

The Fortran (not the foresight) saga: the light and the dark

This paper was published in *Fortran Forum*, Vol 9 No 2, October 1990, pp 23-32.

This PDF version was taken from the HTML version prepared by the author from the text submitted for publication, and appears by permission of the copyright holder. Copyright © 1990 by the Association for Computing Machinery, Inc
http://www.acm.org/pubs/copyright_policy/

The Fortran (not the foresight) saga: the light and the dark

Week by week we read about corporate battles for shares in a market, be it hardware, software or services, about takeover manoeuvrings and buyouts, about rival consortia seeking to dominate areas like Unix or networking. Yet these ongoing stories, hard though they may sometimes be to follow, have the complexity of a Ladybird book, compared with the long-running saga of the Fortran 8X standard.

The story began immediately after the finalisation of the Fortran 77 standard, hence predating the Thatcher era - and it could still outlive it. No longer is it the "8X" story, of course; the US Fortran Standards Committee is now calling it "Fortran 90", but considering it started as "82" and for a period of false dawn was "88", no-one who has survived the whole saga to date would put serious money on that being the final designation. (However, those who remember a novel called "Limbo 90" may regard it as not inappropriate.)

One day someone perhaps will record the history of Fortran 82/8X/88/90/..., but it needs an Escher to depict it or a Borges to describe it. Indeed, Borges, Borgias and Byzantium all come to mind when one looks at the complex, multi-layered interplay of forces which have contributed. The best that can be done here is to indicate the various forces and their interconnections, and briefly describe the issues and forces involved in the main battles that have occurred.

The main groups of forces are the standards committees, corporate bodies, professional institutions, and individuals. These themselves have subgroupings and classifications.

Standards Committees The standards committees, as well as being "forces", in the sense that they have power and hence are players in the game, also provide the battlefields, at least for the set-piece battles. Other battles are fought in public meetings like the Fortran Forums which have been held from time to time, mostly in the US and the UK, to present and discuss issues and developments; in published media like *Computer Weekly*, technical journals, and the newsletter also called *Fortran Forum* (see below); and in recent years through the electronic mail networks and bulletin boards.

Most of the set-piece battles have been fought in, and sometimes between, the US Fortran Standards Committee X3J3, and the international Fortran standards working group, ISO/IEC JTC1/SC22/WG5 ("WG5" for short) or its predecessors. The international committee structure has undergone substantial changes during the life of the saga but not in ways directly relevant; I shall use the current designations, and "WG5" or "WG5-to-be" for the Fortran group.

X3J3 does the technical work of producing the standard (both as a US national and an international standard); WG5 coordinates international comment on the work X3J3 is doing and gives general advice on the direction in which the development of the standard should be heading.

WG5 reports to SC22, the ISO subcommittee responsible for programming languages as a whole. SC22 in turn reports to JTC1, the Joint Technical Committee of the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) responsible for information technology standards generally.

The organisational structure within the USA is rather different, and X3J3 reports to X3, an Accredited Standards Committee of the American National Standards Institute (ANSI) which is the US member body of ISO rather as the British Standards Institution (BSI) is the UK member body. For the purposes of explaining the Fortran saga, X3 is roughly at the same level as JTC1 though not completely equivalent. Hence there is no intermediate level in the USA equivalent to SC22, though there is an SC22 TAG (Technical Advisory Group) whose purpose is to advise X3 on liaison with the ISO body.

The US delegation to WG5 is appointed by (and drawn from) X3J3. Some non-US members of WG5 are also members of X3J3, including some from the UK.

Corporate Bodies Corporate bodies can be divided into suppliers (vendors) and users; in this case it is sufficient to restrict "suppliers" to those who supply Fortran compilers. Corporate users can for some purposes be divided into public sector and private sector, or even commercial, industrial, academic, and government. Both on the supplier and on the user side these corporate bodies vary greatly in size and in attitude.

Professional Institutions

The two professional institutions that have played a supporting role in the saga have been the British Computer Society (BCS) and its US counterpart the Association for Computing Machinery (ACM). Except at one point the ACM's role has been mainly the passive one of distributing information through publishing the newsletter Fortran Forum. The BCS, through its Fortran Specialist Group (FSG) has been more active, receiving reports of X3J3's progress at its meetings, sponsoring public meetings (Fortran Forums) in the UK, and making available, at lower cost than would otherwise have been the case, copies of drafts of the standards for review and comment. Most UK members of WG5 are members of the BCS FSG.

Individuals

At one level, individual people are the least important of the groupings. They are not major economic forces as suppliers or users, they cannot nominate representatives as the other groups can. Nevertheless, the corporate, institutional and standards bodies are composed of individuals, it is individuals that recommend policy, write position papers, argue, discuss, and cast votes. As with most human activities, personalities play an important part.

Even leaving aside affiliations, duties and loyalties, the taxonomy of the various individuals concerned with Fortran standardisation almost deserves a small dissertation of its own. Two classifications need to be noted at the outset. One is between traditionalists and revisionists (otherwise conservatives and radicals, sensibles and irresponsibles, backwoodsmen and progressives, depending on attitude). Most (not all) can be firmly placed on one side or the other of this divide. The other classification is between featurists and generalists - those whose primary interest is in particular features, and those concerned with the language as a whole. The featurists of course have to be subdivided into different interest groups - vector processing, process control and so on. This can include support for specific vendor extensions to Fortran 77.

The point is that the battle lines are redrawn and the groupings rearrange, depending on what issue arises, the way it comes up, and the current context. For example, a featurist who is also a traditionalist may generally support moves to "reduce the size of the language" (a recurrent rallying call throughout the saga) but will change sides if the proposal is to remove a wished-for facility. A featurist who is also a revisionist may, however, be reluctant to support a proposal for a wished-for facility if done in a way which runs counter to general moves to modernise the language.

It is hardly surprising, with such a shifting mix of interests, that it is hard to reach consensus on anything. The gradual shifts in individual membership of committees and organisations over so long a period of time, plus the continual changes in the technology and the marketplace, means that any consensus needs continual refreshment and reaffirmation.

Anyone who has been on any standards committee knows the weary feeling when a recently recruited member brings up an issue or an argument which was debated at length and dealt with years before. Yet people wonder why standards take so long to agree. The wonder is, rather, that they ever get agreed at all. It is a measure of the crucial importance of standards (even if not adequately recognised by everyone) that people nevertheless continue to soldier on.

Let us now turn to some of the highlights of the saga. I am grateful to C.P. Snow and Anthony Powell for providing me, unknowingly, with the titles of the sections and (with John Galsworthy) of the piece as a whole.

Time of hope

In this saga of the light and the dark, in the beginning there was light. Fortran 77, let it not be forgotten (though often it is), was rather a spatchcock affair, agreed because people agreed that they had to produce something even if they could not agree on what it should be. The idea for "Fortran 82" was, OK that's out of the way now let's do it right. The starting point was to be a "core plus modules" architecture able to accommodate featurists and generalists, traditionalists and revisionists.

For a time, things went well. The first ever formal meeting, in London in 1978, of what was to become WG5 (there had been an ad hoc group discussion the previous year at the ISO programming language committee meeting at The Hague) discussed constructively the nature of "core plus modules" and what the modules should be. At the end of 1979 Walt Brainerd of X3J3 (then and still) described the project in this way:

"There has been a significant shift in attitudes concerning the objectives of Fortran standardisation. One reason for this has been the realisation that the language is becoming too large by adding new features and almost never removing any old ones....

"X3J3 has responded to these challenges by proposing a language architecture consisting of a 'core' language plus 'modules'. The core is to be a complete general purpose programming language, but will not contain features that are redundant or obsolete just because they are in Fortran 77....

"One of these modules will contain all of the features in Fortran 77 that are not in the core. This is important in order to support the huge investment in Fortran programs written using these features." [1]

The beauty of core-plus-modules was that it could contain the battles in very small areas. Its weakness was that it crucially depended on achieving consensus on two issues; what should be in the core, and what should be in what became known as the "obsolete features" module - the parts of Fortran 77 not in the core. This developed into a battle between revisionists/generalists, and traditionalists allied to revisionist/featurists who, while not objecting to core-plus-modules or to the concept of obsolescence, could not accept that cherished features should be excluded from the core - or, worse, deemed "obsolete". After a protracted engagement, almost exclusively within X3J3, core-plus-modules was dropped, thus opening up the whole battlefield to anyone. Fortran "82" had by this time become Fortran "8X".

In an update on his earlier discussion, in February 1983, Walt Brainerd was writing:

"The committee seemed to have rejected the idea of a small elegant core. It appears to be the old problem that a committee simply cannot design something small and elegant; one person's frill is another's essential feature, and for some reason a committee almost always seems to vote to include things in this category rather than exclude them. After everything else is done, it still may be possible to identify a small 'core'." [2]
He was to be disappointed.

The acceptance world

During this period the international group had observed, commented, occasionally expressed concern, but generally acquiesced in what was happening in X3J3. Both Fortran 77 and its predecessor, by now known as Fortran 66, had been ANSI standards endorsed by ISO. The only groups organised enough to provide substantial international comment on the Fortran 77 draft had been the BCS FSG and a group in The Netherlands. Even the change of status in 1983 to a formal ISO working group seemed at first to make little change.

The conscience of the rich

However, at the Geneva meeting of WG5-to-be in the spring of 1984, an issue came up which with the benefit of hindsight can be seen as showing straws in the wind. Having now been at work for more than six years, but still with no draft standard in sight, X3J3 felt that public feedback on what had been done so far would be useful, and proposed the publication by X3 of a "Fortran Information Bulletin". This proposal was discussed also at the ISO WG and it was agreed that a FIB ("an unfortunate acronym", I wrote in *Computer Weekly* at the time) be published internationally as well, with identical text.

When this came to X3 for approval, however, the vote was 40 in favour and 2 against - but the two were DEC and IBM. DEC objected unless changes were made which (I wrote at the time) "would ... have undermined all the work that has been done on Fortran 8X". IBM did not mind publication, but only if its objections and those of others were published in the FIB as part of the document.

The first major set-piece battle had begun; battlefield X3, two major vendors against the majority of X3J3.

X3 asked X3J3 if it would accept such conditions. Since the purpose of the FIB was simply to present the current state of development on 8X, both X3J3 and WG5-to-be said that this was unacceptable as a matter of principle. The FIB was a factual document, not propaganda, and those who read it should be free to make up their own minds.

These views were reported back to X3, who decided (despite the large majority) not to approve publication. All standards committees have a duty to try to reach consensus. If agreement cannot be reached, however, a vote of the size of 40 to 2 is usually reckoned enough. Clearly, on X3 some votes were more equal than others.

However, the battle was not quite over. For the first and so far only time in the saga, ACM played an active role, though as a natural extension of its usual role of making information available. A special joint issue of *Fortec Forum* (now *Fortran Forum*) and *SIGNUM Newsletter* (SIGNUM is the ACM Special Interest Group in Numerical Mathematics) published a technical report - which happened to contain all of the same words as the FIB (apart from the front cover) but not including the DEC changes or the IBM objections. Association of Computing Marines to the rescue?

The "straws in the wind"? The light was beginning to give way to the dark; two major vendors had shown that they were willing to go to considerable lengths to block progress in the direction the majority wanted to go; and by backing up X3J3, the ISO working group dimly began to see that it had a bit more power of its own than merely power to comment - if it chose to use it.

The sleep of reason

The ACM report appeared, comments came in, X3J3 continued its work. By the Bonn meeting of WG5-to-be in the summer of 1985, it was clear that the time was at last approaching when a draft standard could be produced, and the meeting passed a resolution commending X3J3's work as "broadly meeting the needs of the Fortran community worldwide".

However, if the shape and size of the emerging 8X was becoming clear, it was also apparent that moves were afoot to mobilise forces to oppose the general trend. The main, overlapping groups were of vendors, led by IBM and DEC but now joined by others (though by no means all); and traditionalists, both user companies and individuals. The vendor group mainly objected to the size of the language, the traditionalist grouping to some new features and also to the concept of "deprecated" features (by now the weakened term for "obsolete").

After some internal skirmishing (since all the various factions were represented), the Bonn meeting added to the general endorsement of the current 8X a statement to the effect that any major "downward revision" would be "detrimental to the interests of Fortran users". This was by way of a warning shot, across the bows of the vendor group in particular. It was also, though in itself no more than reinforcement of approval of what X3J3 had done, by far the strongest statement made by any ISO group up to that time of the direction in which it would like the work to proceed in future.

For some months, progress continued, until a draft was produced, which went for internal ballot to the members of the US and ISO groups for approval for publication as a draft for public review. In both groups, there were people who felt that too many minor deficiencies remained to make it publishable at that stage, but the ISO vote was fairly favourable. In the US, however, there was a powerful group of US objectors, complaining that the language was too large and should be cut down, and brought more into line with "existing practice" (which, being translated, meant what the vendors already supplied).

It was at this point that the event occurred which escalated the disputes over Fortran 8X into a "shooting war", and crucially altered the relationship between X3J3 and WG5. The first consequence was clear at the time; this is not hindsight. The second took longer to become evident, but may eventually be seen of longer term importance.

The normal response to failing to achieve an adequate level of support in such circumstances is to review the objections and seek to resolve them if possible in ways

which will not alienate support among those in favour of the present proposals. However, the X3J3 response was (as more than one observer said at the time) "to go into panic mode". Despite the ISO group resolution at Bonn, wholesale cuts were made, and the result was dubbed "the compromise plan". Large sections of the draft standard were transferred to an appendix of "suggested extensions" or deleted altogether. The structure that had been so painstakingly built over the last few years had, virtually overnight, shrunk to something much more limited in scope, and now lacking facilities into which many had invested long hours in designing and promoting, and in convincing others of their technical merit. Many inconsistencies were revealed when the consequences of this surgery (or butchery, as some would have it) were examined.

In retrospect, it is hard to understand why the architects of the compromise could not see that the inevitable consequence would be not to speed agreement on the eventual standard, but to delay it.

The valley of bones

Outside X3J3 (and one gathers to some extent within it) the initial reaction to the "compromise plan" was of stunned disbelief, followed in many cases by outrage. Revisionists, of course, who on the whole did not support the view that 8X was getting too large, were horrified at the scale of the cuts. Featurists whose own pet facilities had fallen under the axe were also annoyed. However, it became clear that the main source of objections with any organisational structure for expressing it was WG5 (which the ISO group now finally was).

The source of the resentment was not just the cuts themselves, though probably a majority on WG5 were of revisionist tendency. In addition, many of those who could have accepted some or even most of the cuts resented the way that it had been done, without consultation with the international community. Even without the Bonn resolutions of the previous year, this would have been the case. The Bonn statement warning against substantial downward revision now assumed major importance, being not just a readymade stick, but one given extra weight because X3J3 had chosen to ignore it.

During the summer of 1986, the architects of the "compromise plan" continued to clean up the draft 8X in the reduced form; but it became clear that too many regarded the plan as a capitulation rather than a compromise, for it ever to achieve consensus, the only worthwhile aim of compromise. Wherever compromise lay, if it lay anywhere, it was not here.

At the Halifax (Nova Scotia) meeting of WG5 in the late summer of 1986, one or two of the US delegation attempted to argue that the "compromise" was not really a substantial downward revision; but most X3J3 members seemed ready to admit both that it was, and that WG5 should have been consulted.

The revisionists - their ranks enhanced by aggrieved featurists, assured of substantial support from an international community increasingly assertive of its rights - licked their wounds, regrouped, and started the long haul back. Month by month, lost ground was gradually regained - not all, but enough perhaps to achieve the elusive consensus.

Books do furnish a room

The Halifax meeting of WG5 had passed, by an adequate but not overwhelming majority, that reductions on the scale proposed were tolerable only if this served to expedite the appearance of the standard. Even then, it was fairly clear that this was a forlorn hope. The "compromise" had reduced consensus, not increased it. However, there was much more agreement that things could not go on like this, and that the best hope of progress would be to get a document out for public comment, hoping that the resulting feedback would resolve at least some of the issues.

Even that was hard enough to achieve. Through the rest of 1986 and 1987 successive versions of the draft were produced, and X3J3 members and observers got used, every few weeks, to another bulky package arriving to add to the ever-growing pile. Features were reviewed, some reinstated, some modified to look more like "existing practice" (i.e. what vendors already supplied) in an attempt to achieve genuine compromise. By the Liverpool WG5 meeting in August 1987 the document had reached the stage of being submitted to X3 and SC22 for distribution as a draft proposed standard.

By this time, enough remedial and reconstruction work had been done on the erstwhile "compromise" plan that this again in X3 attracted negative votes from vendor interests. However, WG5 still had a short shopping list of required features. One related to data handling at the bit level, a typical minority featurist requirement which had fallen victim to the "compromise", and believe it or not had similarly failed to gain acceptance over a decade earlier for Fortran 77. Another was free form source code, another casualty but by now almost the touchstone of revisionism.

The new men

While the revisionists had been hard at work retrieving lost ground, the traditionalists had not been idle. Electronic mail and meeting reports carry evidence of attempts to drum up opposition to 8X, in the wake of the failure of the "compromise plan" to achieve the desired effect. Scare stories circulated about 8X not being compatible with Fortran 77, despite that being an explicit design aim.

One prime source of scare stories, for example, was the concept of "deprecating" storage association. Another tale going round was that free source form would invalidate "millions of lines" of existing code. That these scares had no foundation did not matter: as the great Dr Goebbels taught us, such stories do gain credence among people who do not bother to check the facts. If one believes that the end justifies the means, the traditionalists responsible for spreading them were able at the end of 1989 to find their justification.

Though hard evidence is much more difficult to come by, there is also reason to suspect that people alarmed by these stories were encouraged to join X3J3 to defend their interests. (X3 rules make it relatively easy to do this and soon obtain full voting rights, provided the company is prepared to pay a small annual fee, currently \$250). But for whatever reason, during the latter part of the 1980s X3J3 grew significantly in size, and most of the new members were of the traditionalist tendency - a factor of importance given X3's requirement of a two thirds majority - and also mostly representing vendors. These new members, as well as not so much on top of technical issues as those who had been working on 8X for a decade or more - understandable and excusable in itself - were also generally unaware of or indifferent to the international dimension, also a factor to prove of importance. X3J3 had always consisted both of members (of all camps, it must be said) who were internationally-minded and who generally attended WG5, and of members who were not. The new men affected both the revisionist/traditionalist balance and the internationalist/nationalist balance.

The international result of the review of the first public draft was that there were several negative votes from member countries of SC22, mostly because of deficiencies, some of which had already been identified by WG5. Individual comments ran into many hundreds, both in USA and outside. The majority of these were negative, on a simple "vote" count, a fact immediately seized on by some of the traditionalist camp as showing that "X3J3 had got it wrong". In fact the negative "votes" consisted of featurist ones ("my pet feature has been left out"), revisionist ones ("the opportunity has been missed to fix this deficiency") and technical ones ("this proposal won't work because ..."), as well as traditionalist ones ("the language is too large/ too revolutionary/ not existing practice/ ..."). Of the traditionalist replies, many were of the blanket form "language too large" or the like, without supporting arguments, and some were virtually or actually exact copies of suggested letters of protest distributed by some of the traditionalists.

Things did not improve when X3J3 got down to considering the public review comments and how to respond to them. Divisions seemed to be deepening, and positions becoming entrenched. At one point, votes were held on five major issues, which between them would have resolved most of the international negative votes. Each one separately won the necessary majority, but when the package as a whole was put forward, it was voted down.

Students of voting methods will not be surprised at this - different coalitions were involved each time. It is an instance of the Condorcet "paradox of voting", which in its simplest form allows, in straight fights, A to beat B, B to beat C, and C to beat A. The consequence, however, was dismay and disarray, and individuals and groups went away to try to come up with complete package proposals.

As the summer of 1988 wore on, numerous plans were advanced and considered, until four remained. Two were traditionalist, and one mildly revisionist. The fourth, produced by five of the most experienced members of X3J3, was the one that came closest to meeting the international comment. However, at its last meeting before WG5 in Paris in September 1988, X3J3 voted down that plan too. At this point, X3J3 itself was close to

reaching an internal solution, in the sense of an adequate majority though by no means a consensus. The price of such a solution was failure to address many of the international concerns. While it is possible that many members of X3J3 had overlooked that aspect in the manoeuvrings to build a majority, at the end of that meeting it was clear that this was a conscious decision of the committee. It even explicitly voted that the fourth plan referred to above was not to be submitted to WG5 for consideration, despite protests that X3J3 did not have the right to set WG5's agenda.

Temporary kings

So WG5 met in Paris in September 1988, presented with three alternative plans to meet the comments from the first public review of 8X, none of which came close to meeting the points of concern in the international negative votes. Members of WG5 had previously complained on occasion that X3J3 paid no more attention to the considered national position of an SC22 member body (often, as in the UK, after extensive public consultation) than to a letter from Harvey P. Digweed from Backwater, Iowa. It rapidly became plain that WG5 was not in the mood to play that game any more.

First, WG5 decided to consider the fourth plan which X3J3 had voted should not be submitted to WG5; then it had presentations of all four, after which it decided to waste no more time on two (the two "traditionalist" plans). The architects of the remaining two plans were asked to get together to produce a combined plan to resolve the international negative votes.

While WG5 had in the past expressed views on technical issues, it had never before embarked on specifically technical work, previously regarded as X3J3's prerogative. However, the mood seemed to be, "if X3J3 can't resolve the negative votes, it is up to us to do it".

By the end of an intense week, WG5 had produced the framework of a new 8X standard that was likely to achieve the support of all member countries (except possibly USA); had laid down a timetable for future progress to a second public review; and had declared that the revised Fortran 8X would be called "Fortran 88". This last was mainly a symbolic assertion that "enough is enough", though the sequence of Fortran 66, 77, 88 was appealing in itself.

More than two years had passed since the 1986 "compromise plan", and this time it was the turn of the traditionalists to retire from the battlefield shell-shocked. Recovery was not helped by SC22 broadly endorsing what WG5 had done, and X3 issuing a directive to X3J3 (there was only one dissenting vote) that it should work towards a single national and international standard - "that is, one document". This last phrase was to take on an added significance a year later.

The Boston meeting of X3J3 in November 1988 was a virtually complete capitulation to these pressures, and work towards a second public draft proceeded mainly if not wholly on the lines laid down by WG5. The mood was apparently more one of anti-

traditionalism than of positive internationalism, though the effect was the same. On the other hand, some explicit anti-internationalist sentiments had started being openly expressed by a small minority of X3J3 members. The high tide of WG5 pre-eminence can now be seen as the change of sides of IBM, who early in 1989 announced that it was prepared to drop its opposition to Fortran 8X in the interests of achieving general agreement. The WG5 timetable slipped, but that is so familiar a phenomenon in standards making, controversial or not, that it was hardly remarked. The July meeting of WG5 in northern Italy was a low-key affair with the resolution most critical of X3J3 merely urging them to get on and finish the job.

Corridors of power

It was the lull before the next storm, as the Fortran saga took yet another twist. The first hint of what was to come appeared the week after the WG5 meeting. It must have been known the previous week, and some thought it significant that it had not been mentioned to WG5.

X3 rules allow occasional meetings of its committees outside the US mainland, and more than once X3J3 had held such meetings immediately following WG5. This time it was Vienna, and there it was learned that a vendor representative on X3J3, who was also a member of X3, had put forward a proposal to X3's Standards Policy and Requirements Committee (SPARC) that the USA change the terms of reference of the domestic ANSI 8X project from being a revision of Fortran 77 to a collateral standard - "extended" Fortran - while retaining Fortran 77.

Now since 1977 the international policy, reaffirmed and strengthened in 1985, has been that the aim of programming language standardisation should be a single worldwide standard for each language. This move was widely seen - by many in X3J3 as well as WG5 - to run counter to at least the spirit of that international agreement, to which the USA had subscribed, and possibly even to the letter of it. Whether it was the letter as well as the spirit, depended as so often on interpretation. If ISO had one standard but the USA two, was this a breach of the agreement in SC22 to seek "a single worldwide standard" - or, as X3 had put it the previous year, "one document"? Before the standard was finally agreed, there would be a difference in status of the 8X project - a revised standard in ISO, a collateral standard in ANSI. Would that constitute a breach of the agreement?

All this at least suggested that SPARC should recommend consultation with SC22 and WG5 over the proposed change of status - which in ISO would also need the approval of JTC1 were it put forward there.

However, SPARC chose not to see it that way. With only one dissident, the committee voted to recommend the change of status, which was therefore forwarded to X3 for endorsement. Though SPARC is technically only an advisory committee, it is very much X3's "inner cabinet" and its recommendations are rarely overturned. Apparently convinced by the argument that a substantial body of US users (and some outside) did not

want Fortran 8X, X3 endorsed the proposal - coincidentally, the ballot closing just before the Berlin meeting of SC22 in September 1989.

It was undoubtedly true that there was a body of users who simply wanted to preserve their existing investment in Fortran code, but saw their future needs elsewhere - C, Ada or whatever. The traditionalist/vendor group had been assiduous in finding such users and stimulating them to object to 8X - if only by copying a "form letter". During the second public review, one went so far as to urge people "you don't have to buy a copy of the draft standard to object, you don't even have to read it ..." This prompted a scathing response on the electronic network from an X3J3 member - a vendor representative, incidentally - remarking on the standards of professionalism and codes of conduct which should be observed in the standards world.

However, none of these needs meant that Fortran 77 had to be retained as a standard - the long continuance of support for Cobol 74 and even Cobol 68 after the appearance of Cobol 85 was proof of that. While the requirement remained, vendors continuing to support Fortran 77 would find a market. Nor was retention of Fortran 77 the only option. The possibility of specifying within the standard a subset of Fortran 8X of more limited power had been raised from time to time earlier, and though the general view in X3J3 and WG5 was that subsets were undesirable (a long road had been travelled since core-plus-modules), nevertheless WG5 at Paris 1988 had agreed that they could be considered if all else failed to achieve consensus.

There are strong grounds for believing that the true motivation for the move - which, though purporting to be on behalf of users, had been promoted by a group of vendors - was that US Federal procurements are normally based on conformance to standards, where standards exist. The clear implication was that vendors would aim to get procurements to specify "Fortran", but leaving them with the choice of which of the two Fortrans to supply. The conformance rules of Fortran 77 are so weak that the vendors could then supply Fortran 77 implementations with vendor-specific extensions which were incompatible with Fortran 8X. The net result would be to wreck all post-Fortran 77 standardisation.

At Berlin, SC22 called on those concerned to work to try to reconcile the conflicting positions on the status of 8X and the retention of Fortran 77. The US delegation, however, stuck to the "party line" that this was a purely domestic US matter, and that if 8X was the same in ISO and ANSI (though differently named and of different status) then the US standards body had fulfilled its international obligations. In strict legal terms, the retention of Fortran 77 was indeed a domestic matter; USA or any other country had the right so to do. In the real world of international interdependence, the continuing existence of Fortran 77 as a US standard had more than domestic US implications.

These manoeuvrings aroused more resentment even than the ill-fated "compromise plan" of 1986. One factor was that this should happen when final agreement had seemed at last in sight. Another was that the 1985 "compromise" had at least been devised inside the Fortran community, and dealt with there - whereas now, a member of the community,

unable to persuade his colleagues on X3J3 and WG5, had sought to get them overruled by higher (and less well-informed) authority, and regardless of the international consequences.

After SC22 in Berlin there was still some hope that X3 might revise its views - there had been enough dissent to the retention of Fortran 77 to require a "reconsideration" vote, and in view of expressed international concerns it was hoped that there would be second thoughts, at least to give time for negotiations.

It was not to be. When X3 met to reconsider, the attitude to consulting X3J3 seemed to be a mixture of "we know what X3J3 thinks" and "we're in charge". The international dimension was dismissed by repeating that it was a purely domestic matter, disregarding clear evidence of international concern and alarm. When asked who believed that Fortran 77 was a subset of Fortran 8X, only four hands went up. One of the four, a long-standing X3J3 member, tried to explain, but someone else said "everything else we've heard suggests that Fortran 77 is not a proper subset of Fortran 8X" and that was the end of the discussion. X3 reaffirmed its decision.

Ironically, a third X3J3 and X3 member who was present later concluded a lengthy reply to the torrent of electronic complaints with "all X3 did was say that Fortran 77 would remain a proper subset of Fortran 8X and in the US only" (my italics). Whatever one makes of this inconsistency, what it was that X3 thought it was doing is beside the point. What matters is that, while Fortran 77 is a subset of Fortran 8X, vendor implementations of Fortran 77 are in general very much not subset implementations of Fortran 8X, nor is there much evidence that they have any intention of making them so.

Annoyance with X3, and the X3J3 members who had engineered this coup, was considerable in the Fortran community, even among X3J3 members outside the revisionist camp. The complete reversal of attitude by X3 from the previous year, amounting virtually to a 180-degree U-turn, has not as yet been satisfactorily explained.

Hearing secret harmonies

The atmosphere of the special additional meeting of WG5 held in London in March 1990 to review the situation was one of determination rather than recrimination. As at Paris eighteen months earlier, it rolled up its collective sleeves and went to work on the draft, which was now dubbed "Fortran 90". Rather than just reviewing it and asking X3J3 to make changes, it made changes - in effect saying to X3J3 "here's what the ISO Fortran 90 standard will be, please check it for accuracy" leaving X3J3 to bring the ANSI Extended Fortran standard into line. To be sure, it was not put quite as starkly as that, and some issues left unresolved at the end of the week were referred to X3J3 to deal with. However, it was quite clear that WG5 (which, remember, contains a good many X3J3 members, who at this meeting took the lead in the technical work) would take a very dim view if X3J3 did not basically accept what had been done.

As this is being written (June 1990) X3J3 has indeed completed a final draft and this is in the last stages of being prepared for publication from ISO headquarters in Geneva as a Draft International Standard (DIS) - after which the technical content is effectively frozen with changes limited to correction of typographical and other errors. Fortran 90 does look to be on its way - though long term survivors of this saga will still be crossing their fingers.

So, are the problems over? Hardly; there are too many unresolved issues:

- Will the US Fortran 90 (ANS Extended Fortran) be identical to ISO Fortran 90 apart from designation? [It is believed and certainly hoped so, but in this area one has learned that nothing should be taken for granted.]
- Will the traditionalist camp, with this coup under their belt, now give Fortran 90 a clear run? [It looks like it so far, but we thought before the coup that the troubles were over.]
- What are the traditionalist/vendor intentions with regard to the future of Fortran 77? [There have been assurances, but also nasty rumours, and a lot of unease.]
- What will be the impact of this unilateral decoupling of national and international standards, and difference in status of Fortran 90 between ANSI and ISO, outside the Fortran community, and could there be adverse implications for other standardisation projects? [At the least, it could lead to yet another review of the national vs. international standards issue.]

Harmonisation and commonality worldwide is the aim of international standardisation. While the differences in status of and attitude towards Fortran 90, and for that matter Fortran 77, remain between the United States and the rest of the world, then that aim will not properly have been fulfilled. The reverberations will continue until that anomaly has been resolved - for example, by revising the US Fortran 77 standard, not in its content but in its conformance rules, so that no extensions to Fortran 77 are allowed unless compatible with the Fortran 90 standard. That is a poor substitute for what ought to have happened, but better than nothing.

However, there is little evidence of any will to do that, and good grounds for fearing that any hint of revising Fortran 77 will be seized upon by some interests to undermine rather than support Fortran 90. While this unhappy state of affairs continues, if there are any harmonies to be heard, they are indeed secret ones.

Last things

This account of the Fortran saga has been written from the point of view of a UK member of WG5 and of an unreconstructed and unashamed revisionist. That no doubt has been evident! Since the author is no Borges, the saga is much simplified but it is hoped that the main events have been covered even if not everyone will agree with the interpretations and views expressed - which are the author's own and do not necessarily reflect those of the UK at WG5 and SC22 level. David Muxworthy, Lawrie Schonfelder and Loren Meissner provided valuable comments and corrections of factual errors on earlier

versions, which are gratefully acknowledged, but they should in no way be held responsible for the views expressed nor for any further errors, which are the author's own. An article based on material drawn from an earlier version appeared in *Computer Weekly* on 3 May 1990.

References

[1] W.S Brainerd, in chapter 2 (*Fortran*) of *Programming language standardisation*, ed. I.D. Hill and B.L. Meek, Ellis Horwood 1980, p. 34

[2] W.S Brainerd, *The programming language standards scene II: Fortran, Computers and Standards* Vol 2 Nos 2-3, 1983, p. 76

Copyright © 1990 by the Association for Computing Machinery, Inc. Please click here for full [ACM copyright notice](#).

Note added in 1995: despite the pessimism shown in the *Hearing secret harmonies* passage, Fortran 90 was indeed published in 1991 as ISO/IEC 1539:1991, the US standard is the same, and moves are in progress to withdraw the parallel US standard for Fortran 77.

HTML version prepared by Brian Meek, 31 July 1995, updated 22 October 1996.
PDF version edited and created by Jane Sleightholme, jane@fortranplus.co.uk, 27th March 2008.