. However, if the Internet seeks to implement formal voting procedures (a la ANSI) it may find the openness afforded by its extensive reliance on electronic communication creates difficulties not confronted by other SDOs.

The Internet process is also interesting for the light it sheds on standardization in the face of uncertainty. SDOs such as ISO which attempt anticipatory standards which are designed to limit the need to modify future standards (i.e., do it all today) may face higher bargaining costs today which may increase the likelihood of cooperation failure[59]. Moreover, these processes appear to create disincentives to share private information efficiently.

On the other hand, the Internet's incremental approach (which resembles a series of sequential short term contracts) may offer a superior approach when the threat of future opportunistic behavior is reduced via the monitoring of customers interested in the maximization of joint-benefits. Under these circumstances, we may realize improved incentives to share private information in the first stage. A more formal presentation of these ideas and further exploration of related questions must await continued refinement of the preliminary model presented here.

---

59 A delay in reaching agreement which is sufficiently long to make a standards irrelevant in the market is an example fo coordination failure.

Bibliography
Compatibility Standardization and the Internet

Axelrod, R., The Evolution of Cooperation, Basic Books, New York, 1984.

Besen, S., "The European Telecommunications Standards Institute", Telecommunications Policy, December 1990, 521-530.

Besen, S. and J. Farrell, "The Role of the ITU in Standardization", Telecommunications Policy, August 1991, 311-321.

Buerger, D., Editorial, CommunicationsWeek, No. 412, July 20, 1992, 42.

Cerf, V., "The Internet Activities Board", RFC1120, September 1989.

Chapin, L., "Charter for the Internet Architecture Board (IAB)", RFC1358, August 1992.

Chapin, L., "The Internet Standards Process", RFC1310, March 1992.

Clark, D., "Towards the Future Internet Architecture", RFC1287, December 1991.

Cook, G., The National Research and Education Network: Whom Shall It Serve?, published by Gordon Cook, 1992.

David, P. and S. Greenstein, "The Economics of Compatibility Standards: an Introduction to Recent Research", Economics of Innovation and New Technology, vol 1, 1990, p3-41.

Desmond, P., "Std. push for TCP/IP miffs OSI backers", Network World, June 18, 1990, page 1.

Farrell, J. and G. Saloner, "Coordination through Committees and Markets", Rand Journal of Economics, Vol. 19, No. 2, Summer 1988, 235-252.

Farrell, J. and C., Shapiro, "Standard Setting in High Definition Television", Brookings Papers on Economic Activity, Microeconomics, (1992) 1-94.

Gibson, S. and S. Girishankar, "IAB Backs Off Address Proposal", CommunicationsWeek, No. 412 (July 20, 1992) 1,65.

Kahin, B. (ed) Building Information Infrastructure: Issues in the Development of the National Research Education Network, McGraw-Hill Primis, New York, 1992.

King, J., "Still Up in the Air -- Interconnected LANs make Integrated Network Management Even Less Possible, but Standards and Vendor Alliances Bring Some Pieces Together", Computerworld, vol 25, no 1, May 27, 1991, 71-72.

Kozel, E., "Commercializing the Internet: Impact on Corporate Users", Telecommunications (International Edition), vol 26, no 1 (January 1992) s11-s14.

Lehr, W., "The Political Economics of Voluntary Standards", mimeo, Columbia University, March 1992a.

Lehr, W., "Compatibility Standards and Industry Competition: Two Case Studies", mimeo, Columbia University, June 1992b.

Lehr, W., "Standardization: Understanding the Process", Journal of the American Society for Information Science, vol. 43, no. 8 (September 1992c) 550-556.

Lottor, M., "Internet Growth", RFC1296, January 1992.

McCubbins, M., R. Noll, and B. Weingast, "Administrative Procedures as Instruments of Political Control", Journal of Law, Economics and Organization, vol 3 (1987) 243-277.

McQuillan, J. "Setting a faster standards-setting pace", Network World, March 11, 1991, 35.

Malkin, G., "Who's Who in the Internet: Biographies of IAB, IESG, and IRSG Members", RFC1336, May 1992.

Messmer, E., "Internet Architect Gives Long-Term View", Network World, v9n20 (May 18, 1992) 37,46.

Milgrom, P. and J. Roberts, "Bargaining Costs, Influence Costs, and the Organization of Economic Activity", Chapter 3 in J. Alt and K. Shepsle (eds), Perspectives on Positive Political Economy, Cambridge University Press, New York, 1990.

Moad, J,, "The Standards Process Breaks Down", Datamation, vol 36, no 18, September 18, 1990, 24-32.

Olson, M., The Logic of Collective Action, Harvard University Press, Cambridge, MA, 1965.

Perkins, D., "Standards", Simple Times, v1.2, May/June 1992, 10

Postel, J., "IAB Official Protocol Standards", RFC1280, March 1992.

Ringling, F., Network Management: Strategic Directions in Local and Wide-Area Communications through 1996, Insight Research Corporation, 1992.

Rose, M., The Simple Book: An Introduction to Management of TCP/IP based Internets, Prentice-Hall, New York, 1991.

Rose, M., The Open Book: A Practical Perspective on OSI, Prentice-Hall, New York, 1990.

Satz, G., "Technical Article: Origins of SNMP", Simple Times, vol 1, no. 1, March/April 1992, 1.

Tannenbaum, A., Computer Networks, Second Edition, Prentice-Hall, Englewood Cliffs, 1988.

Vasquez, W., "Standards groups have their work cut out for them ", Network World, January 15, 1990, 37.

Note: A number of the documents referred to in this bibliography are referenced as "RFCxxxx", which means they are part of the series of on-line documents which are maintained by the Internet Community. These may be accessed via FTP in the sub-directory RFC on the host NIC.DDN.MIL as user Anonymous. The file RFC-Index.txt contains an index to the complete collection of these files (which currenly number over 1,350).