A BROADER VIEW ON SOME FORCES SHAPING STANDARDISATION

Kai Jakobs

TECHNICAL UNIVERSITY OF AACHEN

ABSTRACT: The paper tries to unveil and discuss some rather more economic aspects inherent to standards and standardisation, as seen form a non-economist's perspective. What would happen if it were solely up to the market to standardise? What are the motivations behind corporate activities in standards setting processes, which are, after all, costly and tend to be frustratingly slow? Why would users want to have standards? Answers to these questions are given based on a critical literature review, combined with insights obtained from a recent survey. And finally, how could today's processes be modified to take into account all those economic aspects? A proposal is made in response to this question.

A Brief Introduction

Standardisation may have far-reaching impact on companies and even on full grown economies. Placing money on a technology that eventually fails to become a standard, and to be adopted by the market, may easily lead to the breakdown of a company. Pros and cons of joining the standardisation bandwagon vs trying to push a proprietary solution need to be considered. Standards based products or services may imply price wars and lower revenues, but may also open new markets and widen the customer base. Offering a proprietary solution may yield (or keep, rather) a loyal customer base, but may also result in a technological lock-in and, eventually, marginalisation for a vendor or service provider.

Assuming that standardisation is a desirable goal per se, and that it will take place anyway, the problem of how to select the 'right' standard, or how to standardise on the 'right' system, needs to be addressed. 'Right', of course, means different things to different people, which is why at least the non-technical dimensions of standardisation tend to be very tricky. Something that is 'right' for one country or one company may be disastrous for another. The international and possibly global scale of standards in the field of Information and Communication Technology (ICT) means that players with very different backgrounds from very different economies need to agree on something they deem to be at least more or less 'right'. For obvious reasons, some of which will be outlined below, those problems have been attracting quite some attention, especially against the background of the emerging Global Information Infrastructure (GII).

The remainder of the paper is organised as follows: some potential results to be experienced if standardisation were left to the market alone are discussed in chapter two. In chapter three I will address corporate strategies in standards setting, have a brief glimpse inside the committees, and look at some compatibility issues. Subsequently, some very basic economic issues are discussed in chapter four, and the user's point of view is presented in chapter five. Finally, a consequence of all these deliberations - a proposal for a modified standards setting process - is presented in chapter six.

Let the Market Standardise?

Let us imagine what might happen if the market had to decide upon which technology to standardise. Several results are possible; one, of course, being that an optimal technology (or at least the best alternative available) actually wins. There is no need to discuss this case further. There are other possible outcomes, though. Consider, for example, a situation where different, but roughly equivalent technologies are available, none of which commands sufficient support to establish itself as the 'standard'. It may now well happen that this uncertainty paralyses the market, and that potential buyers postpone their purchases in order not to invest in a losing technology. As a consequence, innovation in that technical domain would come to a near standstill. Clearly, nobody would benefit from a situation like this.

The notion of 'uncertainty' is important here. Standards are but a part of a larger socio economic system, which does exert a certain amount of influence on standards' development. For example, the more recent need for new technologies to be environmentally sound, due to the environmental awareness that came virtually out of the blue. As one result, nuclear power has become less desirable in many parts of the world [Cowan 92]. That is, a standard is subject to path dependencies imposed on it by its broader environment. Unforeseen, and indeed unforeseeable developments may hamper all efforts and may even impose the need to start an activity all over again from scratch. Moreover, in most cases a standard is not a stand-alone document. Rather, it is positioned in a network of other standards (some of them possibly only emerging), which influence the boundary conditions within which it can emerge by laying down, for example, stringent compatibility requirements. Last but not least, early decisions made during the standardisation of a technology itself may have significant impact of later decisions. Selecting the telephone network as the carrier for facsimile transmissions, for example, implicitly pre-defined numbering schemes to be used and possible transmission speeds [Schmidt 92], as well as the need to eventually switch from analogue to digital transmission technology. More general, path dependencies were established at an early stage of the process, which to a considerable degree shaped subsequent developments.

It should be obvious that hardly any sufficiently accurate forecasts can realistically be made regarding future developments. Whilst this uncertainty affects all predictions, it has a particular strong impact on standardisation. Here, big oaks from minor acorns grow. That is, comparably small events may carry great weight; in the absence of a sound basis for judgement and decisions the adoption of a particular technology by just one firm may encourage others to follow suit. If this happens, chances are that an inferior technology will be adopted, which may suit the initial adopter (who will have evaluated the alternatives and selected the technology to best suit his needs), but does not necessarily meet other entities' demands. They, in turn, will then make their choices solely based on the initial adopter's policy decisions. Little, if any experimentation with alternative technologies or systems will occur, which will rapidly be discarded. A similar effect may be observed when decisions to adopt are based only on initial expressions of a technology (e.g. implementations of IT systems). In such cases, a poor first implementation can easily reduce to zero this technology's chances of being adopted, since possibly superficial, implementation-specific shortcomings hide the technology's inherent advantages (for example, X.400 suffered from many inadequate early implementations). Likewise, observable early benefits of a technology will outweigh all other aspects; in particular, higher benefits to be gained from a different technology at some later stage will be ignored. Indeed, these benefits again cannot be identified at all due to the lack of opportunities for experimentation. It follows that the market can - and frequently will - adopt the 'wrong' technology when left on its own. 'Wrong', like 'right', is of course a vague term; a technology may appear to be 'right' for a particular adopter, but at the same time the adoption may have negative impacts on others [Besen 95].

It is most likely that the above course of events could sooner or later be observed if standardisation were left to market forces alone (the DOS operating system was a case in point: one strong player, IBM, chose this system, which did not really represent state-of-the-art at that time, and almost all others followed suit. Obviously, IBM gained significant profits from this development - as did, even more so, Microsoft. In fact, to some extent users gained as well, albeit not from superior technology, but solely from the emerging network externalities). Consequently, some form of co-ordinated standardisation efforts is required. (Prospective) standards surely try to reduce uncertainty by aligning players' views and expectations. Indeed, the pure existence of a standards setting process might suffice to prevent the development outlined above, as it would then be possible to raise expectations that a standard will soon be emerging from this process [David 95].

Corporate Strategies and Their Consequences

The desire to make sure that the 'right' standard emerges normally lies at the heart of firms' involvement in the standards setting process, be it in the 'official' process or in consortium-led activities. Yet, what exactly characterises the 'right', or at least a 'good' standard is far from being clear. Cowan associates a good standard with the attributes 'speed' and 'meet technical requirements' [Cowan 92]. Whilst these characteristics are valuable for winning stakeholders' support, this is a surprisingly narrow focus. Clearly any technical specification should meet technical demands, the issue of speed, however, is popular, yet questionable. In [Jakobs 98], for instance, it has been shown that most users of a particular IT service (e-mail) were not able to identify meaningful requirements, even after over a decade of service usage. It is safe to say that even fewer requirements will typically be available at the beginning of a standards setting activity, an observation warranting certain doubts about the quality of a speedily done specification, which will inevitably be based on completely unverified requirements, potentially far off the mark. Moreover, to meet organisational and societal requirements should clearly play a role in standards setting as well. Regrettably, though, trying to prevent standards from coming into being may also be a motivation for participation.

Standardisation is becoming all the more important with the increasing economic and corporate globalisation. At the same time, standardisation politics change. Strangely, national interests are becoming more important. A domestic standard successfully introduced into the global arena will not least boost the prospects of the domestic economy. Accordingly, governments now have a vested interest in pushing such standards to support domestic firms. These firms, in turn, will look to standards setting for several reasons which are typically, though not necessarily, related to their own economic well-being. Standardisation may thus to some degree be seen as an interface between technical and non-technical (e.g. economic, organisational or social) considerations. That is, standards are not only rooted in technical deliberations, but also result from a process of social interactions between the stakeholders (particularly including governments and user/vendor companies). These dynamic interactions are projected onto the standardisation bodies' committees, where another dimension is added, that of the individual.

Inside the Committees

Schmidt and Werle [Schmidt 92] note that the common engineering background of most committee members will lead to a co-operative situation where all participants will work towards the 'best' technical solution. Yet, this assessment may well be a bit over-simplistic. Whilst most committee members indeed have an engineering background, they do not necessarily solely strive for technical brilliance; if they do, there will be a real danger of over-engineered solutions. It all depends on the respective roles they assume (see e.g. [David 94], [Jakobs 98]). Along a similar, but more realistic line Cowan notes that once the basic choices have been made at a later stage of the process standardisation work becomes more co-operative, now very much resembling joint R&D efforts [Cowan 92]. Information is shared, and tasks are undertaken co-operatively. Resources dedicated by the different stakeholders, i.e. commitment demonstrated through their willingness to conduct quality research, prepare high-quality proposals, and take over responsibilities in the committee are important. Likewise, the technical, diplomatic and political capabilities of their representatives must not be underestimated (see e.g. [David 94]). These commitments and capabilities create asymmetries within the committees which may be exploited by a player.

Other gains than those purely associated with a successful proposal may result from participation in standards setting. Many committee members only participate for reasons related to intelligence gathering. For example, information regarding strategic moves of competitors or recent technical achievements may be gained, yielding a better evaluation of a company's position relative to its competitors. A recent survey showed that about fourteen per cent of committee members belong to that category [Spring 95]. Moreover, a company's reputation may rise due to its commitment to standardisation, which (potential) customers may associate with a dedication to high quality.

Compatibility

Globalisation may have further impact on standards setting. Besen claims that market growth may well lower the need for compatibility, and that variety may be possible without a negative impact on market growth. He goes on to argue that variety itself may be a source of growth [Besen 95]. This is a dangerous proposition, though. Pushing it only a small step further we will find companies introducing variety purely for (their own) growth's sake. Indeed, there is a likelihood that companies will follow strategies of deliberately introducing incompatibilities to tie customers to their systems [David 95]. If this happens, there will be an urgent need for other entities to counter such moves by backing alternative, compatible proposals. To add weight and credibility to this move, this alternative system could be introduced into the standardisation process.

If no compromise can be achieved when competing proposals exist, one possible outcome will be the formation of a new 'standards'-setting consortium established by one of the rival entities. This might also be an explanation for the alarming expansion of the number of 'standards' consortia (see also e.g. [Cargill 95]). A 'Balkanisation' [Besen 95] of the standardisation process, with competing bodies developing competing specifications would be a potential result, which would further contribute to a deliberate introduction of incompatibilities.

Somewhat surprisingly, it has been argued that situations may exist where compatibility is neither practicable nor a socially desirable goal to achieve. ‘Genuine uncertainty’ regarding the best specification is offered as a justification for this proposition [Richardson 97]. Yet, this immediately brings us back to the above discussion on ‘uncertainty’, and the potentially disastrous results that may be expected if ‘standardisation’ is left to the market (which would happen in case of uncertainty with no standard on the horizon).

Keeping to the issue of incompatibilities despite standards we note that interworking units (e.g. gateways) may contribute to lasting incompatibilities, as they enable information exchange across, and maybe even interworking between heterogeneous (i.e. incompatible) networks, albeit typically with a loss of some functionality or information. Whilst good gateways may offer a higher level of interoperability, it must be noted that at the same time the better the gateways the more they serve to entrench incompatibility. Here, little motivation exists to go to great lengths to install an overall compatible, standards-based system since potential gains are comparably small thanks to the good gateways. An almost comic touch is added to this situation when gateways themselves become subject of standardisation efforts.

Some (Very) Basic Economics

Being active in standards setting is a costly business. It has been estimated that the costs for the development of an average IT standard amount to about $ 10,000,000 [Spring 96] - and that is only one standard. Another estimation says that development cost for a ‘major international telecommunications standard’ may amount to some 1,000 person-years of experience, twenty person-years of actual effort, plus $3 million [OTA 92]. JTC1 alone has been producing between forty and fifty standards per year over the last decade [Gibson 95].

Only fairly recently have economists addressed the problems associated with compatibility standards. In their terms major differences exist between standards in ICT (or, more generally speaking, in fields where networking effects occur, as e.g. railroads), and those valid in the ‘rest of the world’ (i.e. primarily where no networking effects can be observed). For the latter, the assumption of ‘decreasing returns’ holds [Farrell 90], that is, benefits derived from producing something decrease with the number of people producing something similar. For instance, the revenues of the sole producer of washing machines may decrease once other companies start offering similar machines. In contrast, increased returns on adoption may be assumed for ICT; the value of an electronic mail service, for instance, will increase potentially manifold with the number of users with whom communication links can be established. The arrival of competitors offering a compatible service will therefore not necessarily result in lower revenues, it may have the opposite effect and contribute to increasing profits due to the bigger market and the resulting increased value of the first system. Thus, given the increasing returns that stem from the global networks of today, ICT clearly has a major strategic implication - as has the underlying standardisation, upon the outcome of which products will be based. Thus, the choice of a standard will have significant impact on the emergence of new technologies, the performance of single companies, and it may affect competitive advantages of whole economies [OECD 91]. Standardisation may therefore be considered by some as a useful vehicle to bring a company or a country in a more favourable position in the market by trying to push proprietary or national standards at the international level. Yet, with the dramatic increase of players and would-be players in the field of standardisation it remains to seen whether the respective values of the single, and sometimes competing, specifications live up to the expectations.

A company trying to push a proprietary solution towards the status of an international standard is probably the foremost association one has when thinking about the economic dimension of standards and standardisation. Significant increases in market shares - and thus potential gains - may be at stake when a product stands to be ennobled by becoming a standard. At the same time this is the ground upon which turf wars within the committees flourish if competitors try to either push their own ideas, propose a ‘neutral’ solution, or just try to impede the whole process in order to prevent any standard in the field in question. According to [Besen 95] four distinct situations are possible:

- Common interests
  There are no competing proposals, and a decision can quickly be reached by consensus. All parties involved attempt to serve the common good.

- Opposed interests
  Each opponent prefers his own proposal to be adopted, but would prefer no standard at all to the adoption of a competitor's proposal. This situation arises when the gains associated with the winning proposal are comparably big compared to the gains of the industry as a whole.
- Overlapping interests
  Again, each opponent prefers his own proposal to be adopted, but would rather have a competitor's proposal adopted than have no standard at all. This may happen if, conversely to the situation outlined above, the whole industry stands to benefit the most from the adoption of a standard (regardless from where it originated) rather than the original proposer.
- Destructive interest
  At least one player prefers not to have any openly available standard at all, and accordingly tries to slow down the process. This player typically is a major vendor largely dominating the market with a proprietary product who would lose market shares if a standard were in place.

Obviously, these alternatives all come down to the question of competition vs co-operation. The path towards competition may eventually lead to a company's dominating market position with a product or service based on its own proprietary specification. Yet, at the same time the virtual absence of other players may render this particular market insignificant. On the other hand, Cupertino establishes a broader market for products or services based on open specifications, created through, and capable of accommodating, a number of different players. As has for instance been shown in [Swann 90], a product that succeeds in creating an environment in which other vendors consider it beneficial to produce compatible products will prove considerably more successful than its competitors. Such compatible products can only emerge if the underlying original specifications have been made public, or if a very liberal licensing policy has been pursued. This example serves to highlight potential benefits to be gained from open specifications, even if the product itself is inferior to its (less open) rivals in terms of functionality provided. Here, the range of products compatible to the original specification strengthen its status as a de-facto 'standard', which in turn triggers the development of even more compliant products. As a result, a bigger market has been established, leading to increasing revenues.

Another popular (and to some degree valid) perception has it that standards, once established, tend to suppress the development of superior technology. This is true, however, only if 'superior' at the same times means 'incompatible', which is not necessarily, albeit often, the case. Customers waiting for the advent of ATM systems, for instance, severely hampered and, in fact, virtually thwarted the take-off of another, earlier high-speed communication system, DQDB. In the same way, an established standard may not only hinder progress, but may also reduce the variety of alternative technologies. After all, that is what compatibility standards are all about. The resulting limited variety of technological options carries the risk of leaving the market stranded with a less-than-optimal solution, which in turn may eventually yield the need for an expensive move towards a better, but incompatible system. At the same time, however, because of its dominant position in the market the winning standard-based system may trigger follow-up developments. For example, a wealth of different applications were soon available once the PC and DOS had established themselves as (de-facto) standards. This example also illustrates another potential economic effect of standardisation: increased price competition. As functionality or other product characteristics have been eliminated as means of competition, prices become even more important. This development will most likely result in price cuts which will, in turn, push the diffusion of the system. It does not really matter here if the standard in question has evolved through sheer market power as a de-facto standard, originated from an industry consortium, or proceeded through 'official' standardisation processes.

Against this background it is no big surprise that, at least initially, the major players in the ICT field were very reluctant when it came to open standardisation. With large customer bases for their proprietary systems they had little incentive to open up this lucrative market to competitors; IBM in the sixties and seventies being a case in point. Such dominant companies, which control the market, or at least major segments of it, have to lose the most from openly available standards. More recently, however, even major players seem to realise that their products hardly stand a chance of dominating an ever growing and increasingly competitive market. Strategic alliances are formed with producers of complementary products, users, and competitors. Even arch-enemies (e.g. Netscape and Microsoft) have agreed to co-operate in certain areas to enable the development of de-facto 'standards'. Ultimately, the desire to open up markets has in many cases led to the formation of consortia, the major, and maybe only, goal of which being the establishment of open specifications, which may eventually be submitted to one of the 'official' standards setting bodies for formal approval.

The User's Side

Thus far, the discussion has been somewhat focussed on the vendor's views on standardisation. Obviously, things look slightly different from the user's perspective. For them, standards serve three major purposes:

- Avoid technological dead-ends
  Users want to avoid purchasing products that eventually leave them stranded with an incompatible technology. A number of issues need to be considered in this context. For instance, it has to be decided if and when a new technology should be purchased, and which one should be selected. Too early adoptions not only bear the risk of adopting a technology that eventually fails in being successful in the market, but also ignore the considerable time and money that have gone into the old technology. It has to be decided if and when to switch from a well-established technology to a new one. Investments in the old technology need to be balanced with the prospective benefits potentially to be gained from this move. On the other hand, late adopters may lose competitive advantage while being stuck with outdated technology.

- Reduce dependency on vendors
  Being locked in into a vendor-specific environment is increasingly becoming a major risk for a user, despite the advantages that can be associated with integrated proprietary solutions. In particular, problems occur if a vendor misses an emerging development, and its users are forced to switch to completely new (and different) systems; a very costly exercise. Accordingly, standard compliant products from a choice of vendors appeal to the users, who can pursue a pick-and-mix purchasing strategy, and also stand to benefit from price cuts as a result of increased competition.

- Promote universality
  Ultimately, users would like to see seamless interoperability between all hardware and software, both internally (between different departments and sites) and externally (with customers and business partners). With the ongoing globalisation of markets this can only be achieved through international standards. Clearly, this holds especially for communications products. Ideally it should not matter at all which vendor or service provider has been selected; interoperability should always be guaranteed. This implies that user needs and requirements are met by the standards (and the implementations). In addition to seamless communication - and the business value this represents in its own right - there is another major economic benefit to be gained: the costs of incompatibility may be tremendous. For instance, in 1980 half of General Motor's automation budget went into the design of specific interfaces between incompatible machines [Foray 95], a situation that would not have occurred if adequate standards had been available in the first place.

An issue closely related to the above is the timing of standards. A typical complaint about today's 'official' standards setting processes has it that standards emerge too late, that they are overtaken by the technological development, especially in the realm of IT, and that accordingly new ways of producing standards need to be found. Whilst it is certainly true that a standard needs to meet its window of opportunity, it is equally true that a specification done hastily bears the risk of producing a technological lock-in, i.e. to standardise on an inferior specification. This may easily happen since the long-term values of a proposed standard are next to impossible to evaluate, potentially giving advantage to proposals with well understood short-term benefits (see above). Thus, a thorough, though lengthy process may well make sense in the long term.

Funding - or rather the lack of it - is another aspect which is of particular importance to the user community. In fact, it is one of the most prominent explanations for users' abstention from standardisation. Active involvement in standardisation not only demands regular participation in meetings; additional time for preparation is also required. A standard worker will not be available to his/her employer for a considerable length of time if the engagement is taken seriously, thus incurring major expenses. Various suggestions have been made if and how funding should be provided to attract more users. Views differ widely in this respect; Rankine claims that no special funding needs to be made available to users because users already are adequately represented on the committees [Rankine 95], whereas it is argued for instance in [Fischer 90] that additional funding should be made available by interested parties (e.g. governments) to enable and promote participation of smaller users.

Very Briefly – A Consequence

A model of the typical standards development process today, as described in the respective bodies' rules, is shown in Figure 1. It pretty much shows a one-way street. A number of open issues emerge here: for instance, it is not entirely clear who initiates a standardisation activity (apart from the formal rules), and on what grounds. That is, the activity may be either reactive or pro-active, it may or may not be based on user requirements, or it may only be supposed to serve a vendor's purposes. Moreover, until well after the completion of a standards project, it cannot be established whether or not a standard is economically viable. Given the huge amounts of money that have to go into the development of a single standard (see e.g. [Spring 95]) it will be disastrous if a standard fails to deliver.

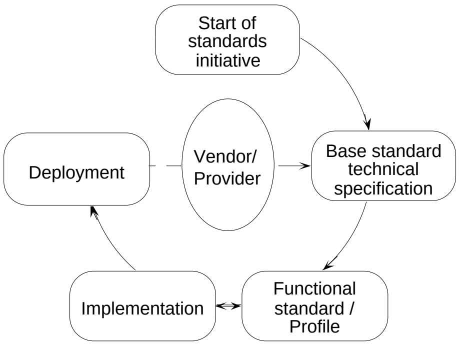

Figure 1: *Development and subsequent deployment of a standard today*

The process exhibits some more potentially severe deficiencies. For one, no dedicated requirements elicitation phase precedes the process. As it currently stands, all requirements are largely made up by committee members, to a considerable degree reflecting vendors' and service providers' interests [???]. Even worse, no formal mechanisms have been established to enable users to feed their working experience directly back into the process. Thus, to actually enable meaningful involvement, modifications to the process described above will be necessary.

Considering the above deliberations, the role users should ideally assume, the comments and insights gathered through the survey, as well as the conclusions outlined above, the model for a specification development process, as depicted in Figure 2, emerges. This model aims primarily at the support of anticipatory standards setting. Yet, in doing so it draws upon an important reactive element, i.e. cycles of user feedback. These cycles allow users to - iteratively - contribute their experiences to the process. Up to now this property has only been available in reactive standardisation, where users had the opportunity to forward their earlier experiences (albeit only informally).

The model describes a two-stage process, with an analysis stage preceding the technical work. During the former, a first compilation and verification of initial requirements from both the technical and the business perspective is performed, the required resources are secured and, if applicable, it is ensured that a window of opportunity will be met (see also [Morell 96]). The model draws upon ideas from the discipline of Participatory Design, which promotes equal participation of all stakeholders in process or system design. In fact, the process will not work properly without equal and balanced participation of all interested parties, and without a common understanding of the problems and issues at hand. This holds for both phases of the model.

Ideally, commitments from all stakeholders, including vendors, service providers and users to implement and to actually use the technology, respectively, would be required at the earliest possible stage. Such commitments would help ensure that requirements will indeed be addressed and that products based on this standard will eventually be purchased. Unfortunately, given the duration of the process, the pace of technological development and the resulting level of insecurity, such strong commitments are highly unlikely. However, a common understanding, together with a certain degree of commitment and trust on the sides of both users and vendors should be feasible.

Several fundamental decisions need to be taken before the actual (technical) standards setting work can commence. First of all, it is crucial to realise the impossibility of solving all potential future problems from the outset, and accordingly not to try and specify an all-embracing standard. Recent experiences show that attempting to specify such standards are bound to fail. For example, most of the largely proactive OSI standards, based on projected user needs, and designed to include all possible options, have not been accepted in the market (see e.g.

[CEU 96], [Wagner 95]). Accordingly, an evolutionary approach has been adopted (very much in line with the ideas underlying the Participatory Design approach. Work is based on a set of initial requirements, specified primarily by those who will actually use the system in the future. Subsequently, the specification can be refined based on experiences made during the deployment phase.

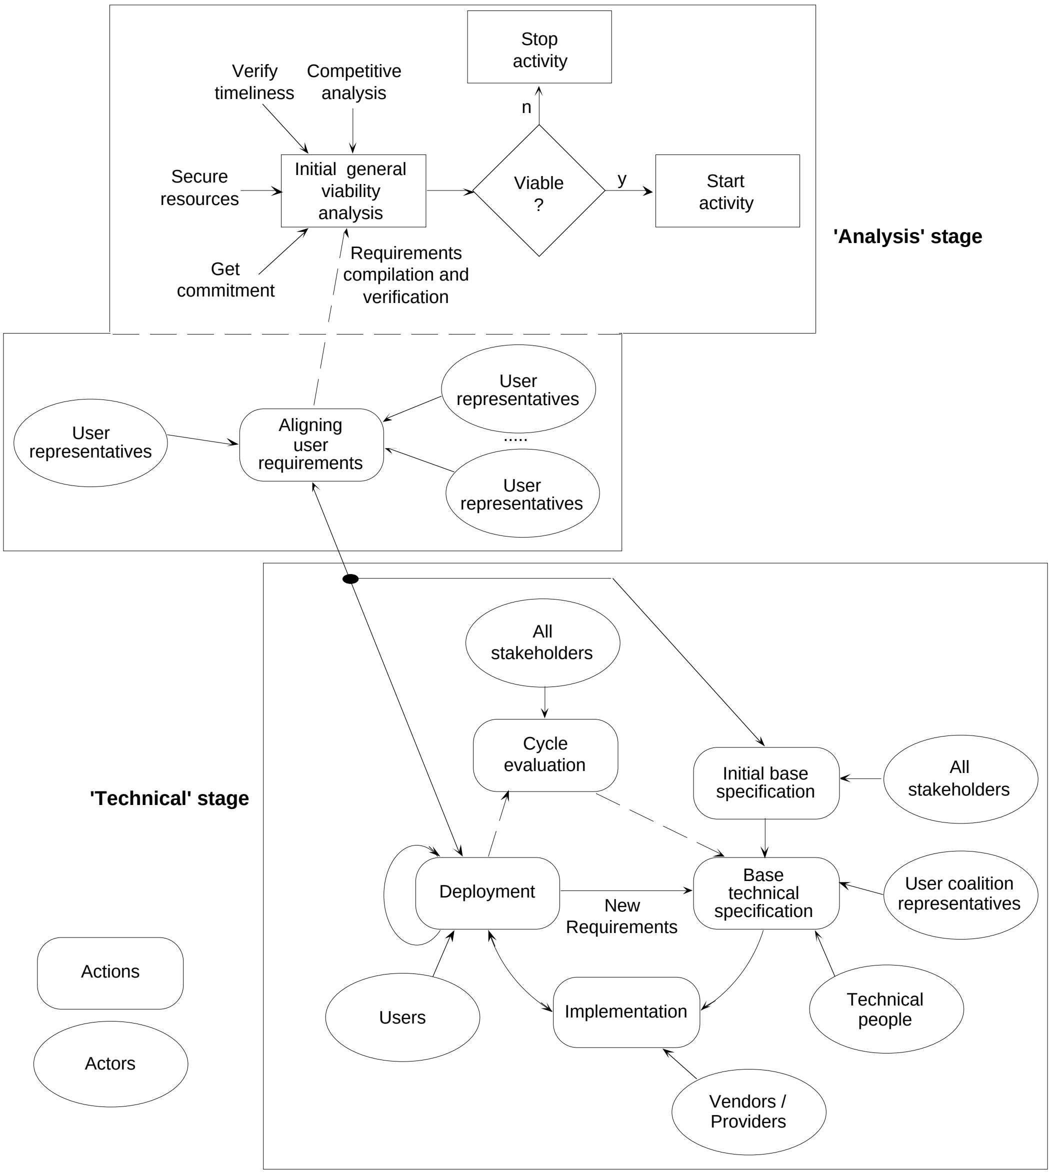

Figure 2: *The Improved Cycle of Specification Development (the Cyclic Stage Model of Standardisation, CSMS)*

However, a survey has shown that the vast majority of corporate users of electronic messaging systems, for example, were not in a position to identify much more than very basic requirements on (standards for) these systems, even after having used comparable services for a lengthy period of time [Jakobs 98]. Obviously, the situation will be even worse at the beginning of a proactive standards setting activity, when very little or no prior experiences at all on the side of the users must be assumed. Thus, the set of initial requirements mentioned above will, in all likelihood, be little else but a comparably sketchy wishlist. Still, despite all its shortcomings, this - inevitably rather pragmatically compiled - list represents the state-of-the-art in user requirements, and will be a far more useful starting point for any standardisation work than - probably largely unfounded - assumptions regarding future 'user needs' which are primarily defined by vendors and service providers anyway, and upon which the standards setting process is typically based today.

Having assembled the initial requirements list, its single items have to be weighed with respect to their perceived importance, potential contradictions have to be sorted out, and finally a catalogue of mandatory requirements has to be agreed upon, to serve as the basis for the subsequent technical work. Furthermore, in the likely absence of well-founded, strong requirements during the early stages of the process, this catalogue will have to be 'living', i.e., requirements will be added and possibly removed as work progresses and practical experiences are gained.

Looking at the implementation of a new standard it needs to be considered if, and to what degree, it will cause changes to an already installed base. If it does, the nature an the extent of these changes must be evaluated carefully. Actor Network Theory (see e.g. [Callon 91]) tells us that a large, well-aligned actor-network - such as e.g. the Internet, or indeed any large, global network - is almost irreversible. It can only be changed into an equally well-aligned network. That is, a step-by-step approach to the introduction of changes is highly recommended, which particularly implies that only very few components of a network may be replaced at a time, and that an alignment phase has to follow each such change before another replacement may be done. Ensuring backward compatibility is a closely related major issue here, as is avoiding installed-base hostility. Yet, as the final outcome of any major changes to a large installed base is not normally predictable, the only realistic way of approaching this problem is to use common sense, to apply a degree of pragmatism, to monitor the transition process closely, and to make sure that changes are indeed carried out gradually.

Above, I have already argued that it is more useful to strive for a reasonably fast first specification based on generic initial requirements, and to enhance it subsequently when real requirements emerge from service use, than to aim at a full-fledged specification from the outset. In parallel, efforts should be undertaken to establish confidence that a viable technology will emerge. Based on the initial requirements compiled, technical committees would then attempt to develop a draft specification, which is returned to the user representatives for review and, eventually, approval. Ideally, the engineers drafting the specifications would come from both sides, vendors and users, as this would help to keep the specifications in line with the requirements available. The group of user representatives should be composed of engineers as well as non-technical people to make sure that all facets of the requirements are met. There may be several iterations, with the proviso that a balance is maintained between evaluation and development. Subsequently, the first version of the final specification can be released for implementation.

During the following deployment phase, operational experiences will be gained within a variety of user environments. Eventually, the experience accumulated will be sufficient to identify shortcomings of the specification. The resulting additional requirements identified will serve as input to a second cycle, during which the specification will be enhanced accordingly. Prior to this stage, the specification will be 'frozen', i.e. no changes may be made. This does not include dealing with 'defect reports', i.e. reports on errors or ambiguities in the original specification, which have to be acted upon immediately.

Some particular characteristics of this process model are worth noting. For instance, it can be applied to industry consortia and official standards setting bodies alike; it does not make any assumptions regarding the process which eventually yields a specification (apart from the crucial role assigned to the user community). This is particularly important since both approaches to standards setting will be needed in the future, largely depending on the type of technology to be standardised. In some areas, such as basic telecommunications infrastructure services (as e.g. X.25, ISDN, and ATM), technical specifications need to be mature and have long-term stability to ensure that potentially large investments in such technology will not become obsolete. Here, formal standards, based on a consensus of all interested parties, will be preferred. On the other hand, in less stable areas (as e.g. PC interfaces and peripherals) agreed specifications will have to be available quickly to establish a common basis for applications, and to avoid competing proprietary specifications. In such cases, where the expected lifespan of a specification may not be too important, less formal processes, as e.g. those typically adopted by industry consortia will be better suited. The model also allows for a consortium to do the initial specification and subsequently have a formal body to transform it into a standard once it has sufficiently matured.

A process based on this model would have considerable advantages over the current one, but some (minor) potential drawbacks must also be conceded. Beginning with the latter, we first note that obviously the 'first round' of the process may take longer than it does today, due to the time required for the viability analysis. Second, the process stands or falls by input from the user community. Consequently, users' attitudes towards standardisation need to be changed. They have to be convinced that active contribution to standardisation is in their interest, not just a waste of scarce resources.

Another issue relates to the need for co-operation between users, primarily during the first stage of the process. It could be argued that companies will be reluctant to publish and discuss their requirements, making them available to competitors, as they might see a danger of revealing too much of their strategic planning. However, successful work done within various consortia, for example, shows that users themselves apparently are not too worried by this issue.

These problems are more than compensated for by the benefits. For one, a viability analysis preceding the technical work will help reduce the number of unsuccessful standardisation activities. As a desirable side effect, this will also free resources, both financial and human, from projects that are not likely to produce any viable output anyway, thus compensating, at least in part, for the additional resources required to develop the viability analysis.

The second major advantage relates to the mechanism provided for user feedback. As pointed out above, the model is based on two assumptions. The first postulates that requirements will emerge over time, and that it may well take them more than a decade to develop (as it did in the case of e-mail). The second one postulates that a standardisation activity is not just a matter of setting up a committee, producing a standard, and disbanding it again (as e.g. the IETF does). Some European research programmes have adopted a similar, although less formal, approach. In these programmes, certain projects serve as testbeds at user sites; experiences gained from these sites are contributed to European standards bodies, typically ETSI. This approach has to some degree been borrowed from the usability domain, where 'learning-by-using' - typically through prototypes - is a popular way of identifying usability deficiencies in a software system [Lindgaard 94]. Reassessments could be done on a regular basis, thus making standards development more reliable, and easing the task of systems planning for the user community. They would ensure the start of new specification activities if and when sufficiently strong new requirements emerge. It follows that the user community must have the right to demand the specification of a new version of a standard.

Summarising the characteristics of the proposed model of a standards setting process, it can be noted that a viability analysis preceding the actual technical work should not only make standardisation more efficient, but should also reduce the number of standards, making life easier for both users and vendors. The feedback and monitor mechanisms for users will significantly contribute to standards that meet actual requirements. The price to pay is primarily constituted by the longer overall process. To compensate for this, the time allocated for the technical specification of a standard should be minimised, to enable timely first implementations.

References

[Besen 95] Besen, F.M.: The standards process in telecommunication and information technology. In: [Hawkins 95].

[Berg 90] Berg, J.L.; Schummy, H. (eds): An Analysis of the Information Technology Standardization Process. North-Holland, 1990.

[Callon 91] Callon, M.: Techno-economic networks and irreversibility. In: Law, J. (ed.): A Sociology of Monsters: Essays on Power, Technology and Domination. Routledge, 1991.

[CEU 96] Commission of the European Union: Standardization and the Global Information Society. http://www.ispo.cec.be/infosoc/legreg/docs/96359.html

[Cargill 95] Cargill, C.F.: A Five-Segment Model for Standardization. In [Kahin 95].

[Cowan 92] Cowan, R.: High Technology and the Economics of Standardization. In: [Dierkes 92].

[David 94] David, P.A.; Monroe, H.K.: Standards Development Strategies Under Incompete Information - Isn't the 'Battle of the Sexes' Really a Revelation Game. MERIT Working Paper 2/94-039, 1994.

[David 95] David, P.A.: Standardization Policies for Network Technologies: The Flux Between Freedom and Order Revisited. In: [Hawkins 95].

[Dierkes 92] Dierkes, M.; Hoffmann, U. (eds): New Technologies at the Outset - Social Forces in the Shaping of Technological Innovations. Campus/Westview, 1992.

[Farrell 90] Farrell, J.: The Economic of Standardization: A Guide for Non-Economists. In [Berg 90].

[Fischer 90] Fischer, P.: The Role of the User in the Standardization Process - Workshop Report. In: [Berg 90].

[Foray 95] Foray, D.: Coalitions and Committees: How Users Get Involved in Information Technology Standardisation. In: [Hawkins 95], 1995.

[Gibson 95] Gibson, R.B.: The Global Standards Process: A Balance of the Old and the New. In: [Kahin 95].

[Hawkins 95] Hawkins, R.W. et al. (eds): Standards, Innovation and Competitiveness. Edward Elgar Publishers, 1995.

[Jakobs 98] Jakobs, K.; Procter, R.; Williams, R.: Telecommunication Standardisation - Do We Really Need the User? Proc. 6th Int. Conf. on Telecommunications, ICT '98, IEE Press, 1998.

[Kahin 95] Kahin, B.; Abbate, J. (eds): Standards Policy for Information Infrastructure. MIT Press, 1995.

[Lindgaard 94] Lindgaard, G.: Usability Testing and System Evaluation. Chapman & Hall, 1994.

[Morell 96] Morell, J.A.; Stewart, S.L.: Standards Development for Information Technology: Best Practices for the United States. ACM StandardView, vol. 4, no. 1, 1996

[Monteiro 98] Monteiro, E.: Scaling Information Infrastructure: The Case of the Next Generation IP in Internet. The Information Society, vol. 14, no. 3, 1998.

[OECD 91] Organization for Economic Co-operation and Development (eds): Information Technology Standards: The Economic Dimension. Information Computer Communications Policy, no. 25, 1991.

[OTA 92] Office of Technology Assesment (eds): Global Standards - Buidling Blocks for the Future. OTA-TCT-512, 1992.

[Rankine 95] Rankine, L.J.: Do we Need a New Standards Systems? In: [Kahin 95].

[Richardson 97] Richardson, G.B.: Economic Analysis, Public Policy and the Software Industry. DRUID Working Paper 97-4, 1997.

[Schmidt 92] Schmidt, S.K.; Werle, R.: The Development of Compatibility Standards in Telecommunications: Conceptual Framework and Theoretical Perspective. In: [Dierkes 92].

[Spring 95] Spring, M.B. et al.: Improving the Standardization Process: Working with Bulldogs an Turtles. In: [Kahi 95].

[Spring 96] Spring, M.B. et al: System Design for an Integrated Document System and Its Impact on Standards Development Efficiency. Proc. IEMC 96, 1996.

[Swann 90] Swann, P.: Standards and the Growth of a Software Network. In: [Berg 90].

[Wagner 95] Wagner, C.S. et al.: Open Systems Standards in Manufacturing: Implications for the National Information Infrastructure. In: [Kahin 95].